

13.3.5 MBMS Interfaces

For MBMS in LTE, two new control plane interfaces have been defined (M3 and M2), as well as one new user plane interface (M1), as shown in Figure 13.2.

13.3.5.1 M3 Interface (MCE – MME)

An Application Part (M3AP) is defined for this interface between the MME and MCE. The M3AP primarily specifies procedures for starting, stopping and updating MBMS sessions. Upon start or modification of an MBMS session, the MME provides the details of the MBMS bearer while the MCE verifies if the MBMS service (or modification of it) can be supported. Point-to-point signalling transport is applied, using the Stream Control Transmission Protocol (SCTP) [4].

13.3.5.2 M2 Interface (MCE – eNodeB)

Like the M3 interface, an application part (M2AP) is defined for this interface between the MCE and eNodeB, again primarily specifying procedures for starting, stopping and updating MBMS sessions. Upon start or modification of an MBMS session, the MCE provides the details of the radio resource configuration that all participating eNodeBs shall apply. In particular, the MCE provides the updated control information to be broadcast by the eNodeBs. SCTP is again used for the signalling transport.

13.3.5.3 M1 Interface (MBMS GW – eNodeB)

Similarly to the S1 and X2 interfaces (see Sections 2.5 and 2.6 respectively), the GTP-U⁴ protocol over UDP⁵ over IP is used to transport MBMS data streams over the M1 interface. IP multicast is used for point-to-multipoint delivery.

13.4 MBMS Single Frequency Network Transmission

One of the targets for MBMS in LTE is to support a cell edge spectral efficiency in an urban or suburban environment of 1 bps/Hz – equivalent to the support of at least 16 mobile TV channels at around 300 kbps per channel in a 5 MHz carrier. This is only achievable by exploiting the special features of the LTE OFDM⁶ air interface to transmit multicast or broadcast data as a multicell transmission over a synchronized ‘single frequency network’: this is known as *Multimedia Broadcast Single Frequency Network (MBSFN)* operation.

13.4.1 Physical Layer Aspects

In MBSFN operation, MBMS data is transmitted simultaneously over the air from multiple tightly time-synchronized cells. A UE receiver will therefore observe multiple versions of the signal with different delays due to the multicell transmission. Provided that the transmissions from the multiple cells are sufficiently tightly synchronized for each to arrive at the UE within

⁴GPRS Tunnelling Protocol – User Plane [5].

⁵User Datagram Protocol.

⁶Orthogonal Frequency Division Multiplexing.

the Cyclic Prefix (CP) at the start of the symbol, there will be no Inter-Symbol Interference (ISI). In effect, this makes the MBSFN transmission appear to a UE as a transmission from a single large cell, and the UE receiver may treat the multicell transmissions in the same way as multipath components of a single-cell transmission without incurring any additional complexity. This is illustrated in Figure 13.3. The UE does not even need to know how many cells are transmitting the signal.

The method of achieving the required tight synchronization between MBSFN transmissions from different eNodeBs is not defined in the LTE specifications; this is left to the implementation of the eNodeBs. Typical implementations are likely to use satellite-based solutions (e.g. GPS⁷) or possibly synchronized backhaul protocols (e.g. [6]).

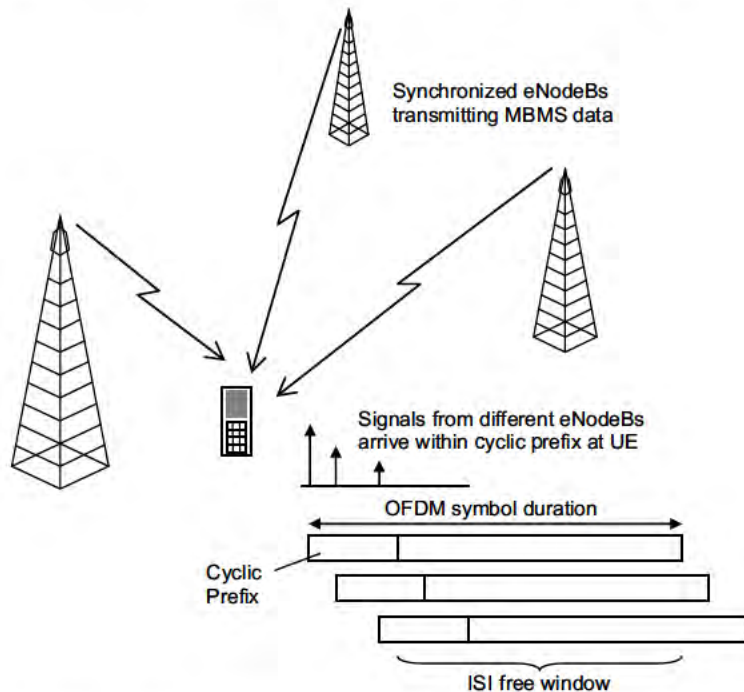


Figure 13.3: ISI-free operation with MBSFN transmission.

This single frequency network reception leads to significant improvements in spectral efficiency compared to UMTS Release 6 MBMS, as the MBSFN transmission greatly enhances the Signal-to-Interference-plus-Noise Ratio (SINR). This is especially true at the cell edge, where transmissions which would otherwise have constituted inter-cell interference are translated into useful signal energy – hence the received signal power is increased at the same time as the interference power being largely removed.

An example of the improvement in performance achievable using MBSFN transmission compared to single-cell point-to-multipoint transmission is shown in Figure 13.4. In this

⁷Global Positioning System.

example, the probability of a randomly located UE not being in outage (defined as MBMS packet loss rate < 1%) is plotted against spectral efficiency of the MBMS data transmissions (a measure of MBMS data rate in a given bandwidth). A hexagonal cell-layout is assumed, with the MBSFN area comprising one, two or three rings around a central cell for which the performance is evaluated. It can be seen that the achievable data rates increase markedly as the size of the MBSFN area is increased and hence the surrounding inter-cell interference is reduced. A 1 km cell radius is assumed, with 46 dBm eNodeB transmission power, 15 m eNodeB antenna height and 2 GHz carrier frequency.




Figure 13.4: Reduction in total downlink resource usage achievable using MBSFN transmission [7]. Reproduced by permission of © 2007 Motorola.

MBSFN data transmission takes place via the Multicast CHannel (MCH) transport channel, which is mapped to the Physical Multicast CHannel (PMCH) introduced in Section 9.2.3.

The basic structure of the Physical Multicast Channel (PMCH) is very similar to the Physical Downlink Shared Channel (PDSCH). However, as the channel in MBSFN operation is in effect a composite channel from multiple cells, it is necessary for the UE to perform a separate channel estimate for MBSFN reception from that performed for reception of data from a single cell. Therefore, in order to avoid the need to mix normal Reference Signals (RSs) and RSs for MBSFN in the same subframe, frequency-division multiplexing of the PMCH and PDSCH is not permitted within a given subframe; instead, time-division multiplexing of unicast and MBSFN data is used – certain subframes are specifically designated as MBSFN subframes, and it is in these subframes that the PMCH would be

transmitted.⁸ Certain subframes are not allowed to be used for MBSFN transmission: in a Frequency Division Duplex (FDD) system, subframes 0, 4, 5 and 9 in each 10 ms radio frame are reserved for unicast transmission in order to avoid disrupting the synchronization signals or paging occasions, and to ensure that sufficient cell-specific RSs are available for decoding the broadcast system information; in a Time Division Duplex (TDD) system, subframes 0, 1, 2, 5 and 6 cannot be MBSFN subframes.

The key differences from PDSCH in respect of the PMCH are as follows:

- The dynamic control signalling (PDCCH and PHICH⁹ – see Section 9.3) cannot occupy more than two OFDM symbols in an MBSFN subframe. The scheduling of MBSFN data on the PMCH is carried out by higher-layer signalling, so the PDCCH is used only for uplink resource grants and not for the PMCH.
- The PMCH always uses the first redundancy version (see Section 10.3.2.4) and does not use Hybrid ARQ.
- The extended CP is used (~17 μ s instead of ~5 μ s) (see Section 5.4.1). As the differences in propagation delay from multiple cells will typically be considerably greater than the delay spread in a single cell, the longer CP helps to ensure that the signals remain within the CP at the UE receivers, thereby reducing the likelihood of ISI. This avoids introducing the complexity of an equalizer in the UE receiver, at the expense of a small loss in peak data rate due to the additional overhead of the longer CP. Note, however, that if the non-MBSFN subframes use the normal CP, then the normal CP is used in the OFDM symbols used for the control signalling at the start of each MBSFN subframe. This results in there being some spare time samples whose usage is unspecified between the end of the last control signalling symbol and the first PMCH symbol, the PMCH remaining aligned with the end of the subframe; the eNodeB may transmit an undefined signal (e.g. a cyclic extension) during this time or it may switch off its transmitter – the UE cannot assume anything about the transmitted signal during these samples.
- The Reference Signal (RS) pattern embedded in the PMCH (designated ‘antenna port 4’) is different from non-MBSFN data transmission, as shown in Figure 13.5. The RSs are spaced more closely in the frequency domain than for non-MBSFN transmission, reducing the separation to every other subcarrier instead of every sixth subcarrier. This improves the accuracy of the channel estimate that can be achieved for the longer delay spreads. The channel estimate obtained by the UE from the MBSFN RS is a composite channel estimate, representing the composite channel from the set of cells transmitting the MBSFN data. (Note, however, that the cell-specific RS pattern embedded in the OFDM symbols carrying control signalling at the start of each MBSFN subframe remains the same as in the non-MBSFN subframes.)

The latter two features are designed so that a UE making measurements of a neighbouring cell does not need to know in advance the allocation of MBSFN and non-MBSFN subframes. The UE can take advantage of the fact that the first two OFDM symbols in all subframes use the same CP duration and RS pattern.

⁸LTE does not currently support dedicated MBMS carriers in which all subframes would be used for MBSFN transmission.

⁹Physical Downlink Control Channel and Physical HARQ Indicator Channel.




Figure 13.5: MBSFN RS pattern for 15 kHz subcarrier spacing.
Reproduced by permission of © 3GPP.

In addition to these enhancements for MBSFN transmission, a second OFDM parameterization is provided in the LTE physical layer specifications, designed for downlink-only multicast/broadcast transmissions. However, this parameterization is not supported in current releases of LTE. As discussed in Section 5.4.1, this parameterization has an even longer CP, double the length of the extended CP, resulting in approximately 33 μ s. This is designed to cater in the future for deployments with very large differences in propagation delay between the signals from different cells (e.g. 10 km). This would be most likely to occur for deployments at low carrier frequencies and large inter-site distances. In order to avoid further increasing the overhead arising from the CP in this case, the number of subcarriers per unit bandwidth is also doubled, giving a subcarrier spacing of 7.5 kHz. The cost of this is an increase in Inter-Carrier Interference (ICI), especially in high-mobility scenarios with a large Doppler spread. There is therefore a trade-off between support for wide-area coverage and support for high mobile velocities. It should be noted, however, that the maximum Doppler shift is lower at the low carrier frequencies that would be likely to be used for a deployment with a 7.5 kHz subcarrier spacing. The absolute frequency spacing of the reference symbols for the 7.5 kHz parameterization is the same as for the 15 kHz subcarrier spacing MBSFN pattern, resulting in a reference symbols on every fourth subcarrier.

13.4.2 MBSFN Areas

A geographical area of the network where MBMS can be transmitted is called an *MBMS Service Area*.

A geographical area where all eNodeBs can be synchronized and can perform MBSFN transmissions is called an *MBSFN Synchronization Area*.

The area within an MBSFN Synchronization Area, covered by a group of cells participating in an MBSFN transmission, is called an *MBSFN Area*. An MBSFN Synchronization Area may support several MBSFN Areas. Moreover, a cell within an MBSFN Synchronization Area may form part of multiple MBSFN Areas, each characterized by different transmitted content and a different set of participating cells.

Figure 13.6 illustrates an example of the usage of MBSFN areas in a network providing two nationwide MBMS services (N1 and N2) as well as two regional MBMS services, R1 and R2 (i.e. with different regional content), across three different regions.

The most natural way to deploy this would be to use four different MBSFN areas, as shown in the figure. It should be noted that the control information related to a service is transmitted by the same set of cells that transmit the data – e.g. MBSFN area N provides both control and data for the two nationwide services. This kind of approach implies that within a geographical area the resources allocated to MBMS need to be semi-statically divided between the MBSFN areas, and UEs that are potentially interested in receiving both nationwide and regional services need to monitor the control information of multiple MBSFN areas. Indeed, UEs are expected to be capable of monitoring control information of multiple MBSFN areas, although simultaneous reception of multiple services is optional regardless of whether they are part of the same MBSFN area. Further details of UE capabilities can be found in Section 13.5.2.




Figure 13.6: Example scenario with overlapping MBSFN areas.

Figure 13.7 illustrates how the same services can be provided without overlapping MBSFN areas. It should be noted that for the nationwide services this approach results in reduced MBSFN transmission gains and potential service interruptions upon change of MBSFN area.




Figure 13.7: Example scenario without overlapping MBSFN areas.

Current releases of LTE do not support dynamic change of the MBSFN area; the counting procedure introduced in Release 10 (see Section 13.6.5) is used only to decide whether or not to use MBSFN transmission for an MBMS service but not to change set of cells used for an MBSFN transmission.

13.5 MBMS Characteristics

13.5.1 Mobility Support

Mobility for MBMS is purely based on the procedures defined for unicast operation. No additional mechanisms are provided to direct UEs to the carrier frequency providing MBMS, nor to reduce MBMS service interruption when moving from one MBSFN area to another.

E-UTRAN can assign the highest cell reselection priority to the carrier frequency providing MBMS. However, in connected mode, when E-UTRAN controls the mobility, the network is aware neither of whether the UE supports MBMS nor of whether it is interested in actually receiving any particular MBMS service.

13.5.2 UE Capabilities and Service Prioritization

All UEs need to support some MBMS-related behaviour, including those that do not support MBMS reception and UEs that conform to the initial Release 8 version of LTE (which does not include MBMS). These UEs must be aware of the MBSFN subframes that are configured, since the RS pattern is different. Release 8/9 UEs not supporting MBMS may also benefit from knowing that no regular downlink resource assignments will occur for them in the MBSFN subframes.

As unicast and MBMS transmissions are time-division multiplexed, parallel reception of unicast and MBMS on the same carrier should be relatively straightforward. Nevertheless, a UE may of course have other limitations preventing parallel support, for example due to memory or display restrictions. Consequently, there is no requirement for a UE that supports MBMS to support simultaneous reception of unicast, nor to support simultaneous reception of multiple MBMS services.

In the case of a UE that cannot receive all the services in which it is interested, it should be able to receive the service (unicast or MBMS) that it prioritizes most. This implies that a UE that is receiving a service should at least be able to detect the start of other services (unless the currently received service is always considered as having the highest priority).

The UE itself performs the prioritization of MBMS services for reception. The E-UTRAN is neither aware of which MBMS services a UE is interested in receiving (a UE may, for example, have successfully received an earlier transmission of a repeated session), nor about the priority ascribed to reception of each service by the UE. The details of prioritization are not specified, being left instead to UE implementation. A UE may, for example, choose to terminate a unicast service at the application layer if it prevents reception of an MBMS service that it has prioritized.

13.6 Radio Access Protocol Architecture and Signalling

13.6.1 Protocol Architecture

The LTE radio protocols have been extended in Release 9 to accommodate MBMS:

- Two new types of logical channels are introduced in Release 9: the Multicast Control Channel (MCCH) and the Multicast Traffic Channel (MTCH) for MBMS control and user data respectively.
- A new transport channel, the Multicast Channel (MCH), is introduced, mapped onto the PMCH to support the multicell MBSFN transmission mode.

User data for an MBMS session may employ one or more radio bearers, each of which corresponds to an MTCH logical channel. The Medium Access Control (MAC) layer may multiplex information from multiple logical channels (MCCH and one or more MTCH) together in one transport block that is carried on the MCH.

The MBMS control information is mainly carried by the MCCH, of which there is one per MBSFN area. The MCCH primarily carries the MBSFNAreaConfiguration message, which provides the radio bearer and (P)MCH configuration details, as well as the list of ongoing MBMS sessions. The remainder of the MBMS control information, consisting of the MCCH configuration and parameters related to the notification of changes in the MCCH information, clearly cannot be carried by the MCCH itself and is instead carried by the Broadcast Control Channel (BCCH), in SIB13 (see Section 3.2.2). In Release 10, a further MCCH message was introduced for the counting procedure, as explained in Section 13.6.5.

Figure 13.8 illustrates the radio protocol extensions introduced to accommodate MBMS.




Figure 13.8: LTE protocol extensions to support MBMS.

13.6.2 Session Start Signalling

As a result of the relative simplicity of MBMS in LTE Release 9, the main signalling procedures are for ‘session start’ and ‘session stop’. The absence of multicast support means that procedures for subscription and joining/leaving are not required; furthermore, service announcement is assumed to be an application-specific procedure, not part of the main MBMS signalling procedures.

Figure 13.9 provides an overview of the session start message sequence.

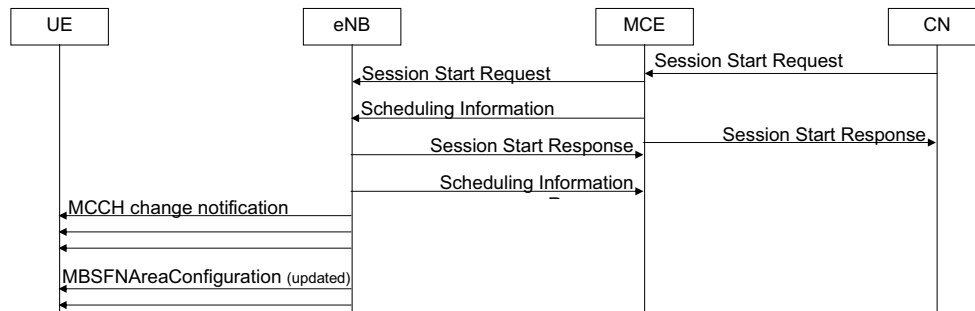


Figure 13.9: The MBMS session start message sequence.

The BM-SC may initiate multiple subsessions with different content, distinguished by flow identifiers. The ‘Session Start Request’ message from the BM-SC includes the following parameters:

- **QoS:** Used for the admission control and radio resource configuration in the MCE.
- **Service area:** Used to decide to which eNodeBs the session start signalling should be forwarded.
- **Session duration:** Used by admission control in the MCE.
- **Time to data transfer:** Used by the MCE to ensure all UEs receive the control information in time.
- **Access indicator:** Indicates in which Radio Access Technologies (RATs) the service will be broadcast.

Upon receiving the session start request, the MCE performs admission control, allocates IP multicast addresses, and allocates and configures the radio resources. The MCE provides the relevant session parameters (e.g. the multicast address) to the eNodeB by means of the ‘Session Start Request’ message. Likewise, the MCE transfers the radio resource configuration within a ‘scheduling information’ message to the eNodeB. Upon receiving these messages, the eNodeB notifies the UEs about a change of MCCH information and subsequently provides the updated MBMS radio resource configuration information within the MBSFNAreaConfiguration message. The eNodeB also joins the transport network IP multicast address.

The actual data transfer starts after a (configurable) duration, such as 10 s.

13.6.3 Radio Resource Control (RRC) Signalling Aspects

13.6.3.1 Scheduling

MBMS does not use the PDCCH for dynamic scheduling. In an MBSFN subframe, the MCH uses all the resources in the frequency domain, so MCH-related scheduling only relates to the subframe allocation in the time domain.

The subframes that carry MCCH are indicated semi-statically by SIB13. For MTCH, the scheduling of which subframes are used for a particular MBMS bearer is somewhat more dynamic: once per ‘MCH Scheduling Period’ (MSP), E-UTRAN indicates which subframes are used for each MTCH. This ‘MCH Scheduling Information’ (MSI) is provided independently for each MCH by means of a MAC control element (see Section 4.4.2.7).

One design goal for LTE MBMS was to avoid long service interruptions when the user switches from one service to another (i.e. channel switching). Assuming that the UE stores the complete MCCH information, the channel switching time is mainly determined by the length of the MSP; if E-UTRAN configures a reasonably low value for this parameter (e.g. 320 ms), users can smoothly switch channels without noticeable delays.

13.6.3.2 MBMS Control Information Validity and Change Notification

As mentioned above, the MBMS control information is mainly carried by the MCCH; the BCCH merely carries the control information needed to acquire the MCCH and to detect MCCH information changes.

Transmission of MBMS control information on the BCCH (i.e. in SIB13) is performed according to the usual procedures for SIBs (see Section 3.2.2.2). The transfer procedures for MBMS control information on the MCCH are similar: control information can only change at specific radio frames with System Frame Number (SFN) given by $\text{SFN mod } m = 0$, where m is the modification period. The MBSFNAreaConfiguration message is repeated a configurable number of times within the modification period.

When an MBMS session is started, E-UTRAN notifies the UEs about an MCCH information change. This change notification is provided a configurable number of times per modification period, by means of a message on the PDCCH using Downlink Control Information (DCI) Format 1C (see Section 9.3.5.1), using a special identifier, the MBMS Radio Network Temporary Identifier (M-RNTI). This message indicates which of the configured MCCHs will change. UEs that are interested in receiving an MBMS session that has not yet started should try receiving the change notification a minimum number of times during the modification period. If the UE knows which MBSFN area(s) will be used to provide the MBMS session(s) it is interested in receiving, it has to meet this requirement only for the corresponding MCCH(s).

Figure 13.10 illustrates a change of MBMS control information, affecting both the BCCH and the MCCH. In this example, E-UTRAN starts to use the updated configuration at an SFN at which the BCCH and MCCH modification period boundaries coincide. The figure illustrates the change notifications provided by E-UTRAN in such a case. Initially, the change notification indicates only that MCCH-2, which has a modification period of 10.24 s, will change. Later, the notification indicates that both MCCH-1 (which has a modification period of 5.12 s) and MCCH-2 will change.




Figure 13.10: An example of a change of MBMS control information.

The change notification procedure is not used to inform UEs about modifications of ongoing sessions. UEs that are receiving an MBMS session are therefore required to acquire, every modification period, the MBMS control information of the MBSFN area used for that session.

13.6.3.3 Radio Resource Allocation

The network broadcasts in SIB13 a bitmap indicating to the UEs the set of subframes that is allocated to the MCCH. This bitmap covers a single radio frame. It is possible to configure a more robust Modulation and Coding Scheme (MCS) for the subframes that include MCCH.

Semi-static resource allocation signalling indicates which subframes are allocated to a particular MBSFN area; this is known as the Common Subframe Allocation (CSA).¹⁰ The CSA is defined by up to eight patterns. Each pattern specifies the allocated subframes by a bitmap, that covers one or four radio frames, and repeats with a periodicity of up to 32 radio frames. The MCCH subframes are included in the CSA.

Next, the subframes assigned to the MBSFN area are allocated to each of the MCHs using the MBSFN area. Within a CSA period, the MCHs are time-multiplexed, with each MCH using one or more subframes which are consecutive within the set of subframes in the CSA; as the subframes are consecutive in this way, only the last allocated subframe is signalled for each MCH. This is known as the MCH Subframe Allocation (MSA).

The MBMS resource allocation is illustrated in Figure 13.11 for an example with two MBSFN areas:

- In even-numbered radio frames, subframes 3 and 6 are allocated, while in odd-numbered radio frames only subframe 6 is allocated;
- Within a period of 80 ms, the MBSFN areas are not interleaved – MBSFN area 1 takes the first 4.5 radio frames while area 2 takes the next 3.5;
- Both MBSFN areas use two MCHs, one with an MSP of 160 ms and one with an MSP of 320 ms;

¹⁰The term ‘common’ here indicates that the CSA subframes are shared by all the MCHs of the MBSFN area.

- Both MBSFN areas use a CSA period of 160 ms;
- Two subframes are allocated to MCCH.

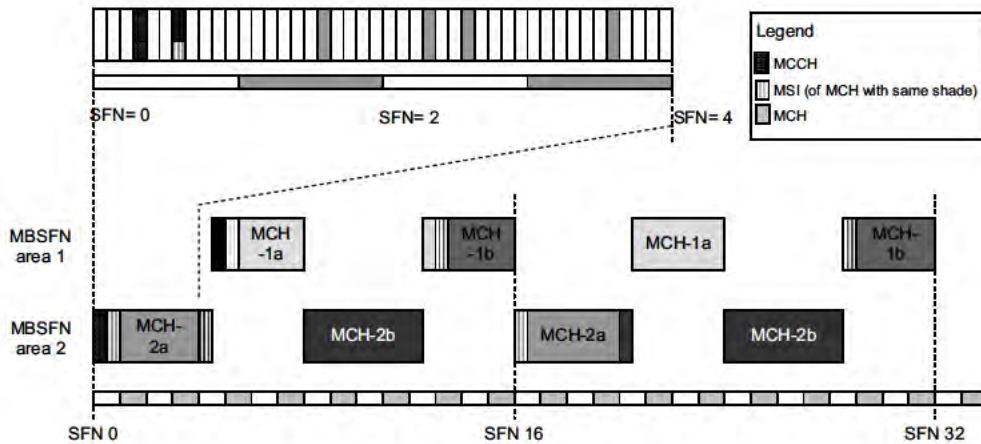


Figure 13.11: An example of MBMS resource allocation.

Every MSP, a UE that is receiving, for example, MCH-2b via MBSFN area 2 acquires the initial subframes to acquire MCCH (if present in this MSP), the MSI MAC control element (see Section 13.6.3.1) and subsequently the subframes dynamically allocated to the service in question.

13.6.4 Content Synchronization

As explained in Section 13.4, the basic property of MBSFN transmission is that all participating cells transmit exactly the same content in a synchronized manner so it appears as one transmission to the UE. Mechanisms are therefore provided to ensure synchronization of the MBMS content – i.e. to ensure that all participating eNodeBs include the same MBMS control information and data in the corresponding time-synchronized subframes.

For the MBMS control information, whenever the MCE updates the control information it indicates the modification period from which the updated control information applies by means of a parameter called ‘MCCH update time’.

For the synchronization of MBMS user data, a specific protocol was designed: the SYNC protocol as specified in [8]. Each MBMS bearer operates a separate instance of the SYNC protocol, which uses a timestamp to indicate the time at which the eNodeB should start the transmission of a first packet belonging to a synchronization sequence. The timestamp is based on a common time reference and should be set taking into account transfer delays that may occur. All packets that the BM-SC allocates to a synchronization sequence are given the same timestamp. The MSP is the same as, or a multiple of, the synchronization sequence values applicable for the services carried by the MCH in question.

Before considering the details of the SYNC protocol, it is important to understand the overall architecture, which can be summarized as follows:

- The BM-SC performs traffic shaping: it discards packets as necessary to ensure that, for each synchronization sequence, the actual bit rate of each MBMS bearer does not exceed the guaranteed bit rate (see Section 2.4). The BM-SC also decides in which synchronization sequence each packet should be transmitted.¹¹
- The MCE semi-statically configures which services are multiplexed together, as well as the radio resources allocated to the relevant MCH. The MCE also semi-statically configures the order in which the services are scheduled.
- The participating eNodeBs first buffer all the packets of an entire MSP. The eNodeBs then dynamically schedule the services, taking into account the order pre-defined by the MCE. If the BM-SC allocates more packets to an MSP than actually fit into the allocated radio resources, the eNodeBs discard some of the packets, starting with the last packets of the service scheduled last according to the pre-defined order (so-called tail-dropping). Finally the eNodeBs compile the MSI. The resulting MAC control element is transmitted at the start of the MSP.

Figure 13.12 illustrates the SYNC protocol by means of a typical sequence of the different SYNC Protocol Data Unit (PDU) types.




Figure 13.12: An example of a SYNC protocol PDU sequence.

¹¹In a future release of LTE, it may become possible for the BM-SC to perform 'radio-aware traffic shaping' – i.e. knowing which services share a given radio resource and their respective bit rates, the BM-SC may handle temporary rate fluctuations by moving some packets to a previous or later scheduling MSP.

The BM-SC transfers the data allocated to one synchronization sequence by means of a number of PDUs of ‘type 1’. These PDUs all include the same timestamp, a packet number that is reset at the start of the synchronization sequence and increases for every subsequent packet, and an elapsed octet counter that indicates the number of octets transferred since the start of the synchronization sequence.

After transferring all the user data, the BM-SC may provide additional information by means of a PDU ‘type 0’ or ‘type 3’. These PDUs may be repeated to increase the likelihood the PDU is received correctly. PDU type 0 not only includes the regular timestamp, the packet number and the elapsed octet counter but also indicates the total number of packets and the total elapsed octets within the synchronization sequence. PDU type 3 provides the same information, but in addition provides the length of each packet in the synchronization sequence.

The eNodeB has a buffer corresponding to one MSP, in which it inserts the received packets in accordance with the pre-defined service order. If one or more non-consecutive packets are lost, the SYNC protocol provides the eNodeB with all the information required to continue filling the transmission buffer correctly. In such cases, the eNodeB only needs to mute the subframes affected by the lost packets (rather than muting all subsequent subframes until the end of the MSP). The eNodeB may be unable to do this if consecutive packets are lost (unless a PDU type 3 is received), as the size of the layer 2 headers depends on the precise construction of the transport blocks, which depends on the packet sizes.

13.6.5 Counting Procedure

A counting procedure was introduced in Release 10 that enables E-UTRAN to determine how many UEs are receiving, or are interested in receiving, an MBMS service via an MBSFN Radio Bearer (MRB) (i.e. using the MBSFN mode of operation). Compared to UMTS, the counting procedure is limited in the following respects:

- Only UEs of Release 10 or later that are in connected mode respond to a counting request from the network;
- The counting request does not include the session identity, with the result that UEs respond to a counting request even if they have already received the upcoming session (in case of session repetition).

Moreover, E-UTRAN can only interpret the counting responses properly when it knows whether or not the service is available via unicast, because if the service is available, all interested UEs should be in connected mode to receive the service, whereas if the service is not available, only a fraction of the interested UEs may be in connected mode.

Based on the counting results, E-UTRAN may decide to start or stop MBSFN transmission of a given MBMS service. Dynamic change of the MBSFN area (i.e. dynamic control of the cells that participate in the MBSFN transmission of a service) is not supported.

Figure 13.13 provides an overview of a typical message sequence for MBMS counting. The sequence shown is based on the following assumptions:

- Whenever a session of an MBMS service is scheduled to take place (e.g. according to the Electronic Programme Guide (EPG)), it is accessible via unicast transmission;
- Upon start of the session’s schedule, the BM-SC allocates a Temporary Mobile Group Identifier (TMGI) and provides a session start indication to E-UTRAN;

- E-UTRAN waits a certain time (e.g. 10 s) before initiating the counting procedure; this delay allows UEs that are interested in receiving the service to enter connected mode to receive the service via unicast.




Figure 13.13: MBMS counting upon session start.

The counting request messages relate to a single MBSFN area and can cover up to 16 services. E-UTRAN can issue one ‘MBMSCountingRequest’ message per MBSFN area in a modification period, and UEs only respond once in a modification period. In the event of E-UTRAN repeating the same counting request (i.e. including the same services) in a subsequent modification period, the UEs respond again.

Counting may also be performed to deactivate the service in case, at a later point in time, the number of interested UEs has dropped. In such a case, the MBSFN transmission of the service may be suspended temporarily; if later the number of interested UEs increases again, the MBSFN transmission of the service may be resumed. This is illustrated in Figure 13.14.

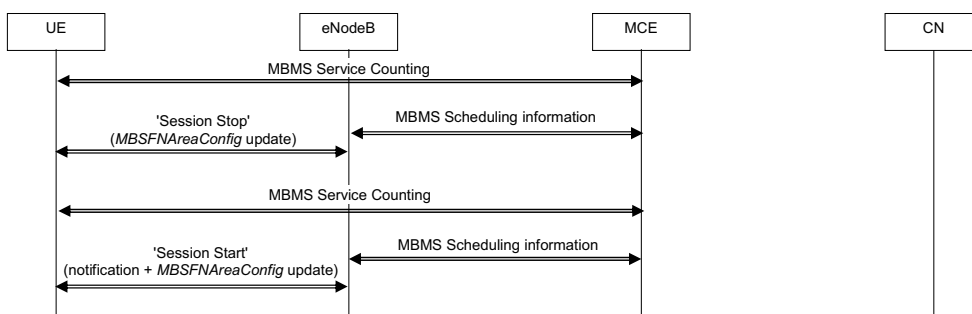


Figure 13.14: Suspension and resumption of MBSFN transmission based on counting.

It should be noted that normally an eNodeB does not initiate the MBMS change notification procedure without a session start; MBMS service resumption is an exception.

13.7 Public Warning Systems

Cellular networks are commonly used to broadcast public warnings and emergency information during events such as natural disasters (e.g. earthquakes, tsunamis or floods). There are two Public Warning System (PWS) variants: the Earthquake and Tsunami Warning System (ETWS), predominantly used in Japan, and the Commercial Mobile Alert Service (CMAS), mainly used in the Americas.

As MBMS was not part of the first release of LTE, it was agreed to support public warnings by means of the BCCH.

ETWS provides ‘primary notifications’, which are short notifications delivered within 4 seconds, as well as ‘secondary notifications’ that are less time-critical and provide more detailed information such as where to get help. Both the primary and secondary notifications, of which there may be multiple, are optional. In case of an emergency, E-UTRAN does not need to delay the scheduling of ETWS notifications until the next BCCH modification period. It notifies UEs (in both idle and connected modes) about the presence of an ETWS notification by means of a paging message including a field named ‘etws-Indication’. Upon detecting this indicator, an ETWS-capable UE attempts to receive SIB10, which is used to transfer the ETWS primary notification. The UE also checks if SIB11 is scheduled, which is used to carry an ETWS secondary notification (or a segment of one). Upon receiving the primary and/or secondary ETWS notifications (depending on which is scheduled), the UE may stop acquisition of system information for ETWS.

CMAS only employs short text messages, for a single type of warning message. In an emergency situation, multiple CMAS alerts may need to be broadcast. However, E-UTRAN only schedules a subsequent CMAS alert after completely transmitting all segments of the previous CMAS warning message. In the same way as for ETWS, E-UTRAN need not delay the scheduling of CMAS warning messages until the next BCCH modification period, and paging is used for notification. Unlike ETWS, however, SIB12 is used to transfer a CMAS warning message (or a segment of one), and the UE should continue SIB12 acquisition upon receiving a complete CMAS warning message.

Both ETWS and CMAS incorporate security mechanisms, including means to verify the integrity of the warning message and to authenticate the message source. The use of these mechanisms is subject to regulatory requirements.

13.8 Comparison of Mobile Broadcast Modes

MBMS is not the only technology which is capable of efficient delivery of multimedia content to mobile consumers, and it is therefore instructive to consider briefly some of the relative strengths and weaknesses of the different general approaches.

13.8.1 Delivery by Cellular Networks

The key feature of using a cellular network for the content delivery is that the services can be deployed using a network operator’s existing network infrastructure. Furthermore, it does not necessarily require the allocation or acquisition of additional spectrum beyond that to which the operator already has access. However, the advantage of reusing the cellular network and

spectrum for multimedia broadcasting is also its main drawback: it reduces the bandwidth available for other mobile services.

It can be noted that mobile operators have been providing mobile television over UTRAN networks for a number of years already; this has however generally been based on point-to-point unicast connections; this is feasible in lightly loaded networks, but, as explained in Section 13.2, it does not make the most efficient use of the radio resources. MBMS transmission from a single source entity to multiple recipients on a single delivery channel allows radio resources to be shared, thus solving the capacity issues associated with the use of unicast transmissions for streaming of bandwidth-hungry services such as mobile television.

In general, the use of a cellular network for MBMS also has the advantage of bringing the capability to send ‘personalized’ content to groups of a few users, as well as the possibility of a built-in application-layer return channel, which is useful for the provision of interactive services. However, such functionality is not supported in current releases of LTE.

13.8.2 Delivery by Broadcast Networks

Mobile broadcast services may alternatively be provided by standalone broadcast systems such as Digital Video Broadcasting-Terrestrial (DVB-T) and Digital Audio Broadcasting (DAB). Originally developed for fixed receivers, specific mobile versions of these standards have been developed, namely Digital Video Broadcasting-Handheld (DVB-H) and Digital Mobile Broadcasting (DMB) respectively.

These systems typically assume the use of a relatively small number of higher-powered transmitters designed to cover a wide geographical area. For broadcast service provision, user capacity constraints do not require the use of small cells.

Such systems can therefore achieve relatively high data rates with excellent wide-area coverage. However, they clearly also require dedicated spectrum in which to operate, the construction and operation of a network of transmitters additional to that provided for the cellular network, and the presence of additional hardware in the mobile terminals.

Standalone broadcast systems are also, by definition, ‘downlink-only’: they do not provide a direct return channel for interactive services, although it may be possible to use the cellular network for this purpose.

13.8.3 Services and Applications

Both cellular delivery (MBMS) and standalone broadcast methods enable the delivery of a variety of multimedia services and applications, the most fashionable being mobile television. This gives network operators the opportunity to differentiate their service offerings.

In some ways, cellular delivery (MBMS) and standalone broadcast (e.g. DVB-H) systems can be seen as complementary. The mobile television industry and market are changing rapidly, and may be better described as part of the mobile entertainment industry, where new market models prevail. An example of this is the concept of the ‘long tail of content’ [9] which describes the consumption of products or services across a large population, as illustrated in Figure 13.15: given a large availability of choice, a few very popular products or services will dominate the market for the majority of consumers; a large number of other products or services will only find a niche market.

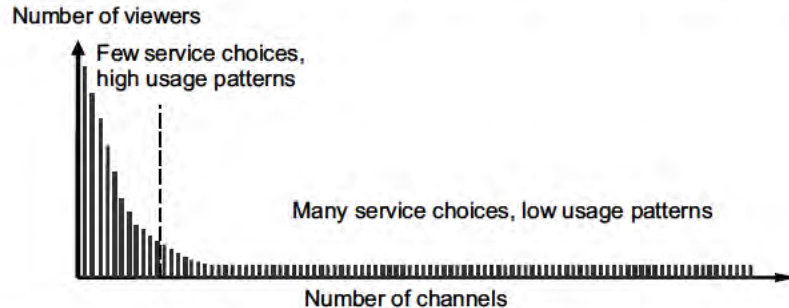


Figure 13.15: The 'Long Tail of Content'.

This is a very appropriate representation of the usage of broadcast television, whether fixed or mobile: viewing figures for a few main channels (typically national broadcasters) are high, while usage is low for a profusion of specialist channels that gather a limited yet dedicated audience.

The deployment of mobile broadcast technology worldwide (whether by cellular MBMS or standalone broadcast) is very much dependent on the availability of spectrum and on the offer of services in each country, while infrastructure costs and content agreements between network operators and broadcasters are also significant factors.

What matters ultimately to the user is the quality of reception on their terminal, and mobility. Deployment of MBMS in LTE networks can be a highly efficient solution, especially by exploiting the use of Single Frequency Network operation based on OFDM – which is also the transmission method used by DVB-H and DMB.

References¹²

- [1] 3GPP Technical Specification 23.246, 'Multimedia Broadcast/Multicast Service (MBMS); Architecture and Functional Description (Release 6)', www.3gpp.org.
- [2] 3GPP Technical Specification 23.228, 'IP Multimedia Subsystem (IMS); Stage 2 (Release 8)', www.3gpp.org.
- [3] Internet Engineering Task Force, 'RFC3261 Session Initiation Protocol', www.ietf.org.
- [4] Internet Engineering Task Force, 'RFC4960 Stream Control Transmission Protocol', www.ietf.org.
- [5] 3GPP Technical Specification 29.060, 'General Packet Radio Service (GPRS); GPRS Tunnelling Protocol (GTP) across the Gn and Gp interface (Release 8)', www.3gpp.org.
- [6] IEEE1588, 'IEEE 1588 Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems', <http://ieee1588.nist.gov>.
- [7] Motorola, 'R1-070051: Performance of MBMS Transmission Configurations', www.3gpp.org, 3GPP TSG RAN WG1, meeting 47bis, Sorrento, Italy, January 2007.
- [8] 3GPP Technical Specification 25.446, 'MBMS synchronisation protocol (SYNC)', www.3gpp.org.
- [9] D. Anderson, *The Long Tail: Why the Future of Business is Selling Less of More*. New York, USA: Hyperion, 2006.

¹²All web sites confirmed 1st March 2011.

Part III

Physical Layer for Uplink

14

Uplink Physical Layer Design

Robert Love and Vijay Nangia

14.1 Introduction

While many of the requirements for the design of the LTE uplink physical layer and multiple-access scheme are similar to those of the downlink, the uplink also poses some unique challenges. Some of the desirable attributes for the LTE uplink include:

- Orthogonal uplink transmission by different User Equipment (UEs), to minimize intra-cell interference and maximize capacity;
- Flexibility to support a wide range of data rates, and to enable the data rate to be adapted to the Signal-to-Interference-plus-Noise Ratio (SINR);
- Sufficiently low Peak-to-Average Power Ratio (PAPR) (or Cubic Metric (CM) – see Section 21.3.3) of the transmitted waveform, to avoid excessive cost, size and power consumption of the UE Power Amplifier (PA);
- Ability to exploit the frequency diversity afforded by the wideband channel (up to 20 MHz), even when transmitting at low data rates;
- Support for frequency-selective scheduling;
- Support for advanced multiple-antenna techniques, to exploit spatial diversity and enhance uplink capacity.

The multiple-access scheme selected for the LTE uplink so as to fulfil these principle characteristics is Single Carrier-Frequency Division Multiple Access (SC-FDMA).

A major advantage of SC-FDMA over the Direct-Sequence Code Division Multiple Access (DS-SS) scheme used in UMTS is that it achieves intra-cell orthogonality even in frequency-selective channels. SC-FDMA avoids the high level of intra-cell interference associated with DS-SS which significantly reduces system capacity and limits the use

LTE – The UMTS Long Term Evolution: From Theory to Practice, Second Edition.

Stefania Sesia, Issam Toufik and Matthew Baker.

© 2011 John Wiley & Sons, Ltd. Published 2011 by John Wiley & Sons, Ltd.

of adaptive modulation. A code-multiplexed uplink also suffers the drawback of an increased CM/PAPR if multi-code transmission is used from a single UE.

The use of OFDMA (Orthogonal Frequency Division Multiple Access) for the LTE uplink would have been attractive due to the possibility for full uplink-downlink commonality. In principle, an OFDMA scheme similar to the LTE downlink could satisfy all the uplink design criteria listed above, except for low CM/PAPR. As outlined in Section 5.2.2.1, much research has been conducted in recent years on methods to reduce the CM/PAPR of OFDM; however, in general the effectiveness of these methods requires careful practical evaluation against their associated complexity and/or overhead (for example, in terms of additional signalling or additional transmission bandwidth used to achieve the CM/PAPR reduction).

SC-FDMA combines the desirable characteristics of OFDM outlined in Section 5.2 with the low CM/PAPR of single-carrier transmission schemes.

Like OFDM, SC-FDMA divides the transmission bandwidth into multiple parallel subcarriers, with the orthogonality between the subcarriers being maintained in frequency-selective channels by the use of a Cyclic Prefix (CP) or guard period. The use of a CP prevents Inter-Symbol Interference (ISI) between SC-FDMA information blocks. It transforms the linear convolution of the multipath channel into a circular convolution, enabling the receiver to equalize the channel simply by scaling each subcarrier by a complex gain factor as explained in Chapter 5.

However, unlike OFDM, where the data symbols directly modulate each subcarrier independently (such that the amplitude of each subcarrier at a given time instant is set by the constellation points of the digital modulation scheme), in SC-FDMA the signal modulated onto a given subcarrier is a linear combination of all the data symbols transmitted at the same time instant. Thus in each symbol period, all the transmitted subcarriers of an SC-FDMA signal carry a component of each modulated data symbol. This gives SC-FDMA its crucial single-carrier property, which results in the CM/PAPR being significantly lower than pure multicarrier transmission schemes such as OFDM.

14.2 SC-FDMA Principles

14.2.1 SC-FDMA Transmission Structure

An SC-FDMA signal can, in theory, be generated in either the time domain or the frequency domain [1]. Although the two techniques are duals and ‘functionally’ equivalent, in practice, the time-domain generation is less bandwidth-efficient due to time-domain filtering and associated requirements for filter ramp-up and ramp-down times [2, 3]. Nevertheless, we describe both approaches here to facilitate understanding of the principles of SC-FDMA in both domains.

14.2.2 Time-Domain Signal Generation

Time-domain generation of an SC-FDMA signal is shown in Figure 14.1.

The input bit stream is mapped into a single-carrier stream of QPSK or QAM¹ symbols, which are grouped into symbol-blocks of length M . This may be followed by an optional repetition stage, in which each block is repeated L times, and a user-specific frequency shift,

¹QPSK: Quadrature Phase Shift Keying; QAM: Quadrature Amplitude Modulation.




Figure 14.1: SC-FDMA time-domain transmit processing.

by which each user’s transmission may be translated to a particular part of the available bandwidth. A CP is then inserted. After filtering (e.g. with a Root-Raised Cosine (RRC) pulse-shaping filter), the resulting signal is transmitted.

The repetition of the symbol blocks results in the spectrum of the transmitted signal only being non-zero at certain subcarrier frequencies, i.e. every L^{th} subcarrier as shown in the example in Figure 14.2.




Figure 14.2: Distributed transmission with equal spacing between occupied subcarriers.

Since such a signal occupies only one in every L subcarriers, the transmission is said to be ‘distributed’ and is one way of providing a frequency-diversity gain.

When no symbol-block repetition is performed ($L = 1$), the signal occupies consecutive subcarriers² and the transmission is said to be ‘localized’. Localized transmissions are beneficial for supporting frequency-selective scheduling or inter-cell interference coordination (as explained in Section 12.5). Localized transmission may also provide frequency diversity if the set of consecutive subcarriers is hopped in the frequency domain.

Different users’ transmissions, using different repetition factors or bandwidths, remain orthogonal on the uplink when the following conditions are met:

- The users occupy different sets of subcarriers. This may in general be accomplished either by introducing a user-specific frequency shift (for localized transmissions) or alternatively by arranging for different users to occupy interleaved sets of subcarriers

²The occupied bandwidth then depends on the symbol rate.

(for distributed transmissions). The latter method is known in the literature as Interleaved Frequency Division Multiple Access (IFDMA) [4].

- The received signals are properly synchronized in time and frequency.
- The CP is longer than the sum of the delay spread of the channel and any residual timing synchronization error between the users.

14.2.3 Frequency-Domain Signal Generation (DFT-S-OFDM)

Generation of an SC-FDMA signal in the frequency domain uses a Discrete Fourier Transform-Spread-OFDM (DFT-S-OFDM) structure [5–7] as shown in Figure 14.3.

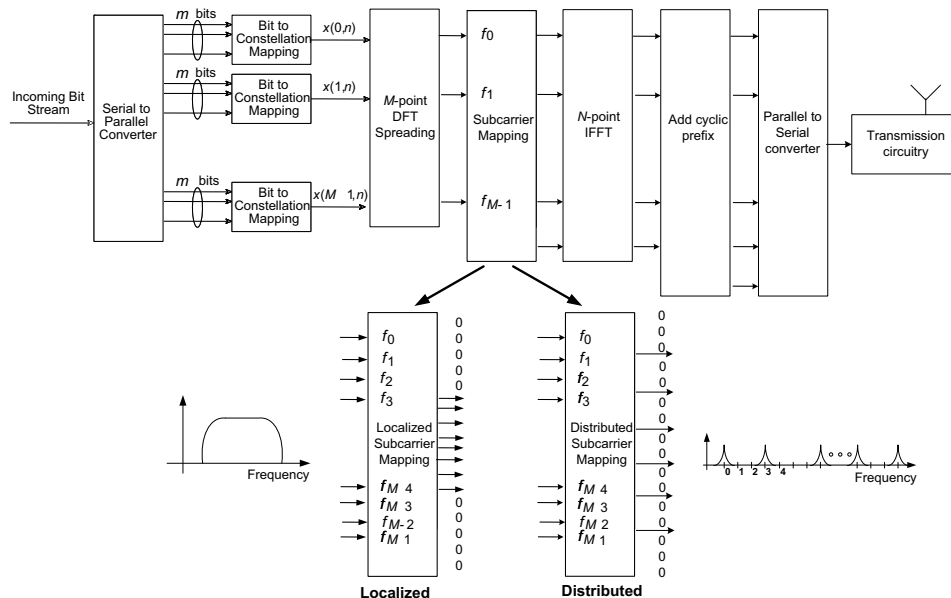


Figure 14.3: SC-FDMA frequency-domain transmit processing (DFT-S-OFDM) showing localized and distributed subcarrier mappings.

The first step of DFT-S-OFDM SC-FDMA signal generation is to perform an M -point DFT operation on each block of M QAM data symbols. Zeros are then inserted among the outputs of the DFT in order to match the DFT size to an N -subcarrier OFDM modulator (typically an Inverse Fast Fourier Transform (IFFT)). The zero-padded DFT output is mapped to the N subcarriers, with the positions of the zeros determining to which subcarriers the DFT-precoded data is mapped.

Usually N is larger than the maximum number of occupied subcarriers, thus providing for efficient oversampling and ‘sinc’ ($\sin(x)/x$) pulse-shaping. The equivalence of DFT-S-OFDM and a time-domain-generated SC-FDMA transmission can readily be seen by considering the

case of $M = N$, where the DFT operation cancels the IFFT of the OFDM modulator resulting in the data symbols being transmitted serially in the time domain. However, this simplistic construction would not provide any oversampling or pulse-shape filtering.

As with the time-domain approach, DFT-S-OFDM is capable of generating both localized and distributed transmissions:

- **Localized transmission.** The subcarrier mapping allocates a group of M adjacent subcarriers to a user. $M < N$ results in zeros being appended to the output of the DFT spreader resulting in an upsampled/interpolated version of the original M QAM data symbols at the IFFT output of the OFDM modulator. The transmitted signal is thus similar to a narrowband single carrier with a CP (equivalent to time-domain generation with repetition factor $L = 1$) and ‘sinc’ pulse-shaping filtering.
- **Distributed transmission.** The subcarrier mapping allocates M equally spaced subcarriers (e.g. every L^{th} subcarrier). $(L - 1)$ zeros are inserted between the M DFT outputs, and additional zeros are appended to either side of the DFT output prior to the IFFT ($ML < N$). As with the localized case, the zeros appended on either side of the DFT output provide upsampling or ‘sinc’ interpolation, while the zeros inserted between the DFT outputs produce waveform repetition in the time domain. This results in a transmitted signal similar to time-domain IFDMA with repetition factor L and ‘sinc’ pulse-shaping filtering.

Like the time-domain SC-FDMA signal generation (in Section 14.2.2), orthogonality between different users with different data rate requirements can be achieved by assigning each user a unique set of subcarriers. The CP structure is the same as for the time-domain signal generation, and therefore the same efficient FDE techniques can be employed at the receiver [8,9].

It is worth noting that, in principle, any unitary matrix can be used instead of the DFT for the spreading operation with similar performance [10]. However, the use of non-DFT spreading would result in increased CM/PAPR since the transmitted signal would no longer have the single carrier characteristic.

14.3 SC-FDMA Design in LTE

Having outlined the key principles of SC-FDMA transmission, we now explain the application of SC-FDMA to the LTE uplink.

14.3.1 Transmit Processing for LTE

Although the frequency-domain generation of SC-FDMA (DFT-S-OFDM) is functionally equivalent to the time-domain SC-FDMA signal generation, each technique requires a slightly different parameterization for efficient signal generation [3]. The pulse-shaping filter used in the time domain SC-FDMA generation approach in practice has a non-zero excess bandwidth, resulting in bandwidth efficiency which is smaller than that achievable with the frequency-domain method with its inherent ‘sinc’ (zero excess bandwidth) pulse-shaping filter which arises from the zero padding and IFFT operation. The non-zero excess bandwidth pulse-shaping filter in the time-domain generation also requires ramp-up and ramp-down times of 3–4 samples duration, while for DFT-S-OFDM there is no explicit pulse-shaping

filter, resulting in a much shorter ramp time similar to OFDM. However, the pulse-shaping filter in the time-domain generation does provide the benefit of reduced CM by approximately 0.25–0.5 dB compared to DFT-S-OFDM, as shown in Figure 14.4. Thus there is a trade-off between bandwidth efficiency and CM/PAPR reduction between the time- and frequency-domain SC-FDMA generation methods.




Figure 14.4: Cubic metric comparison of time-domain SC-FDMA generation (IFDMA), frequency-domain SC-FDMA generation (DFT-S-OFDM) and OFDM.

Frequency-domain signal generation for the LTE uplink has a further benefit in that it allows a very similar parameterization to be adopted as for the OFDM downlink, including the same subcarrier spacing, number of occupied subcarriers in a given bandwidth, and CP lengths. This provides maximal commonality between uplink and downlink, including for example the same clock frequency.

For these reasons, the SC-FDMA parameters chosen for the LTE uplink have been optimized under the assumption of frequency-domain DFT-S-OFDM signal generation.

An important feature of the LTE SC-FDMA parameterization is that the numbers of subcarriers which can be allocated to a UE for transmission are restricted such that the DFT size in LTE can be constructed from multiples of 2, 3 and/or 5. This enables efficient, low-complexity mixed-radix FFT implementations.

14.3.2 SC-FDMA Parameters for LTE

The same basic transmission resource structure is used for the uplink as for the downlink: a 10 ms radio frame is divided into ten 1 ms subframes each consisting of two 0.5 ms slots. As LTE SC-FDMA is based on the same fundamental processing as OFDM, it uses the same 15 kHz subcarrier spacing as the downlink. The uplink transmission resources are also defined in the frequency domain (i.e. before the IFFT), with the smallest unit of resource being a Resource Element (RE), consisting of one SC-FDMA data block length on one subcarrier. As in the downlink, a Resource Block (RB) comprises 12 REs in the frequency domain for a duration of 1 slot, as detailed in Section 6.2. The LTE uplink SC-FDMA physical layer parameters for Frequency Division Duplex (FDD) and Time Division Duplex (TDD) deployments are detailed in Table 14.1.

Table 14.1: LTE uplink SC-FDMA physical layer parameters.

Parameter	Value	Comments
Subframe duration	1 ms	
Slot duration	0.5 ms	
Subcarrier spacing	15 kHz	
SC-FDMA symbol duration	66.67 μ s	
CP duration	Normal CP:	5.2 μ s first symbol in each slot, 4.69 μ s all other symbols
	Extended CP:	16.67 μ s all symbols
Number of symbols per slot	Normal CP: 7	
	Extended CP: 6	
Number of subcarriers per RB	12	

Two CP durations are supported – a normal CP of duration 4.69 μ s and an extended CP of 16.67 μ s, as in the downlink (see Section 5.4.1). The extended CP is beneficial for deployments with large channel delay-spread characteristics, and for large cells.

The 1 ms subframe allows a 1 ms scheduling interval (or Transmission Time Interval (TTI)), as for the downlink, to enable low latency. However, one difference from the downlink is that the uplink coverage is more likely to be limited by the maximum transmission power of the UE. In some situations, this may mean that a single Voice-over-IP (VoIP) packet, for example, cannot be transmitted in a 1 ms subframe with an acceptable error rate. One solution to this is to segment the VoIP packet at higher layers to allow it to be transmitted over several subframes. However, such segmentation results in additional signalling overhead for each segment (including resource allocation signalling and Hybrid ARQ acknowledgement signalling). A more efficient technique for improving uplink VoIP coverage at the cell edge is to use the so-called *TTI bundling*, where a single transport block from the MAC layer is transmitted repeatedly in multiple consecutive subframes, with only one set of signalling messages for the whole transmission. The LTE uplink allows groups of 4 TTIs to be ‘bundled’ in this way, in addition to the normal 1 ms TTI.

In practice in LTE Releases 8 and 9, all the uplink data transmissions are localized, using contiguous blocks of subcarriers. This simplifies the transmission scheme. Frequency-diversity can still be exploited by means of frequency hopping, which can occur both within one subframe (at the boundary between the two slots) and between subframes. In the case of frequency hopping within a subframe, the channel coding spans the two transmission frequencies, and therefore the frequency-diversity gain is maximized through the channel decoding process. The only instance of distributed transmission in the LTE uplink (using an IFDMA-like structure) is for the Sounding Reference Signals (SRSSs) which are transmitted to enable the eNodeB to perform uplink frequency-selective scheduling; these are discussed in Section 15.6.

Like the downlink, the LTE uplink supports scalable system bandwidths from approximately 1.4 MHz up to 20 MHz with the same subcarrier spacing and symbol duration for all bandwidths. The uplink scaling for the bandwidths supported in LTE is shown in

Table 14.2: LTE Uplink SC-FDMA parametrization for selected carrier bandwidths.

	Carrier bandwidth (MHz)					
	1.4	3	5	10	15	20
FFT size	128	256	512	1024	1536	2048
Sampling rate: $M/N \times 3.84$ MHz	1/2	1/1	2/1	4/1	6/1	8/1
Number of subcarriers	72	180	300	600	900	1200
Number of RBs	6	15	25	50	75	100
Bandwidth efficiency (%)	77.1	90	90	90	90	90

Table 14.2. Note that the sampling rates resulting from the indicated FFT sizes are designed to be small rational multiples of the UMTS 3.84 MHz chip rate, for ease of implementation in a multimode UE.

Note that in the OFDM downlink parameter specification, the d.c. subcarrier is unused in order to support direct conversion (zero IF³) architectures. In contrast, no unused d.c. subcarrier is possible for SC-FDMA (as shown in Table 14.2) as it can affect the low CM/PAPR property of the transmit signal.

14.3.3 d.c. Subcarrier in SC-FDMA

Direct conversion transmitters and receivers can introduce distortion at the carrier frequency (zero frequency or d.c. in baseband), for example arising from local oscillator leakage [11]. In the LTE downlink, this is addressed by the inclusion of an unused d.c. subcarrier. However, for the uplink, the same solution would cause significant degradation to the low CM property of the transmitted signal. In order to minimize the impact of such distortion on the packet error rate and the CM/PAPR, the d.c. subcarrier of the SC-FDMA signal is modulated in the same way as all the other subcarriers but the subcarriers are all frequency-shifted by half a subcarrier spacing (± 7.5 kHz), resulting in an offset of 7.5 kHz relative to d.c. Thus two subcarriers straddle the d.c. location, and hence the amount of distortion affecting any individual RB is halved. This is illustrated in Figure 14.5 for deployments with even and odd numbers of RBs across the system bandwidth.

14.3.4 Pulse Shaping

As explained in Sections 14.2.3 and 14.3.1, one of the benefits of frequency-domain processing for SC-FDMA is that, in principle, there is no need for explicit pulse-shaping thanks to the implicit ‘sinc’ pulse-shaping. Nevertheless, an additional explicit pulse-shaping filter can further reduce the CM/PAPR, but at the expense of spectral efficiency (due to the resulting non-zero excess filter bandwidth similar to that of the time domain SC-FDMA signal generation described in Section 14.2.2).

The use of pulse-shaping filters, such as RRC⁴ [12] or Kaiser window [13], was considered for LTE in order to reduce CM/PAPR, especially for lower-order modulations such as QPSK

³Intermediate Frequency.

⁴Roor-Raised-Cosine.

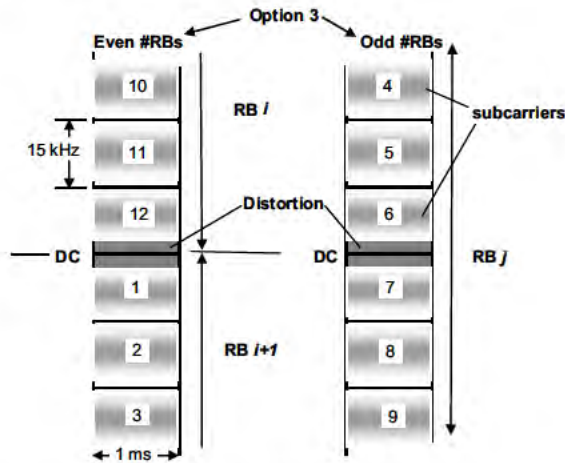


Figure 14.5: 7.5 kHz subcarrier shift for LTE system bandwidths with even or odd numbers of RBs.

to enhance uplink coverage. However, the bandwidth expansion caused by pulse-shaping filtering would require the code rate to be increased for a given data rate, and, in order to provide a net performance benefit, power boosting would be required with the size of the boost being approximately equal to the CM reduction achieved by the pulse-shaping.

Simulation results show that the CM benefits obtained from the use of pulse shaping (RRC or Kaiser window) are small. Moreover, pulse-shaping could not be applied to reference symbols, in order to prevent degradation of the channel estimation.

Finally, as discussed in Chapter 21, CM is not a problem for UEs operating at a maximum power level for resource allocations of up to 10 RBs. This means that there is no benefit possible from pulse-shaping given that UEs at the cell edge would often have less than 10 RBs allocated.

Therefore, no pulse-shaping is used in LTE.

14.4 Summary

The important properties of the SC-FDMA transmission scheme used for the LTE uplink are derived from its multicarrier OFDM-like structure with single-carrier characteristic. The multicarrier-based structure gives the LTE uplink the same robustness against ISI as the LTE downlink, with low-complexity frequency-domain equalization being facilitated by the CP. At the same time, the DFT-based precoding ensures that the LTE uplink possesses the low CM required for efficient UE design. Crucially, the LTE uplink is designed to be orthogonal in the frequency domain between different UEs, thus virtually eliminating the intra-cell interference associated with CDMA.

The parameters of the LTE uplink are designed to ensure maximum commonality with the downlink, and to facilitate frequency-domain DFT-S-OFDM signal generation.

The localized resource allocation scheme of the LTE uplink allows both frequency-selective scheduling and the exploitation of frequency diversity, the latter being achieved by means of frequency hopping.

The following chapters explain the application of SC-FDMA in LTE in more detail, showing how reference signals for channel estimation, data transmissions and control signalling are multiplexed together.

Enhancements to the uplink transmission scheme in Release 10 for LTE-Advanced are described in Section 28.3.6.

References⁵

- [1] Motorola, ‘R1-050397: EUTRA Uplink Numerology and Frame Structure’, www.3gpp.org, 3GPP TSG RAN WG1, meeting 41, Athens, Greece, May 2005.
- [2] Motorola, ‘R1-050584: EUTRA Uplink Numerology and Design’, www.3gpp.org, 3GPP TSG RAN WG1, meeting 41bis, Sophia Antipolis, France, June 2005.
- [3] Motorola, ‘R1-050971: Single Carrier Uplink Options for E-UTRA’, www.3gpp.org, 3GPP TSG RAN WG1, meeting 42, London, UK, August 2005.
- [4] U. Sorger, I. De Broeck and M. Schnell, ‘Interleaved FDMA – A New Spread-Spectrum Multiple-Access Scheme’, in *Proc. IEEE International Conference on Communications*, pp. 1013–1017, 1998.
- [5] K. Bruninghaus and H. Rohling, ‘Multi-carrier Spread Spectrum and its Relationship to Single-carrier Transmission’, in *Proc. IEEE Vehicular Technology Conference*, May 1998.
- [6] D. Galda and H. Rohling, ‘A Low Complexity Transmitter Structure for OFDM-FDMA Uplink Systems’, in *Proc. IEEE Vehicular Technology Conference*, May 2002.
- [7] R. Dinis, D. Falconer, C. T. Lam and M. Sabbaghian, ‘A Multiple Access Scheme for the Uplink of Broadband Wireless Systems’, in *Proc. IEEE Global Telecommunications Conference*, November 2004.
- [8] H. Sari, G. Karam and I. Jeanclaude, ‘Frequency-Domain Equalization of Mobile Radio and Terrestrial Broadcast Channels’, in *Proc. IEEE Global Telecommunications Conference*, November 1994.
- [9] D. Falconer, S. L. Ariyavisitakul, A. Benyamin-Seeyar and B. Eidson, ‘Frequency Domain Equalization for Single-carrier Broadband Wireless Systems’. *IEEE Communications Magazine*, Vol. 40, pp. 58–66, April 2002.
- [10] V. Nangia and K. L. Baum, ‘Experimental Broadband OFDM System – Field Results for OFDM and OFDM with Frequency Domain Spreading’, in *Proc. IEEE International Conference on Vehicular Technology*, Vancouver, BC, Canada, September 2002.
- [11] Motorola, ‘R1-062061: Uplink DC Subcarrier Distortion Considerations in LTE’, www.3gpp.org, 3GPP TSG RAN WG1, meeting 46, Tallinn, Estonia, August 2006.
- [12] NTT DoCoMo, NEC, and SHARP, ‘R1-050702: DFT-spread OFDM with Pulse Shaping Filter in Frequency Domain in Evolved UTRA Uplink’, www.3gpp.org, 3GPP TSG RAN WG1, meeting 42, London, UK, August 2005.
- [13] Huawei, ‘R1-051434: Optimum family of spectrum-shaping functions for PAPR reduction in SC-FDMA’, www.3gpp.org, 3GPP TSG RAN WG1, meeting 43, Seoul, Korea, November 2005.

⁵All web sites confirmed 1st March 2011.

15

Uplink Reference Signals

Robert Love and Vijay Nangia

15.1 Introduction

As in the downlink, the LTE Single-Carrier Frequency Division Multiple Access (SC-FDMA) uplink incorporates Reference Signals (RSs) for data demodulation and channel sounding. In this chapter, the design principles behind these RSs are explained, including in particular features related to interference randomization and coordination, and the flexible configuration of channel sounding.

The roles of the uplink RSs include enabling channel estimation to aid coherent demodulation, channel quality estimation for uplink scheduling, power control, timing estimation and direction-of-arrival estimation to support downlink beamforming. Two types of RS are supported on the uplink:

- **DeModulation RS (DM-RS)**, associated with transmissions of uplink data on the Physical Uplink Shared CHannel (PUSCH) and/or control signalling on the Physical Uplink Control CHannel (PUCCH).¹ These RSs are primarily used for channel estimation for coherent demodulation.
- **Sounding RS (SRS)**, not associated with uplink data and/or control transmissions, and primarily used for channel quality determination to enable frequency-selective scheduling on the uplink.

The uplink RSs are time-multiplexed with the data symbols. The DM-RSs of a given UE² occupy the same bandwidth (i.e. the same Resource Blocks (RBs)) as its PUSCH/PUCCH data transmission. Thus, the allocation of orthogonal (in frequency) sets of RBs to different

¹See Chapter 16 for details of the uplink physical channels.

²User Equipment.

UEs for data transmission automatically ensures that their DM-RSs are also orthogonal to each other.

If configured by higher-layer signalling, SRSs are transmitted on the last SC-FDMA symbol in a subframe; SRS can occupy a bandwidth different from that used for data transmission, in order to allow wider bandwidth sounding. UEs transmitting SRSs in the same subframe can be multiplexed via either Frequency or Code Division Multiplexing (FDM or CDM respectively), as explained in Section 15.6.

Desirable characteristics for the uplink RSs include:

- Constant amplitude in the frequency domain for equal excitation of all the allocated subcarriers for unbiased channel estimates;
- Low Cubic Metric (CM) in the time domain (at worst no higher than that of the data transmissions; a lower CM for RSs than the data can be beneficial in enabling the transmission power of the RSs to be boosted at the cell-edge);
- Good autocorrelation properties for accurate channel estimation;
- Good cross-correlation properties between different RSs to reduce interference from RSs transmitted on the same resources in other (or, in some cases, the same) cells.

The following sections explain how these characteristics are achieved in LTE.

15.2 RS Signal Sequence Generation

The uplink RSs in LTE are mostly based on Zadoff–Chu (ZC) sequences [1, 2]. The fundamental structure and properties of these sequences are described in Section 7.2.1.

These sequences satisfy the desirable properties for RSs mentioned above, exhibiting 0 dB CM, ideal cyclic autocorrelation and optimal cross-correlation. The cross-correlation property results in the impact of an interfering signal being spread evenly in the time domain after time-domain correlation of the received signal with the desired sequence; this results in more reliable detection of the significant channel taps. However, in practice the CM of a ZC sequence is degraded at the Nyquist sampling rate because of the presence of unused guard subcarriers at each end of the sequence (due to the number of occupied RS subcarriers being less than the IFFT³ size of the OFDM⁴ modulator) and results in the ZC sequence effectively being oversampled in the time domain.

The RS sequence length, N_p , is equal to the number of assigned subcarriers M_{sc}^{RS} , which is a multiple of the number of subcarriers per RB, $N_{sc}^{RB} = 12$, i.e.

$$N_p = M_{sc}^{RS} = m \cdot N_{sc}^{RB}, \quad 1 \leq m \leq N_{RB}^{UL} \quad (15.1)$$

where N_{RB}^{UL} is the uplink system bandwidth in terms of RBs.

The length- N_p RS sequence is directly applied (without Discrete Fourier Transform (DFT) spreading) to N_p RS subcarriers at the input of the IFFT as shown in Figure 15.1 [3]. Recall that a ZC sequence of odd-length N_{ZC} is given by

$$a_q(n) = \exp\left[-j2\pi q \frac{n(n+1)/2 + ln}{N_{ZC}}\right] \quad (15.2)$$

³Inverse Fast Fourier Transform

⁴Orthogonal Frequency Division Multiplexing

where $q = 1, \dots, N_{ZC} - 1$ is the ZC sequence index (also known as the root index), $n = 0, 1, \dots, N_{ZC} - 1$, and $l = 0$ in LTE.




Figure 15.1: Transmitter structure for SC-FDMA reference signals. Note that no DFT spreading is applied to the RS sequence.

In LTE, N_{ZC} is selected to be the largest prime number smaller than or equal to N_p . The ZC sequence of length N_{ZC} is then cyclically extended to the target length N_p as follows:

$$\bar{r}_q(n) = a_q(n \bmod N_{ZC}), \quad n = 0, 1, \dots, N_p - 1 \tag{15.3}$$

The cyclic extension in the frequency domain preserves the constant amplitude property. Cyclic extension of the ZC sequences is used rather than truncation, as in general it provides better CM characteristics. For sequence lengths of three or more RBs, this provides at least 30 sequences with CM smaller than or close to that of QPSK.

However, for the shortest sequence lengths, suitable for resource allocations of just one or two RBs, only a small number of low-CM extended ZC sequences is available (six and 12 sequences respectively with CM less than that of QPSK). Therefore, in order to provide at least as large a number of sequences as for the 3 RB case, 30 special RS sequences are defined in LTE for resource allocations of one or two RBs. These special sequences are QPSK rather than ZC-based sequences and were obtained from computer searches so as to have constant modulus in the frequency-domain, low CM, low memory and complexity requirements and good cross-correlation properties.

The QPSK RS sequences are given by,

$$\bar{r}(n) = e^{j\varphi(n)\pi/4}, \quad n = 0, 1, \dots, M_{sc}^{RS} - 1, \tag{15.4}$$

where M_{sc}^{RS} is the number of subcarriers to which the sequence is mapped, and the values of $\varphi(n)$ can be found in [4, Tables 5.5.1.2-1 and 5.5.1.2-2].

15.2.1 Base RS Sequences and Sequence Grouping

In order for a cell to support uplink transmissions of different bandwidths, it is necessary to assign a cell at least one *base RS sequence* for each possible RB allocation size. Multiple RS sequences for each allocation size are then derived from each base sequence by means of different cyclic time shifts, as explained in Section 15.2.2.

The smallest number of available base sequences is for resource allocations of three RBs, where, as noted above, only 30 extended ZC sequences exist. As a result, the complete set of available base sequences across all RB allocation sizes is divided into 30 non-overlapping *sequence-groups*. A cell is then assigned one of the sequence-groups for uplink transmissions. For each resource allocation size up to and including five RBs, each of the 30 sequence-groups contains only one base sequence, since for five RBs (i.e. sequences of length 60) only 58 extended ZC-sequences are available. For sequence lengths greater than five RBs, more extended ZC-sequences are available, and therefore each of the 30 sequence-groups contains two base sequences per resource allocation size; this is exploited in LTE to support *sequence hopping* (within the sequence-group) between the two slots of a subframe.

The base sequences for resource allocations larger than three RBs are selected such that they are the sequences with high cross-correlation to the single 3 RB base sequence in the sequence-group [5]. Since the cross-correlation between the 3 RB base sequences of different sequence-groups is low due to the inherent properties of the ZC sequences, such a method for assigning the longer base sequences to sequence-groups helps to ensure that the cross-correlation between sequence-groups is kept low, thus reducing inter-cell interference.

The v base RS sequences of length 3 RBs or larger (i.e. $M_{sc}^{RS} \geq 36$) assigned to a sequence-group u are given by,

$$\bar{r}_{u,v}(n) = a_q(n \bmod N_{ZC}), \quad n = 0, 1, \dots, M_{sc}^{RS} - 1 \quad (15.5)$$

where $u \in \{0, 1, \dots, 29\}$ is the sequence-group number, v is the index of the base sequence of length M_{sc}^{RS} within the sequence-group u , and is given by

$$v = \begin{cases} 0, 1 & \text{for } M_{sc}^{RS} \geq 72 \\ 0 & \text{otherwise} \end{cases} \quad (15.6)$$

N_{ZC} is the largest prime number smaller than M_{sc}^{RS} , and q is the root ZC sequence index (defined in [4, Section 5.5.1.1]).

15.2.2 Orthogonal RS via Cyclic Time-Shifts of a Base Sequence

UEs which are assigned to different RBs transmit RSs in these RBs and hence achieve separability of the RSs via FDM. However, in certain cases, UEs can be assigned to transmit on the same RBs, for example in the case of uplink multi-user MIMO⁵ (sometimes also referred to as Spatial Division Multiple Access (SDMA) or ‘Virtual MIMO’). In these cases, the RSs can interfere with each other, and some means of separating the RSs from the different transmitters is required. Using different base sequences for different UEs transmitting in the same RBs is not ideal due to the non-zero cross-correlation between the base sequences which can degrade the channel estimation at the eNodeB. It is preferable that

⁵Multiple-Input Multiple-Output.

the RS signals from the different UEs are fully orthogonal. In theory, this could be achieved by FDM of the RSs within the same RBs, although this would reduce the RS sequence length and the number of different RS sequences available; this would be particularly undesirable for low-bandwidth transmissions.

Therefore in LTE, orthogonality between RSs occupying the same RBs is instead provided by exploiting the fact that the correlation of a ZC sequence with any Cyclic Shift (CS) of the same sequence is zero (see Section 7.2.1). As the channel impulse response is of finite duration, different transmitters can use different cyclic time shifts of the same base RS sequence, with the RSs remaining orthogonal provided that the cyclic shifts are longer than the channel impulse response.

If the RS SC-FDMA symbol duration is T_p and the channel impulse response duration is less than T_{cs} , then up to T_p/T_{cs} different transmitters can transmit in the same symbol, with different cyclic shift values, with separable channel estimates at the receiver. For example, Figure 15.2 shows that if $T_p/T_{cs} = 4$ and there are four transmitters, then each transmitter $t \in \{1, \dots, 4\}$ can use a cyclic $(t - 1)T_p/4$ of the same base sequence. At the eNodeB receiver, by correlating the received composite signal from the different transmitters occupying the same RBs with the base sequence, the channel estimates from the different transmitters are separable in the time domain [6].




Figure 15.2: Illustration of cyclic time shift orthogonality of RS signals.

Since a cyclic time shift is equivalent to applying a phase ramp in the frequency domain, the representation of a base sequence with cyclic shift, α , in the frequency-domain is given by

$$r_{u,v}^{(\alpha)}(n) = e^{j\alpha n} \bar{r}_{u,v}(n) \tag{15.7}$$

where $\bar{r}_{u,v}(n)$ is the base (or unshifted) sequence of sequence-group u , with base sequence index v within the sequence-group, $\alpha = 2\pi n_t/P$ with n_t the cyclic time shift index for transmitter t , and P is the number of equally spaced cyclic time shifts supported.

In LTE, 12 equally spaced cyclic time shifts are defined for the DM-RS on the PUSCH and PUCCH. This allows for delay spreads up to 5.55 μ s.

The degree of channel estimate separability at the receiver between different cyclic time shifts depends in practice on the (circular) distance between the shifts (as well as the received power differences between the transmitters). Cyclic time shifts spaced the furthest apart experience the least cross-talk between the channel estimates (for example, arising from practical issues such as channel estimation filtering and finite sampling granularity). Thus, when the number of UEs using different cyclic time shifts is less than the number of cyclic

time shifts supported (P), it is beneficial to assign cyclic time shifts with the largest possible (circular) separation; this is supported in LTE.

In LTE the cyclic time shifts also always hop between the two slots in a subframe, for inter-cell interference randomization (see Section 15.4).

15.3 Sequence-Group Hopping and Planning

LTE supports both *RS sequence-group hopping* and *RS sequence-group planning* modes of system deployment [7, 8]; the mode is configurable by Radio Resource Control (RRC) signalling.

The sequence-group assigned to a cell is a function of its physical-layer cell identity (cell-ID) and can be different for PUSCH and PUCCH transmissions. A UE acquires knowledge of the physical layer cell-ID from its downlink synchronization signals, as described in Section 7.2.

15.3.1 Sequence-Group Hopping

The sequence-group hopping mode of deployment can be enabled in a cell by a 1-bit broadcast signalling parameter called ‘groupHoppingEnabled’. This mode actually consists of a combination of hopping and shifting of the sequence-group according to one of 504 sequence-group hopping/shifting patterns corresponding to the 504 unique cell-IDs [9]. Since there are 30 base sequence-groups, $17 (= \lceil 504/30 \rceil)$ unique sequence-group hopping patterns of length 20 are defined (corresponding to the duration of a radio frame with 20 slots), each of which can be offset by one of 30 sequence-group shift offsets. The sequence-group number u depends on the sequence-group hopping pattern f_{gh} and the sequence-group shift offset f_{ss} as defined in [4, Section 5.5.1.3]. The sequence-group hopping pattern changes u from slot to slot in a pseudo-random manner, while the shift offset is fixed in all slots. Both f_{gh} and f_{ss} depend on the cell-ID.

The sequence-group hopping pattern f_{gh} is obtained from a length-31 Gold sequence generator (see Section 6.3) [10, 11], of which the second constituent M-sequence is initialized at the beginning of each radio frame by the sequence-group hopping pattern index of the cell. Up to 30 cell-IDs can have the same sequence-group hopping pattern (e.g. part of a planned coordinated cell cluster), with different sequence-group shift offsets being used to minimize RS collisions and inter-cell interference. The same sequence-group hopping pattern is used for PUSCH DM-RS, SRS and PUCCH DM-RS transmissions. The sequence-group shift offset can be different for PUSCH and PUCCH.

For PUSCH, it should be possible to assign cell-IDs such that the same sequence-group hopping pattern and sequence-group shift offset, and hence the same base sequences, are used in adjacent cells. This can enable the RSs from UEs in adjacent cells (for example, the cells of the same eNodeB) to be orthogonal to each other by using different cyclic time shifts of the same base sequence. Therefore in LTE the sequence-group shift offset for PUSCH is explicitly configured by a cell-specific 5-bit broadcast signalling parameter, ‘groupAssignmentPUSCH’. As the sequence-group shift offset is a function of the cell-ID, the overhead for signalling one of the 504 sequence-group hopping patterns for PUSCH is reduced from nine to five bits, such that $f_{ss} = (\text{cell-ID} \bmod 30 + \Delta_{ss}) \bmod 30$, where $\Delta_{ss} \in \{0, \dots, 29\}$ is indicated by ‘groupAssignmentPUSCH’.

For PUCCH transmissions, as described in Section 16.3, the same RBs at the edge of the system bandwidth are normally used by all cells. Thus, in order to randomize interference on the PUCCH between neighbouring cells which are using the same sequence-group hopping pattern, the sequence-shift offset for PUCCH is simply given by cell-ID mod 30. Similarly, for interference randomization on the SRS transmissions which occur in the same SC-FDMA symbol (see Section 15.6), the same sequence-group shift offset as PUCCH is used.

In Section 15.2.1, it was explained that there are two base sequences per sequence-group for each RS sequence length greater than 60 (5 RBs), with the possibility of interference randomization by sequence-hopping between the two base sequences at the slot boundary in the middle of each subframe. If sequence-group hopping is used, the base sequence automatically changes between each slot, and therefore additional sequence hopping within the sequence group is not needed; hence only the first base sequence in the sequence group is used if sequence-group hopping is enabled.

15.3.2 Sequence-Group Planning

If sequence-group hopping is disabled, the same sequence-group number u is used in all slots of a radio frame and is simply obtained from the sequence group shift offset, $u = f_{ss}$.

Since 30 sequence-groups are defined (see Section 15.2.1), planned sequence-group assignment is possible for up to 30 cells in LTE. This enables neighbouring cells to be assigned sequence groups with low cross-correlation to reduce RS interference, especially for small RB allocations.⁶

An example of sequence-group planning with a conventional six-sequence-group reuse plan is shown in Figure 15.3 [12]. The same sequence-group number (and hence base sequences) are used in the three cells of each eNodeB, with different cyclic time shifts assigned to each cell (only 3 cyclic time shifts, D1, D2 and D3 are shown in this example).

With sequence-group planning, sequence hopping within the sequence group between the two slots of a subframe for interference randomization can be enabled by a 1-bit cell-specific parameter, 'sequenceHoppingEnabled'. The base sequence index for $M_{sc}^{RS} \geq 72$ used in slot n_s is then obtained from the length-31 Gold sequence generator. In order to enable the use of the same base RS sequence (and hopping pattern) in adjacent cells for PUSCH, the pseudo-random sequence generator is initialized at the beginning of each radio frame by the sequence-group hopping pattern index (based on part of the cell-ID), offset by the PUSCH sequence-group shift index of the cell (see [4, Section 5.5.1.4]).

The same hopping pattern within the sequence-group is used for PUSCH DM-RS, SRS, and PUCCH DM-RS.

Further interference randomization is provided by the cyclic time shift hopping, which is always enabled in LTE as discussed in the following section.

15.4 Cyclic Shift Hopping

Cyclic time-shift hopping is always enabled for inter-cell interference randomization for PUSCH and PUCCH transmissions. For PUSCH with $P = 12$ evenly spaced cyclic time shifts, the hopping is between the two slots in a subframe, with the cyclic shift (α in

⁶Power-limited cell-edge UEs are likely to have small RB allocations.




Figure 15.3: Example of RS sequence-group planning – six-sequence-group reuse plan.

Equation (15.7)) for a UE being derived in each slot from a combination of a 3-bit cell-specific broadcast cyclic time shift offset parameter, a 3-bit cyclic time shift offset indicated in each uplink scheduling grant and a pseudo-random cyclic shift offset obtained from the output of the length-31 Gold sequence generator (see [4, Section 5.5.2.1.1]).

As mentioned in Section 15.3, it should be possible to use different cyclic time shifts in adjacent cells with the same sequence-group (e.g. the cells of the same eNodeB), in order to support orthogonal RS transmissions from UEs in different cells. This requirement is similar to the case for initialization of the pseudo-random sequence generator for sequence hopping within a sequence-group, where the same hopping patterns are needed in neighbouring cells. Thus, the initialization of the PUSCH DM-RS cyclic shift pseudo-random sequence generator is the same as that for the sequence hopping pattern generator in Section 15.3.2, initialized every radio frame (see [4, Section 5.5.1.4]).

In the case of PUCCH transmission, cyclic time-shift hopping (among the $P = 12$ evenly spaced cyclic time shifts) is performed per SC-FDMA symbol, with the cyclic shift α for a given SC-FDMA symbol in a given slot being derived (as specified in [4, Section 5.4]) from a combination of the assigned PUCCH resource index (see Section 16.3) and the output of the length-31 Gold sequence generator. In order to randomize interference on the PUCCH arising from the fact that the same band-edge RBs are used for PUCCH transmissions in all cells (see Section 16.3), the pseudo-random sequence generator is initialized at the beginning of each radio frame by the cell-ID.

In addition, to achieve intra-cell interference randomization for the PUCCH DM-RS, the cyclic time shift used in the second slot is hopped such that UEs which are assigned adjacent cyclic time shifts in the first slot use non-adjacent cyclic time shifts (with large separation) in the second slot [13]. A further benefit of using a different cyclic time shift in each slot is that the non-ideal cross-correlation between different base RS sequences is averaged (as the cross-correlation is not constant for all time lags).

15.5 Demodulation Reference Signals (DM-RS)

The DM-RSs associated with uplink PUSCH data or PUCCH control transmissions from a UE are primarily provided for channel estimation for coherent demodulation and are therefore present in every transmitted uplink slot.

RSs could in theory be concentrated in one position in each slot, or divided up and positioned in multiple locations in each slot. Two alternatives were considered in the design of the DM-RSs in LTE, as shown in Figure 15.4:

- One RS symbol per slot, having the same duration as a data SC-FDMA symbol (sometimes referred to as a ‘Long Block’ (LB)), with the RS symbol having the same subcarrier spacing as the data SC-FDMA symbols;
- Two RS symbols per slot, each of half the duration of a data SC-FDMA symbol (sometimes referred to as a ‘Short Block’ (SB)), with the subcarrier spacing in the RS symbols being double that of the data SC-FDMA symbols (i.e. resulting in only six subcarriers per RB in the RS symbols).

As can be seen from Figure 15.4, it would be possible to support twice as many SB RSs in a slot in the time domain compared to the LB RS structure with the same number of data SC-FDMA symbols. However, the frequency resolution of the SB RS would be half that of the LB RS due to the subcarrier spacing being doubled.




Figure 15.4: Example of slot formats considered with (a) one Long Block RS per slot and (b) two Short Block RSs per slot.

Figures 15.5 and 15.6 show a performance comparison of LB and SB RSs for a resource allocation bandwidth of 1 RB for a duration of one subframe (1 ms), for UE speeds of 30 km/h and 250 km/h respectively. The performance of LB RS is similar to that for SB




Figure 15.5: Demodulation performance comparison for Long Block RS and Short Block RS structure. Code rate $r = 1/2$, GSM Typical Urban (TU) channel model, 30 km/h, 2 GHz carrier frequency.




Figure 15.6: Demodulation performance comparison for Long Block RS and Short Block RS structure. Code rate $r = 1/2$, GSM Typical Urban (TU) channel model, 250 km/h, 2 GHz carrier frequency.

RS for medium speeds, with some degradation at high speeds. However, with LB RS the signal parameterization remains the same as that of the downlink OFDM. LB RSs also have the advantage of providing longer sequences for a given bandwidth allocation (due to there being twice as many RS subcarriers), and thus a larger number of RS sequences with desirable characteristics.

Therefore the LB RS structure of Figure 15.4(a) was adopted in LTE for the PUSCH DM-RS.

The exact position of the single PUSCH DM-RS symbol in each uplink slot depends on whether the normal or extended CP is used, as shown in Figure 15.7. For the case of the normal CP with seven SC-FDMA symbols per slot, the PUSCH DM-RS occupies the centre (i.e. fourth) SC-FDMA symbol. With six SC-FDMA symbols per slot in the case of the extended CP, the third SC-FDMA symbol is used. For PUCCH transmission, the position and number of DM-RS depend on the type of uplink control information being transmitted, as discussed in Section 16.3.




Figure 15.7: LTE uplink subframe configuration for PUSCH DM RS: (a) normal CP; (b) extended CP.

The DM-RS occupies the same RBs as the RB allocation for the uplink PUSCH data or PUCCH control transmission. Thus, the RS sequence length, M_{sc}^{RS} , is equal to the number of subcarriers allocated to the UE for PUSCH or PUCCH transmissions. Further, since the PUSCH RB allocation size is limited to multiples of two, three and/or five RBs (as explained in Section 14.3.1), the DM-RS sequence lengths are also restricted to the same multiples.

As discussed in Section 15.2.2, the DM-RS SC-FDMA symbol supports 12 cyclic time shifts with a spacing of 5.55 μ s.

To support inter-cell interference randomization, cyclic time-shift hopping is always enabled for DM-RS as detailed in Section 15.4.

Release 10 enhancements to the uplink DM-RS to support non-contiguous PUSCH resource allocation and uplink Single-User MIMO are explained in Sections 28.3.6.2 and 29.2.1 respectively.

15.6 Uplink Sounding Reference Signals (SRS)

The SRSs, which are not associated with uplink data and/or control transmission, are primarily used for channel quality estimation to enable frequency-selective scheduling on the uplink. However, they can be used for other purposes, such as to enhance power control or to support various start-up functions for UEs not recently scheduled. Some examples include initial Modulation and Coding Scheme (MCS) selection, initial power control for data transmissions, timing advance, and ‘frequency semi-selective scheduling’ in which the frequency resource is assigned frequency-selectively for the first slot of a subframe and hops pseudo-randomly to a different frequency in the second slot [14].

15.6.1 SRS Subframe Configuration and Position

The subframes in which SRSs are transmitted by any UE within the cell are indicated by cell-specific broadcast signalling. A 4-bit cell-specific ‘srsSubframeConfiguration’ parameter indicates 15 possible sets of subframes in which SRS may be transmitted within each radio frame (see [4, Section 5.5.3.3]). This configurability provides flexibility in adjusting the SRS overhead depending on the deployment scenario. A 16th configuration switches the SRS off completely in the cell, which may for example be appropriate for a cell serving primarily high-speed UEs.

The SRS transmissions are always in the last SC-FDMA symbol in the configured subframes, as shown in Figure 15.8. Thus the SRS and DM-RS are located in different SC-FDMA symbols. PUSCH data transmission is not permitted on the SC-FDMA symbol designated for SRS, resulting in a worst-case sounding overhead (with an SRS symbol in every subframe) of around 7%.

15.6.2 Duration and Periodicity of SRS Transmissions

The eNodeB in LTE may either request an individual SRS transmission from a UE or configure a UE to transmit SRS periodically until terminated; a 1-bit UE-specific signalling parameter, ‘duration’, indicates whether the requested SRS transmission is single or periodic. If periodic SRS transmissions are configured for a UE, the periodicity may be any of 2, 5, 10, 20, 40, 80, 160 or 320 ms; the SRS periodicity and SRS subframe offset within the period




Figure 15.8: Uplink subframe configuration with SRS symbol.

in which the UE should transmit its SRS are configured by a 10-bit UE-specific dedicated signalling parameter called ‘srs-ConfigIndex’.

In Release 10, a mechanism for dynamically triggering an aperiodic SRS transmission by means of the PDCCH is introduced; this is explained in Section 29.2.2.

15.6.3 SRS Symbol Structure

In order to support frequency-selective scheduling between multiple UEs, it is necessary that SRS from different UEs with different sounding bandwidths can overlap. In order to support this, Interleaved FDMA (IFDMA, introduced in Section 14.2) is used in the SRS SC-FDMA symbol, with a RePetition Factor (RPF) of 2. The (time-domain) RPF is equivalent to a frequency-domain decimation factor, giving the spacing between occupied subcarriers of an SRS signal with a comb-like spectrum. Thus, RPF = 2 implies that the signal occupies every 2nd subcarrier within the allocated sounding bandwidth as shown by way of example in Figure 15.9. Using a larger RPF could in theory have provided more flexibility in how the bandwidth could be allocated between UEs, but it would have reduced the sounding sequence length (for a given sounding bandwidth) and the number of available SRS sequences (similar to the case for DM-RS). Therefore the RPF is limited to 2.

Due to the IFDMA structure of the SRS symbol, a UE is assigned, as part of its configurable SRS parameters, the ‘transmissionComb’ index (0 or 1) on which to transmit the SRS. The RS sequences used for the SRS are the same as for the DM-RS, resulting in the SRS sequence length being restricted to multiples of two, three and/or five times the RB size. In addition, the SRS bandwidth (in RBs) must be an even number, due to the RPF of 2 and the minimum SRS sequence length being 12. Therefore, the possible SRS bandwidths, N_{RB}^{SRS} (in number of RBs), and the SRS sequence length, M_{sc}^{SRS} , are respectively given by,

$$\begin{aligned}
 N_{RB}^{SRS} &= 2^{(1+\alpha_2)} \cdot 3^{\alpha_3} \cdot 5^{\alpha_5} \\
 M_{sc}^{SRS} &= \frac{1}{2} \cdot N_{RB}^{SRS} \cdot 12
 \end{aligned}
 \tag{15.8}$$

where $\alpha_2, \alpha_3, \alpha_5$ is a set of positive integers. Similarly to the DM-RS, simultaneous SRS can be transmitted from multiple UEs using the same RBs and the same offset of the comb, using different cyclic time shifts of the same base sequence to achieve orthogonal separation (see Section 15.2.2). For the SRS, eight (evenly spaced) cyclic time shifts per SRS comb are supported (see [4, Section 5.5.3.1]), with the cyclic shift being configured individually for each UE.

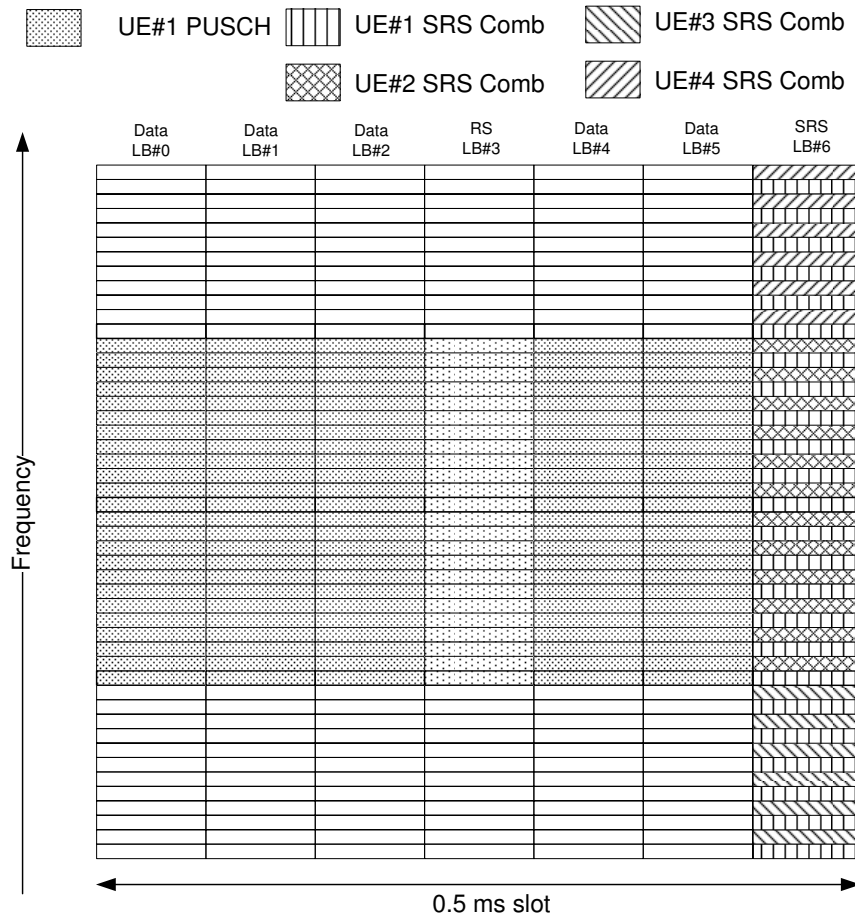


Figure 15.9: SRS symbol structure with RPF = 2.

15.6.3.1 SRS Bandwidths

Some of the factors which affect the SRS bandwidth are the maximum power of the UE, the number of supportable sounding UEs, and the sounding bandwidth needed to benefit from uplink channel-dependent scheduling. Full bandwidth sounding provides the most complete channel information when the UE is sufficiently close to the eNodeB, but degrades as the path-loss increases when the UE cannot further increase its transmit power to maintain the transmission across the full bandwidth. Full bandwidth transmission of SRS also limits the number of simultaneous UEs whose channels can be sounded, due to the limited number of cyclic time shifts (eight cyclic time shifts per SRS comb as explained above).

To improve the SNR and support a larger number of SRSs, up to four SRS bandwidths can be simultaneously supported in LTE depending on the system bandwidth. To provide

flexibility with the values for the SRS bandwidths, eight sets of four SRS bandwidths are defined for each possible system bandwidth. RRC signalling indicates which of the eight sets is applicable in the cell by means of a 3-bit cell-specific parameter ‘srs-BandwidthConfig’. This allows some variability in the maximum SRS bandwidths, which is important as the SRS region does not include the PUCCH region near the edges of the system bandwidth (see Section 16.3), which is itself variable in bandwidth. An example of the eight sets of four SRS bandwidths applicable to uplink system bandwidths in the range 40–60 RBs is shown in Table 15.1 (see [4, Table 5.5.3.2-2]).

Table 15.1: SRS BandWidth (BW) configurations for system bandwidths 40–60 RBs (see [4, Table 5.5.3.2-2]). Reproduced by permission of © 3GPP.

Configuration	Number of RBs			
	SRS-BW 0	SRS-BW 1	SRS-BW 2	SRS-BW 3
0	48	24	12	4
1	48	16	8	4
2	40	20	4	4
3	36	12	4	4
4	32	16	8	4
5	24	4	4	4
6	20	4	4	4
7	16	4	4	4

The specific SRS bandwidth to be used by a given UE is configured by a further 2-bit UE-specific parameter, ‘srs-Bandwidth’.

As can be seen from Table 15.1, the smallest sounding bandwidth supported in LTE is 4 RBs. A small sounding bandwidth of 4 RBs provides for higher-quality channel information from a power-limited UE. The sounding bandwidths are constrained to be multiples of each other, i.e. following a tree-like structure, to support frequency hopping of the different narrowband SRS bandwidths (see [4, Section 5.5.3.2]). Frequency hopping can be enabled or disabled for an individual UE based on the value of the parameter ‘freqDomainPosition’. The tree structure of the SRS bandwidths limits the possible starting positions for the different SRS bandwidths, reducing the overhead for signalling the starting position to 5 bits (signalled to each UE by the parameter ‘freqDomainPosition’).

Table 15.2 summarizes the various SRS configurable parameters which are signalled to a UE [15].

15.7 Summary

The uplink reference signals provided in LTE fulfil an important function in facilitating channel estimation and channel sounding. The ZC-based sequence design can be seen to be a good match to this role, with constant amplitude in the frequency domain and the ability to provide a large number of sequences with zero or low correlation. This enables both interference randomization and interference coordination techniques to be employed in LTE

Table 15.2: Uplink SRS configurable parameters.

Sounding RS parameter name	Significance	Signalling type
srs-BandwidthConfig	Maximum SRS bandwidth in the cell	Cell-specific
srs-SubframeConfig	Sets of subframes in which SRS may be transmitted in the cell	Cell-specific
srs-Bandwidth	SRS transmission bandwidth for a UE	UE-specific
freqDomainPosition	Frequency-domain position	UE-specific
srs-HoppingBandwidth	Frequency hop size	UE-specific
duration	Single SRS or periodic	UE-specific
srs-ConfigIndex	Periodicity and subframe offset	UE-specific
transmissionComb	Transmission comb offset	UE-specific
$n_{\text{SRS}}^{\text{CS}}$	Cyclic shift	UE-specific

system deployments, as appropriate to the scenario. A high degree of flexibility is provided for configuring the reference signals, especially for the sounding reference signals, where the overhead arising from their transmission can be traded off against the improvements in system efficiency which may be achievable from frequency-selective uplink scheduling.

Enhancements to the uplink RSs for LTE-Advanced are outlined in Sections 28.3.6.2 and 29.2.1 for the DM-RS and in Section 29.2.2 for the SRS.

References⁷

- [1] D. C. Chu, ‘Polyphase Codes With Good Periodic Correlation Properties’. *IEEE Trans. on Information Theory*, pp. 531–532, July 1972.
- [2] B. M Popovic, ‘Generalized Chirp-like Polyphase Sequences with Optimal Correlation Properties’. *IEEE Trans. on Information Theory*, Vol. 38, pp. 1406–1409, July 1992.
- [3] Motorola, ‘R1-060878: EUTRA SC-FDMA Uplink Pilot/Reference Signal Design & TP’, www.3gpp.org, 3GPP TSG RAN WG1, meeting 44bis, Athens, Greece, March 2006.
- [4] 3GPP Technical Specification 36.211, ‘Evolved Universal Terrestrial Radio Access (E-UTRA); Physical Channels and Modulation’, www.3gpp.org.

⁷All web sites confirmed 1st March 2011.

- [5] Huawei, LG Electronics, NTT DoCoMo, and Panasonic, 'R1-080576: Way Forward on the Sequence Grouping for UL DM RS', www.3gpp.org, 3GPP TSG RAN WG1, meeting 51bis, Sevilla, Spain, January 2008.
- [6] K. Fazel and G. P. Fettweis, *Multi-Carrier Spread-Spectrum*. Kluwer Academic Publishers, Dordrecht, Holland, 1997.
- [7] Alcatel-Lucent, Ericsson, Freescale, Huawei, LGE, Motorola, Nokia, Nokia-Siemens Networks, NTT DoCoMo, Nortel, Panasonic, Qualcomm, and TI, 'R1-072584: Way Forward for PUSCH RS', www.3gpp.org, 3GPP TSG RAN WG1, meeting 49, Kobe, Japan, May 2007.
- [8] Alcatel-Lucent, Ericsson, Freescale, Huawei, LGE, Motorola, Nokia, Nokia-Siemens Networks, NTT DoCoMo, Nortel, Panasonic, Qualcomm, and TI, 'R1-072585: Way forward for PUCCH RS', www.3gpp.org, 3GPP TSG RAN WG1, meeting 49, Kobe, Japan, May 2007.
- [9] NTT DoCoMo, 'R1-074278: Hopping and Planning of Sequence Groups for Uplink RS', www.3gpp.org, 3GPP TSG RAN WG1, meeting 50bis, Shanghai, China, October 2007.
- [10] Motorola, 'R1-080719: Hopping Patterns for UL RS', www.3gpp.org, 3GPP TSG RAN WG1, meeting 52, Sorrento, Italy, February 2008.
- [11] Qualcomm Europe, Ericsson, Motorola, Samsung, Panasonic, and NTT DoCoMo, 'R1-081133: WF on UL DM-RS Hopping Pattern Generation', www.3gpp.org, 3GPP TSG RAN WG1, meeting 52, Sorrento, Italy, February 2008.
- [12] Motorola, 'R1-071341: Uplink Reference Signal Planning Aspects', www.3gpp.org, 3GPP TSG RAN WG1, meeting 48bis, St. Julian's, Malta, March 2007.
- [13] Panasonic, Samsung, and ETRI, 'R1-080983: Way forward on the Cyclic Shift Hopping for PUCCH', www.3gpp.org, 3GPP TSG RAN WG1, meeting 52, Sorrento, Italy, February 2008.
- [14] Motorola, 'R1-073756: Benefit of Non-Persistent UL Sounding for Frequency Hopping PUSCH', www.3gpp.org, 3GPP TSG RAN WG1, meeting 50, Athens, Greece, August 2007.
- [15] Ericsson, 'R1-082199: Physical-layer parameters to be configured by RRC', www.3gpp.org, 3GPP TSG RAN WG1, meeting 53, Kansas City, USA, May 2008.

16

Uplink Physical Channel Structure

Robert Love and Vijay Nangia

16.1 Introduction

The LTE Single-Carrier Frequency Division Multiple Access (SC-FDMA) uplink provides separate physical channels for the transmission of data and control signalling, the latter being predominantly to support the downlink data transmissions. The detailed structure of these channels, as explained in this chapter, is designed to make efficient use of the available frequency-domain resources and to support effective multiplexing between data and control signalling.

LTE Release 8 also incorporates uplink multiple antenna techniques, including closed-loop antenna selection and Spatial Division Multiple Access (SDMA) or Multi-User Multiple-Input Multiple-Output (MU-MIMO). More advanced multiple-antenna techniques, including closed-loop spatial multiplexing for Single-User MIMO (SU-MIMO) and transmit diversity for control signalling are introduced in Release 10 for LTE-Advanced, as explained in Section 29.4.

The physical layer transmissions of the LTE uplink comprise three physical channels and two signals:

- PRACH – Physical Random Access CHannel (see Chapter 17);
- PUSCH – Physical Uplink Shared CHannel (see Section 16.2);
- PUCCH – Physical Uplink Control CHannel (see Section 16.3);
- DM-RS – DeModulation Reference Signal (see Section 15.5);
- SRS – Sounding Reference Signal (see Section 15.6).

LTE – The UMTS Long Term Evolution: From Theory to Practice, Second Edition.

Stefania Sesia, Issam Toufik and Matthew Baker.

© 2011 John Wiley & Sons, Ltd. Published 2011 by John Wiley & Sons, Ltd.

The uplink physical channels, and their relationship to the higher-layer channels, are summarized in Figure 16.1.




Figure 16.1: Uplink physical channels and their mapping to higher layers.

16.2 Physical Uplink Shared Data Channel Structure

The Physical Uplink Shared CHannel (PUSCH), which carries data from the Uplink Shared Channel (UL-SCH) transport channel, uses DFT-Spread OFDM (DFT-S-OFDM), as described in Chapter 14. The transmit processing chain is shown in Figure 16.2. As explained in Chapter 10, the information bits are first channel-coded with a turbo code of mother code rate $r = 1/3$, which is adapted to a suitable final code rate by a rate-matching process. This is followed by symbol-level channel interleaving which follows a simple ‘time-first’ mapping [1] – in other words, adjacent data symbols end up being mapped first to adjacent SC-FDMA symbols in the time domain, and then across the subcarriers (see [2, Section 5.2.2.8]). The coded and interleaved bits are then scrambled by a length-31 Gold code (as described in Section 6.3) prior to modulation mapping, DFT-spreading, subcarrier mapping¹ and OFDM modulation. The signal is frequency-shifted by half a subcarrier prior to transmission, to avoid the distortion caused by the d.c. subcarrier being concentrated in one Resource Block (RB), as described in Section 14.3.3. The modulations supported are QPSK, 16QAM and 64QAM (the latter being only for the highest categories of User Equipment (UE) – Categories 5 and 8 (see Sections 1.3.4 and 27.5)).

¹Only localized mapping (i.e. to contiguous sets of Resource Blocks (RBs)) is supported for PUSCH and PUCCH transmissions in Releases 8 and 9. In Release 10, mapping to two clusters of RBs is also supported – see Section 28.3.6.2.




Figure 16.2: Uplink physical data channel processing.

The baseband SC-FDMA transmit signal for SC-FDMA symbol ℓ is given by the following expression (see [3, Section 5.6]),

$$s_\ell(t) = \sum_{k=-\lfloor N_{\text{RB}}^{\text{UL}} N_{\text{sc}}^{\text{RB}}/2 \rfloor}^{k=-\lceil N_{\text{RB}}^{\text{UL}} N_{\text{sc}}^{\text{RB}}/2 \rceil - 1} a_{k,\ell} \exp[j2\pi(k + 1/2)\Delta f(t - N_{\text{CP},\ell} T_s)] \quad (16.1)$$

for $0 \leq t < (N_{\text{CP},\ell} + N)T_s$, where $N_{\text{CP},\ell}$ is the number of samples of the Cyclic Prefix (CP) in SC-FDMA symbol ℓ (see Section 14.3), $N = 2048$ is the Inverse Fast Fourier Transform (IFFT) size, $\Delta f = 15$ kHz is the subcarrier spacing, $T_s = 1/(N \cdot \Delta f)$ is the sampling interval, $N_{\text{RB}}^{\text{UL}}$ is the uplink system bandwidth in RBs, $N_{\text{sc}}^{\text{RB}} = 12$ is the number of subcarriers per RB, $k^{(-)} = k + \lfloor N_{\text{RB}}^{\text{UL}} N_{\text{sc}}^{\text{RB}}/2 \rfloor$ and $a_{k,\ell}$ is the content of subcarrier k on symbol ℓ . For the PUSCH, the SC-FDMA symbol $a_{k,\ell}$ is obtained by DFT-spreading the QAM data symbols, $[d_{0,\ell}, d_{1,\ell}, \dots, d_{M_{\text{sc}}^{\text{PUSCH}-1,\ell}}]$ to be transmitted on SC-FDMA symbol ℓ (see [3, Section 5.3.3]),

$$a_{k,\ell} = \frac{1}{\sqrt{M_{\text{sc}}^{\text{PUSCH}}}} \sum_{i=0}^{M_{\text{sc}}^{\text{PUSCH}-1} - 1} d_{i,\ell} e^{-j2\pi i k / M_{\text{sc}}^{\text{PUSCH}}} \quad (16.2)$$

for $k = 0, 1, 2, \dots, M_{\text{sc}}^{\text{PUSCH}} - 1$, where $M_{\text{sc}}^{\text{PUSCH}} = M_{\text{RB}}^{\text{PUSCH}} \cdot N_{\text{sc}}^{\text{RB}}$ and $M_{\text{RB}}^{\text{PUSCH}}$ is the allocated PUSCH bandwidth in RBs.

As explained in Section 4.4.1, a Hybrid Automatic Repeat reQuest (HARQ) scheme is used, which in the uplink is synchronous, using N -channel stop and wait. This means that retransmissions occur in specific periodically occurring subframes (HARQ channels). Further details of the HARQ operation are given in Section 10.3.2.5.

16.2.1 Scheduling on PUSCH

In the LTE uplink, both frequency-selective scheduling and non-frequency-selective scheduling are supported. The former is based on the eNodeB exploiting available channel knowledge to schedule a UE to transmit using specific RBs in the frequency domain where the UEs experience good channel conditions. The latter does not make use of frequency-specific channel knowledge, but rather aims to benefit from frequency diversity during the transmission of each transport block. The possible techniques supported in LTE are discussed in more detail below. Intermediate approaches are also possible.

16.2.1.1 Frequency-Selective Scheduling

With frequency-selective scheduling, the same localized² allocation of transmission resources is typically used in both slots of a subframe – there is no frequency hopping during a subframe. The frequency-domain RB allocation and the Modulation and Coding Scheme (MCS) are chosen based on the location and quality of an above-average gain in the uplink channel response [4]. In order to enable frequency-selective scheduling, timely channel quality information is needed at the eNodeB. One method for obtaining such information in LTE is by uplink channel sounding using the SRS described in Section 15.6. The performance of frequency-selective scheduling using the SRS depends on the sounding bandwidth and the quality of the channel estimate, the latter being a function of the transmitted power spectral density used for the SRS. With a large sounding bandwidth, link quality can be evaluated on a larger number of RBs. However, this is likely to lead to the SRS being transmitted at a lower power density, due to the limited total UE transmit power, and this reduces the accuracy of the estimate for each RB within the sounding bandwidth especially for cell-edge UEs. Conversely, sounding a smaller bandwidth can improve channel estimation on the sounded RBs but results in missing channel information for certain parts of the channel bandwidth, thus risking exclusion of the best quality RBs. As an example, it is shown in [5] that, at least for a bandwidth of 5 MHz, frequency-selective scheduling based on full-band sounding outperforms narrower bandwidth sounding.

16.2.1.2 Frequency-Diverse or Non-Selective Scheduling

There are cases when no, or limited, frequency-specific channel quality information is available, for example because of SRS overhead constraints or high Doppler conditions. In such cases, it is preferable to exploit the frequency diversity of LTE's wideband channel.

In LTE, frequency hopping of a localized transmission is used to provide frequency-diversity.³ Two hopping modes are supported – hopping only between subframes (inter-subframe hopping) or hopping both between and within subframes (inter- and intra-subframe hopping). These modes are illustrated in Figure 16.3. Cell-specific broadcast signalling is used to configure the hopping mode via the parameter 'Hopping-mode' (see [6, Section 8.4]).

In case of intra-subframe hopping, a frequency hop occurs at the slot boundary in the middle of a subframe; this provides frequency diversity within a codeword (i.e. within a single transmission of transport block). On the other hand, inter-subframe hopping provides frequency diversity between HARQ retransmissions of a transport block, as the frequency allocation hops every allocated subframe.

Two methods are defined for the frequency hopping allocation (see [6, Section 8.4]): either a pre-determined pseudo-random frequency hopping pattern (see [3, Section 5.3.4]) or an explicit hopping offset signalled in the UL resource grant on the Physical Downlink Control Channel (PDCCH). For uplink system bandwidths less than 50 RBs, the size of the hopping offset (modulo the system bandwidth) is approximately half the number of RBs available for PUSCH transmissions (i.e. $\lfloor N_{RB}^{PUSCH}/2 \rfloor$), while for uplink system bandwidths of 50 RBs or more, the possible hopping offsets are $\lfloor N_{RB}^{PUSCH}/2 \rfloor$ and $\pm \lfloor N_{RB}^{PUSCH}/4 \rfloor$ (see [6, Section 8.4]).

²Localized means that allocated RBs are consecutive in the frequency domain.

³Note that there are upper limits on the size of resource allocation with which frequency hopping can be used – see [6, Section 8.4].




Figure 16.3: Uplink physical data channel processing.

Signalling the frequency hop via the uplink resource grant can be used for frequency semi-selective scheduling [7], in which the frequency resource is assigned selectively for the first slot of a subframe and frequency diversity is also achieved by hopping to a different frequency in the second slot. In some scenarios this may yield intermediate performance between that of fully frequency-selective and fully non-frequency-selective scheduling; this may be seen as one way to reduce the sounding overhead typically needed for fully frequency-selective scheduling.

In Release 10, another method of frequency-diverse scheduling is introduced, using dual-cluster PUSCH resource allocations; this is explained in detail in Section 28.3.6.2.

16.2.2 PUSCH Transport Block Sizes

The transport block size for a PUSCH data transmission is signalled in the corresponding resource grant on the PDCCH (DCI Format 0, or, for uplink SU-MIMO in Release 10, DCI Format 4 – see Section 9.3.5.1). Together with the indicated modulation scheme, the transport block size implies the code rate. The available transport block sizes are given in [6, Section 7.1.7.2].

In most cases, a generally linear range of code rates is available for each resource allocation size. One exception is an index which allows a transport block size of 328 bits in a single RB allocation with QPSK modulation, which corresponds to a code rate greater than unity. This is primarily designed to support cell-edge Voice-over-IP (VoIP) transmissions: by using only one RB per subframe, the UE’s power spectral density is maximized for good coverage; ‘TTI bundling’ (see Section 14.3.2), whereby the transmission is repeated in four

consecutive subframes, together with typically three retransmissions at 16 ms intervals, then enables both Chase combining gain and Incremental Redundancy (IR) gain to be achieved.

16.3 Uplink Control Channel Design

In general, uplink control signalling in mobile communications systems can be divided into two categories:

- **Data-associated control signalling** is control signalling which is always transmitted together with uplink data and is used in the processing of that data. Examples include transport format indications, ‘New Data’ Indicators (NDIs) and MIMO parameters.
- **Control signalling not associated with data** is transmitted independently of any uplink data packet. Examples include HARQ Acknowledgements (ACK/NACK) for downlink data packets, Channel Quality Indicators (CQIs) to support link adaptation (see Section 10.2), and MIMO feedback such as Rank Indicators (RIs) and Pre-coding Matrix Indicators (PMIs) (see Section 11.2.2.4) for downlink transmissions. Scheduling Requests (SRs) for uplink transmissions also fall into this category (see Section 4.4.2.2).

In LTE, the low signalling latency afforded by the short subframe duration of 1 ms, together with the orthogonal nature of the uplink multiple access scheme which necessitates centralized resource allocation, make it appropriate for the eNodeB to be in full control of the uplink transmission parameters. Consequently uplink data-associated control signalling is not necessary in LTE, as the relevant information is already known to the eNodeB. Therefore only data-non-associated control signalling exists in the LTE uplink.

When simultaneous uplink PUSCH data and control signalling are scheduled, the control signalling is normally multiplexed together with the data prior to the DFT spreading⁴, in order to preserve the single-carrier low Cubic Metric (CM) property of the uplink transmission. The uplink control channel, PUCCH, is used by a UE to transmit any necessary control signalling only in subframes in which the UE has not been allocated any RBs for PUSCH transmission. In the design of the PUCCH, special consideration was given to maintaining a low CM [8].

16.3.1 Physical Uplink Control Channel (PUCCH) Structure

The control signalling on the PUCCH is transmitted in a frequency region that is normally configured to be on the edges of the system bandwidth.

In order to minimize the resources needed for transmission of control signalling, the PUCCH in LTE is designed to exploit frequency diversity: each PUCCH transmission in one subframe comprises a single (0.5 ms) RB at or near one edge of the system bandwidth, followed (in the second slot of the subframe) by a second RB at or near the opposite edge of the system bandwidth, as shown in Figure 16.4; together, the two RBs are referred to as a *PUCCH region*. This design can achieve a frequency diversity benefit of approximately 2 dB compared to transmission in the same RB throughout the subframe.

⁴In Release 10, the possibility of simultaneous transmission of PUSCH and PUCCH is introduced; this is explained in Section 28.3.2.




Figure 16.4: A PUCCH region.

At the same time, the narrow bandwidth of the PUCCH in each slot (only a single RB) maximizes the power per subcarrier for a given total transmission power (see Figure 16.5) and therefore helps to fulfil stringent coverage requirements.




Figure 16.5: The link budget of a two-slot narrowband transmission exceeds that of a one-slot wider-band transmission, given equal coding gain.

Positioning the control regions at the edges of the system bandwidth has a number of advantages, including the following:

- The frequency diversity achieved through frequency hopping is maximized by allowing hopping from one edge of the band to the other.
- Out-Of-Band (OOB) emissions are smaller if a UE is only transmitting on a single RB per slot compared to multiple RBs. The PUCCH regions can serve as a kind of guard band between the wider-bandwidth PUSCH transmissions of adjacent carriers and can therefore improve coexistence [9].
- Using control regions on the band edges maximizes the achievable PUSCH data rate, as the entire central portion of the band can be allocated to a single UE. If the control regions were in the central portion of a carrier, a UE bandwidth allocation would be

limited to one side of the control region in order to maintain the single-carrier nature of the signal, thus limiting the maximum achievable data rate.

- Control regions on the band edges impose fewer constraints on the uplink data scheduling, both with and without inter-/intra-subframe frequency hopping.

The number of RBs (in each slot) that can be used for PUCCH transmission within the cell is N_{RB}^{HO} (parameter ‘pusch-HoppingOffset’). This is indicated to the UEs in the cell through broadcast signalling. Note that the number of PUCCH RBs per slot is the same as the number of PUCCH regions per subframe. Some typical expected numbers of PUCCH regions for different LTE bandwidths are shown in Table 16.1.

Table 16.1: Typical numbers of PUCCH regions.

Bandwidth (MHz)	Number of RBs per subframe	Number of PUCCH regions
1.4	2	1
3	4	2
5	8	4
10	16	8
20	32	16

Figures 16.6 and 16.7 respectively show examples of even and odd numbers of PUCCH regions being configured in a cell.

In the case of an even number of PUCCH regions (Figure 16.6), both RBs of each RB-pair (e.g. RB-pair 2 and RB-pair $N_{RB}^{UL} - 3$) are used for PUCCH transmission. However, for the case of an odd number of PUCCH regions (Figure 16.7), one RB of an RB-pair in each slot is not used for PUCCH (e.g. one RB of RB-pair 2 and RB-pair $N_{RB}^{UL} - 3$ is unused); the eNodeB

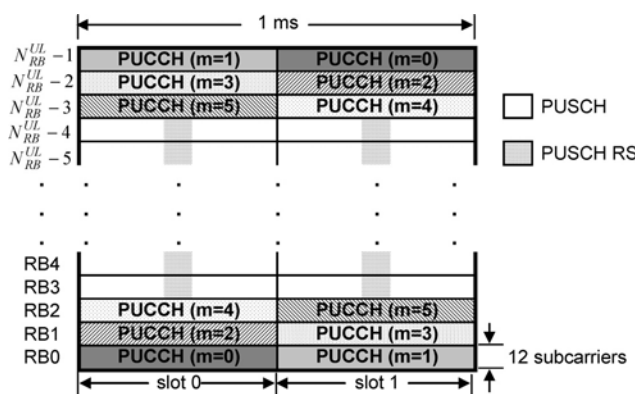


Figure 16.6: PUCCH uplink control structure with an even number of ‘PUCCH Control Regions’ ($N_{RB}^{PUCCH} = 6$).




Figure 16.7: Example of an odd number of PUCCH regions ($N_{RB}^{HO} = 5$).

may schedule a UE with an intra-subframe frequency hopping PUSCH allocation to exploit these unused RBs.

Alternatively, a UE can be assigned a localized allocation which includes the unused RB-pair (e.g. RB-pair 2 or RB-pair $N_{RB}^{UL} - 3$). In this case, the UE will transmit PUSCH data on both RBs of the RB-pair, assuming that neither of the RBs are used for PUCCH by any UE in the subframe. Thus, the eNodeB scheduler can appropriately schedule PUSCH transmission on the PUCCH RBs when they are under-utilized.

The eNodeB may also choose to schedule low-power PUSCH transmission (e.g. from UEs close to the eNodeB) in the outer RBs of the configured PUCCH region, while the inner PUCCH region is used for PUCCH signalling. This can provide further reduction in OOB emissions which is necessary in some frequency bands, by moving higher-power PUCCH transmission (e.g. those from cell-edge UEs) slightly away from the edge of the band.

16.3.1.1 Multiplexing of UEs within a PUCCH Region

Control signalling from multiple UEs can be multiplexed into a single PUCCH region using orthogonal Code Division Multiplexing (CDM). In some scenarios this can have benefits over a pure Frequency Division Multiplexing (FDM) approach, as it reduces the need to limit the power differentials between the PUCCH transmissions of different UEs. One technique to provide orthogonality between UEs is by using cyclic time shifts of a sequence with suitable properties, as explained in Section 15.2.2. In a given SC-FDMA symbol, different cyclic time shifts of a waveform (e.g. a Zadoff–Chu (ZC) sequence as explained in Section 7.2.1) are modulated with a UE-specific QAM symbol carrying the necessary control signalling information, with the supported number of cyclic time shifts determining the number of UEs which can be multiplexed per SC-FDMA symbol. As the PUCCH RB spans 12 subcarriers, and assuming the channel is approximately constant over the RB (i.e. a single-tap channel), the LTE PUCCH supports up to 12 cyclic shifts per PUCCH RB.

For control information transmissions with a small number of control signalling bits, such as 1- or 2-bit positive/negative acknowledgements (ACK/NACK), orthogonality is achieved

between UEs by a combination of cyclic time shifts within an SC-FDMA symbol and SC-FDMA symbol time-domain spreading with orthogonal spreading codes, i.e. modulating the SC-FDMA symbols by elements of an orthogonal spreading code [10]. CDM of multiple UEs is used rather than Time Domain Multiplexing (TDM) because CDM enables the time duration of the transmission to be longer, which increases the total transmitted energy per signalling message in the case of a power-limited UE.

Thus, the LTE PUCCH control structure uses frequency-domain code multiplexing (different cyclic time shifts of a base sequence) and/or time-domain code multiplexing (different orthogonal block spreading codes), thereby providing an efficient, orthogonal control channel which supports small payloads (up to 22 coded bits) from multiple UEs simultaneously, together with good operational capability at low SNR.

16.3.2 Types of Control Signalling Information and PUCCH Formats

The Uplink Control Information (UCI) can consist of:

- Scheduling Requests (SRs) (see Section 4.4.2.2).
- HARQ ACK/NACK in response to downlink data packets on the Physical Downlink Shared CHannel(PDSCH); one ACK/NACK bit is transmitted in the case of single-codeword downlink transmission while two ACK/NACK bits are used in the case of two-codeword downlink transmission.
- Channel State Information (CSI), which includes Channel Quality Indicators (CQIs) as well as the MIMO-related feedback consisting of RIs and PMI. 20 bits per subframe are used for the CSI.

The amount of UCI a UE can transmit in a subframe depends on the number of SC-FDMA symbols available for transmission of control signalling data (i.e. excluding SC-FDMA symbols used for RS transmission for coherent detection of the PUCCH). The PUCCH supports seven different formats (with an eighth added in Release 10), depending on the information to be signalled. The mapping between PUCCH formats and UCI is shown in Table 16.2 (see [6, Section 10.1] and [3, Table 5.4-1]).

Table 16.2: Supported uplink control information formats on PUCCH.

PUCCH Format	Uplink Control Information (UCI)
Format 1	Scheduling request (SR) (unmodulated waveform)
Format 1a	1-bit HARQ ACK/NACK with/without SR
Format 1b	2-bit HARQ ACK/NACK with/without SR
Format 2	CSI (20 coded bits)
Format 2	CSI and 1- or 2-bit HARQ ACK/NACK for extended CP only
Format 2a	CSI and 1-bit HARQ ACK/NACK (20 + 1 coded bits)
Format 2b	CSI and 2-bit HARQ ACK/NACK (20 + 2 coded bits)
Format 3	Multiple ACK/NACKs for carrier aggregation: up to 20 ACK/NACK bits plus optional SR, in 48 coded bits; see Section 28.3.2.1 for details.

The physical mapping of the PUCCH formats to the PUCCH regions is shown in Figure 16.8.




Figure 16.8: Physical mapping of PUCCH formats to PUCCH RBs or regions.

It can be seen that the PUCCH CSI formats $2/2a/2b$ are mapped and transmitted on the band-edge RBs (e.g. PUCCH region $m = 0, 1$) followed by a mixed PUCCH RB (if present, e.g. region $m = 2$) of CSI format $2/2a/2b$ and SR/HARQ ACK/NACK format $1/1a/1b$, and then by PUCCH SR/HARQ ACK/NACK format $1/1a/1b$ (e.g. region $m = 4, 5$). The number of PUCCH RBs available for use by CSI format $2/2a/2b$, N_{RB}^{UL} , is indicated to the UEs in the cell by broadcast signalling.

16.3.3 Channel State Information Transmission on PUCCH (Format 2)

The PUCCH CSI channel structure (Format 2) for one slot with normal CP is shown in Figure 16.9. SC-FDMA symbols 1 and 5 are used for DM-RS transmissions in the case of the normal CP (while in the case of the extended CP only one RS is transmitted, on SC-FDMA symbol 3).

The number of RS symbols per slot results from a trade-off between channel estimation accuracy and the supportable code rate for the UCI bits. For a small number of UCI bits with a low SNR operating point (for a typical 1% target error rate), improving the channel estimation accuracy by using more RS symbols is more beneficial than being able to use a lower channel code rate. However, with larger numbers of UCI bits the required SNR operating point increases, and the higher code rate resulting from a larger overhead of RS symbols becomes more critical, thus favouring fewer RS symbols. In view of these factors, two RS symbols per slot (in case of normal CP) was considered to provide the best trade-off in terms of performance and RS overhead, given the payload sizes required.

10 CSI bits are channel coded with a rate $1/2$ punctured $(20, k)$ Reed-Muller code (see [2, Section 5.2.3.3]) to give 20 coded bits, which are then scrambled (in a similar

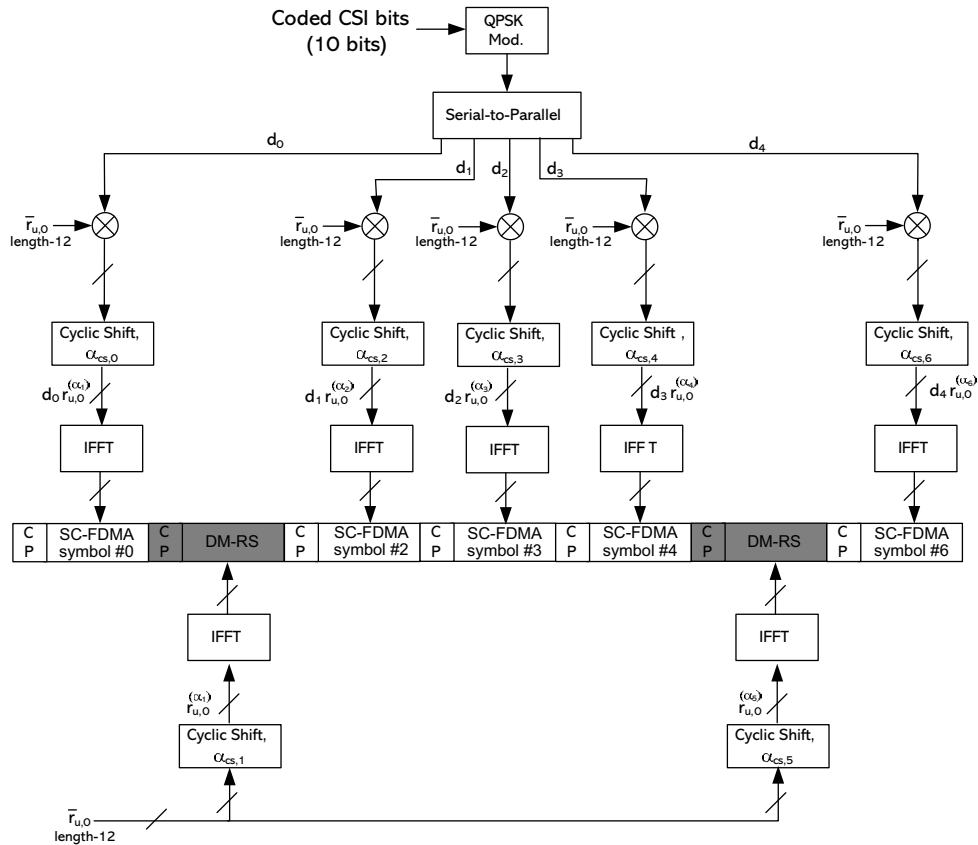


Figure 16.9: CSI channel structure for PUCCH format 2/2a/2b with normal CP for one slot.

way to PUSCH data with a length-31 Gold sequence) prior to QPSK constellation mapping. One QPSK modulated symbol is transmitted on each of the 10 SC-FDMA symbols in the subframe by modulating a cyclic time shift of the base RS sequence of length 12 prior to OFDM modulation. The 12 equally spaced cyclic time shifts allow 12 different UEs to be orthogonally multiplexed on the same PUCCH format 2 RB.

The DM-RS signal sequence (on the 2nd and 6th SC-FDMA symbols for the normal CP and the 4th symbol for the extended CP) is similar to the frequency-domain CSI signal sequence but without the CSI data modulation.

In order to provide inter-cell interference randomization, cell-specific symbol-level cyclic time-shift hopping is used, as described in Section 15.4. For example, the PUCCH cyclic time-shift index on SC-FDMA symbol l in even slots n_s is obtained by adding (modulo-12) a pseudo-random cell-specific PUCCH cyclic shift offset to the assigned cyclic time shift n_{RS}^{PUCCH} . Intra-cell interference randomization is achieved by cyclic time-shift remapping in the second slot as explained in Section 15.4.

A UE is semi-statically configured by higher layer signalling to report periodically different CQI, PMI, and RI types (see Section 10.2.1) on the CSI PUCCH using a PUCCH resource index $n_{\text{PUCCH}}^{(2)}$, which indicates both the PUCCH region and the cyclic time shift to be used. The PUCCH region m used for the PUCCH format 2/2a/2b transmission (see Figure 16.8) is given by (see [3, Section 5.4.3])

$$m = \left\lfloor \frac{n_{\text{PUCCH}}^{(2)}}{12} \right\rfloor \quad (16.3)$$

and the assigned cyclic time shift, $n_{\text{RS}}^{\text{PUCCH}}$, is given by

$$n_{\text{RS}}^{\text{PUCCH}} = n_{\text{PUCCH}}^{(2)} \bmod 12. \quad (16.4)$$

16.3.4 Multiplexing of CSI and HARQ ACK/NACK from a UE on PUCCH

The simultaneous transmission of HARQ ACK/NACK and CSI on the PUCCH can be enabled by UE-specific higher layer signalling.⁵ If simultaneous transmission is not enabled and the UE needs to transmit HARQ ACK/NACK on the PUCCH in the same subframe in which a CSI report has been configured, the CSI is dropped and only HARQ ACK/NACK is transmitted using the transmission structure detailed in Section 16.3.5.

In subframes where the eNodeB scheduler allows for simultaneous transmission of CSI and HARQ ACK/NACK on the PUCCH from a UE, the CSI and the 1- or 2-bit HARQ ACK/NACK information needs to be multiplexed in the same PUCCH RB, while maintaining the low CM single-carrier property of the signal. The method used to achieve this is different for the cases of normal CP and extended CP, as described in the following sections.

16.3.4.1 Normal CP (Format 2a/2b)

The transmission structure for CSI is the same as described in Section 16.3.3. In order to transmit a 1- or 2-bit HARQ ACK/NACK together with CSI, the HARQ ACK/NACK bits (which are not scrambled) are BPSK/QPSK modulated as shown in Figure 16.10, resulting in a single HARQ ACK/NACK modulation symbol, d_{HARQ} . A positive acknowledgement (ACK) is encoded as a binary ‘1’ and a negative acknowledgement (NACK) is encoded as a binary ‘0’ (see [2, Section 5.2.3.4]).

The single HARQ ACK/NACK modulation symbol, d_{HARQ} , is then used to modulate the second RS symbol (SC-FDMA symbol 5) in each slot – i.e. ACK/NACK is signalled using the RS. This results in PUCCH formats 2a/2b. It can be seen from Figure 16.10 that the modulation mapping is such that a NACK (or NACK, NACK in the case of two downlink MIMO codewords) is mapped to +1, resulting in a default NACK in case neither ACK nor NACK is transmitted (so-called Discontinuous Transmission (DTX)), as happens if the UE fails to detect the downlink grant on the PDCCH. In other words, a DTX (no RS modulation) is interpreted as a NACK by the eNodeB, triggering a downlink retransmission.

As one of the RSs in a slot is modulated by the HARQ ACK/NACK symbol, a variety of ACK/NACK and CSI detection schemes are possible. In low-Doppler environments with

⁵Note that Release 10 also provides the option of configuring simultaneous PUCCH and PUSCH transmission from a UE – see Section 28.3.2.

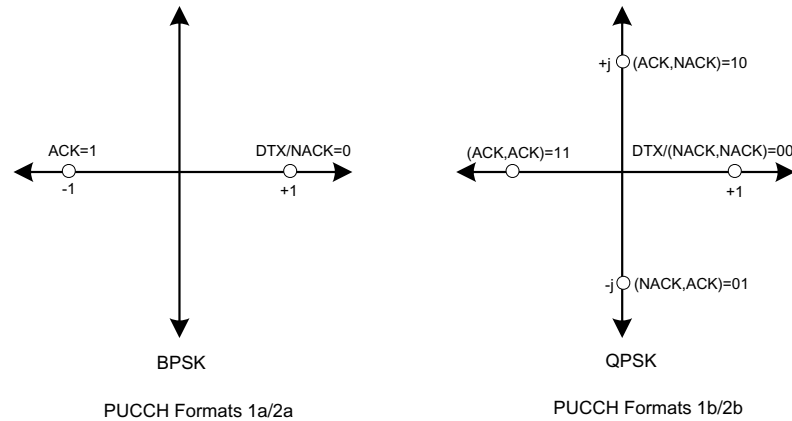


Figure 16.10: Constellation mapping for HARQ ACK/NACK.

little channel variation over the 0.5 ms slot, coherent detection of ACK/NACK and CSI can be achieved by using only the first RS symbol in the slot as the phase reference. Alternatively, to improve the channel estimation quality for CSI detection, an estimate of the HARQ ACK/NACK symbol can be used to undo the modulation on the second RS in the slot so that both RS symbols can be used for channel estimation and demodulation of CSI. In high-Doppler environments in which significant channel variations occur over a slot, relying on a single RS symbol for coherent detection degrades the performance of the ACK/NACK and CSI demodulation. In such cases, blind decoding or multiple hypothesis testing of the different ACK/NACK combinations can be used to decode the ACK/NACK and CSI, selecting the hypothesis that maximizes the correlation between the received signal and the estimated CSI [11] (i.e. a Maximum Likelihood detection).

16.3.4.2 Extended CP (Format 2)

In the case of the extended CP (with one RS symbol per slot), the 1- or 2-bit HARQ ACK/NACK is jointly encoded with the CSI resulting in a $(20, k_{\text{CQI}} + k_{\text{ACK/NACK}})$ Reed-Muller-based block code. A 20-bit codeword is transmitted on the PUCCH using the channel structure described in Section 16.3.3. The joint coding of the ACK/NACK and CSI is performed as shown in Figure 16.11. The largest number of information bits supported by the block code is 13, corresponding to $k_{\text{CSI}} = 11$ CSI bits and $k_{\text{ACK/NACK}} = 2$ bits (for two-codeword transmission in the downlink).

16.3.5 HARQ ACK/NACK Transmission on PUCCH (Format 1a/1b)

The PUCCH channel structure for HARQ ACK/NACK transmission with no CSI is shown in Figure 16.12 for one slot with normal CP. Three (two in case of extended CP) SC-FDMA symbols are used in the middle of the slot for RS transmission, with the remaining four SC-FDMA symbols being used for ACK/NACK transmission. Due to the small number of ACK/NACK bits, three RS symbols are used to improve the channel estimation accuracy for a lower SNR operating point than for the CSI structure described in Section 16.3.3.




Figure 16.11: Joint coding of HARQ ACK/NACK and CSI for extended CP.

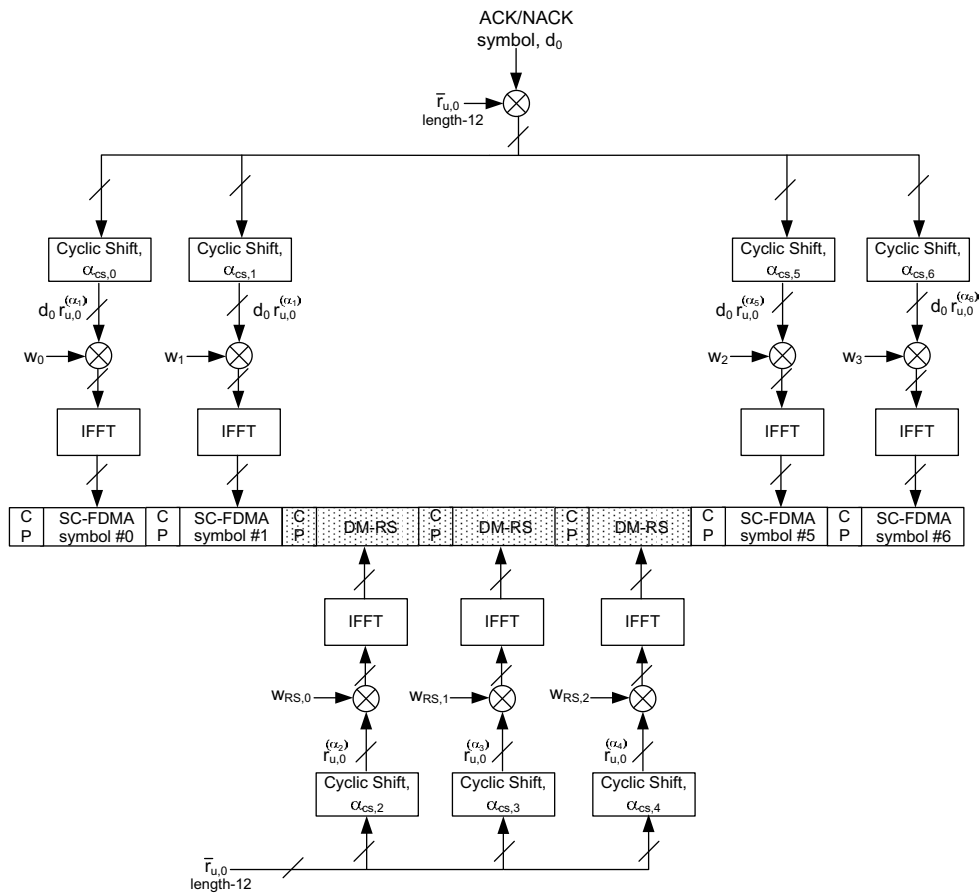


Figure 16.12: ACK/NACK structure – users are multiplexed using different cyclic shifts and time-domain spreading.

Both 1- and 2-bit acknowledgements are supported using BPSK and QPSK modulation respectively. The HARQ ACK/NACK bits are BPSK/QPSK modulated according to the modulation mapping shown in Figure 16.10 (see [3, Table 5.4.1-1]) resulting in a single HARQ ACK/NACK modulation symbol. An ACK is encoded as a binary ‘1’ and a NACK as a binary ‘0’ (see [2, Section 5.2.3.4]). The modulation mapping is the same as the mapping for 1- or 2-bit HARQ ACK/NACK when multiplexed with CSI for PUCCH formats 2a/2b. The modulation symbol is scrambled on a per-slot basis by either 1 or $e^{(-j\pi/2)}$ depending on the PUCCH resource index.

As in the case of CSI transmission, the one BPSK/QPSK modulated symbol (which is phase-rotated by 90 degrees in the second slot) is transmitted on each SC-FDMA data symbol by modulating a cyclic time shift of the base RS sequence of length 12 (i.e. frequency-domain CDM) prior to OFDM modulation. In addition, as mentioned in Section 16.3.1.1, time-domain spreading with orthogonal (Walsh–Hadamard or DFT) spreading codes is used to code-division-multiplex UEs. Thus, a large number of UEs (data and RSs) can be multiplexed on the same PUCCH RB using frequency-domain and time-domain code multiplexing. The RSs from the different UEs are multiplexed in the same way as the data SC-FDMA symbols.

For the cyclic time-shift multiplexing, the number of cyclic time shifts supported in an SC-FDMA symbol for PUCCH HARQ ACK/NACK RBs is configurable by a cell-specific higher-layer signalling parameter $\Delta_{\text{shift}}^{\text{PUCCH}} \in \{1, 2, 3\}$, indicating 12, 6, or 4 shifts respectively (see [3, Section 5.4.1]). The value selected by the eNodeB for $\Delta_{\text{shift}}^{\text{PUCCH}}$ can be based on the expected delay spread in the cell.

For the time-domain spreading CDM, the length-2 and length-4 orthogonal block spreading codes are based on Walsh–Hadamard codes, and the length-3 spreading codes are based on DFT codes as shown in Table 16.3. The number of supported spreading codes is limited by the number of RS symbols, as the multiplexing capacity of the RSs is smaller than that of the data symbols due to the smaller number of RS symbols. Therefore a subset of size s orthogonal spreading codes of a particular length L ($s \leq L$) is used depending on the number of RS SC-FDMA symbols. For the normal CP with four data SC-FDMA symbols and three supportable orthogonal time spreading codes (due to there being three RS symbols), the indices 0, 1, 2 of the length-4 orthogonal spreading codes are used for the data time-domain block spreading.

Table 16.3: Time-domain orthogonal spreading code sequences.
Reproduced by permission of © 3GPP.

Orthogonal code sequence index	Length 2	Length 3	Length 4
	Walsh–Hadamard	DFT	Walsh–Hadamard
0	$[+1 \ +1]$	$[+1 \ +1 \ +1]$	$[+1 \ +1 \ +1 \ +1]$
1	$[+1 \ -1]$	$[1 \ e^{j2\pi/3} \ e^{j4\pi/3}]$	$[+1 \ -1 \ +1 \ -1]$
2	N/A	$[1 \ e^{j4\pi/3} \ e^{j2\pi/3}]$	$[+1 \ -1 \ -1 \ +1]$
3	N/A	N/A	$[+1 \ +1 \ -1 \ -1]$

Similarly, for the extended CP case with four data SC-FDMA symbols but only two RS symbols, orthogonal spreading code indices 0 and 2 of length 4 are used for the data block spreading codes. For the length-4 orthogonal codes, the code sequences used are such that subsets of the code sequences result in the minimum inter-code interference in high-Doppler conditions where generally the orthogonality between the code sequences breaks down [12]. Table 16.4 summarizes the time-domain orthogonal spreading code lengths (i.e. spreading factors) for data and RS. The number of supportable orthogonal spreading codes is equal to the number of RS SC-FDMA symbols, N_{RS}^{PUCCH} .

Table 16.4: Spreading factors for time-domain orthogonal spreading codes for data and RS for PUCCH formats 1/1a/1b for normal and extended CP.

	Normal CP		Extended CP	
	Data, N_{SF}^{PUCCH}	RS, N_{RS}^{PUCCH}	Data, N_{SF}^{PUCCH}	RS, N_{RS}^{PUCCH}
Spreading factor	4	3	4	2

It should be noted that it is possible for the transmission of HARQ ACK/NACK and SRS to be configured in the same subframe. If this occurs, the eNodeB can also configure (by cell-specific broadcast signalling) the way in which these transmissions are to be handled by the UE. One option is for the ACK/NACK to take precedence over the SRS, such that the SRS is not transmitted and only HARQ ACK/NACK is transmitted in the relevant subframe, according to the PUCCH ACK/NACK structure in Figure 16.12. The alternative is for the eNodeB to configure the UEs to use a shortened PUCCH transmission in such subframes, whereby the last SC-FDMA symbol of the ACK/NACK (i.e. the last SC-FDMA symbol in the second slot of the subframe is not transmitted; this is shown in Figure 16.13).

This maintains the low CM single-carrier property of the transmitted signal, by ensuring that a UE never needs to transmit both HARQ ACK/NACK and SRS symbols simultaneously, even if both signals are configured in the same subframe. If the last symbol of the ACK/NACK is not transmitted in the second slot of the subframe, this is known as a *shortened* PUCCH format, as shown in Figure 16.14.⁶ For the shortened PUCCH, the length of the time-domain orthogonal block spreading code is reduced by one (compared to the first slot shown in Figure 16.12). Hence, it uses the length-3 DFT basis spreading codes in Table 16.3 in place of the length-4 Walsh–Hadamard codes.

The frequency-domain HARQ ACK/NACK signal sequence on SC-FDMA symbol n is defined in [3, Section 5.4.1].

The number of HARQ ACK/NACK resource indices $N_{PUCCH, RB}^{(1)}$ corresponding to cyclic-time-shift/orthogonal-code combinations that can be supported in a PUCCH RB is given by

$$N_{PUCCH, RB}^{(1)} = c \cdot P, \quad c = \begin{cases} 3 & \text{normal cyclic prefix} \\ 2 & \text{extended cyclic prefix} \end{cases} \quad (16.5)$$

⁶Note that configuration of SRS in the same subframe as channel quality information or SR is not valid. Therefore the shortened PUCCH formats are only applicable for PUCCH formats 1a and 1b.




Figure 16.13: A UE may not simultaneously transmit on SRS and PUCCH or PUSCH, in order to avoid violating the single-carrier nature of the signal. Therefore, a PUCCH or PUSCH symbol may be punctured if SRS is transmitted.

where c is equal to the number of RS symbols, $P = 12/\Delta_{\text{shift}}^{\text{PUCCH}}$, and $\Delta_{\text{shift}}^{\text{PUCCH}} \in \{1, 2, 3\}$ is the number of equally spaced cyclic time shifts supported. For example, with the normal CP and $\Delta_{\text{shift}}^{\text{PUCCH}} = 2$, ACK/NACKs from 18 different UEs can be multiplexed in one RB.

As in the case of CSI (see Section 16.3.3), cyclic time shift hopping (described in Section 15.4) is used to provide inter-cell interference randomization.

In the case of semi-persistently scheduled downlink data transmissions on the PDSCH (see Section 4.4.2.1) without a corresponding downlink grant on the control channel PDCCH, the PUCCH ACK/NACK resource index $n_{\text{PUCCH}}^{(1)}$ to be used by a UE is semi-statically configured by higher layer signalling. This PUCCH ACK/NACK resource is used for ACK/NACK transmission corresponding to initial HARQ transmission. For dynamically-scheduled downlink data transmissions (including HARQ retransmissions for semi-persistent data) on the PDSCH (indicated by downlink assignment signalling on the PDCCH), the PUCCH HARQ ACK/NACK resource index $n_{\text{PUCCH}}^{(1)}$ is implicitly determined based on the index of the first Control Channel Element (CCE, see Section 9.3) of the PDCCH message.

The PUCCH region m used for the HARQ ACK/NACK with format 1/1a/1b transmission for the case with no mixed PUCCH region (shown in Figure 16.8), is given by

$$m = \left\lfloor \frac{n_{\text{PUCCH}}^{(1)}}{N_{\text{PUCCH, RB}}^{(1)}} \right\rfloor + N_{\text{RB}}^{(2)} \quad (16.6)$$

where $N_{\text{RB}}^{(2)}$ is the number of RBs that are available for PUCCH formats 2/2a/2b and is a cell-specific broadcast parameter (see [3, Section 5.4.3]).

The PUCCH resource index $n^{(1)}(n_s)$, corresponding to a combination of a cyclic time shift and orthogonal code ($n_{\text{RS}}^{\text{PUCCH}}$ and n_{oc}), within the PUCCH region m in even slots is given by

$$n^{(1)}(n_s) = n_{\text{PUCCH}}^{(1)} \bmod N_{\text{PUCCH, RB}}^{(1)} \quad \text{for } n_s \bmod 2 = 0 \quad (16.7)$$

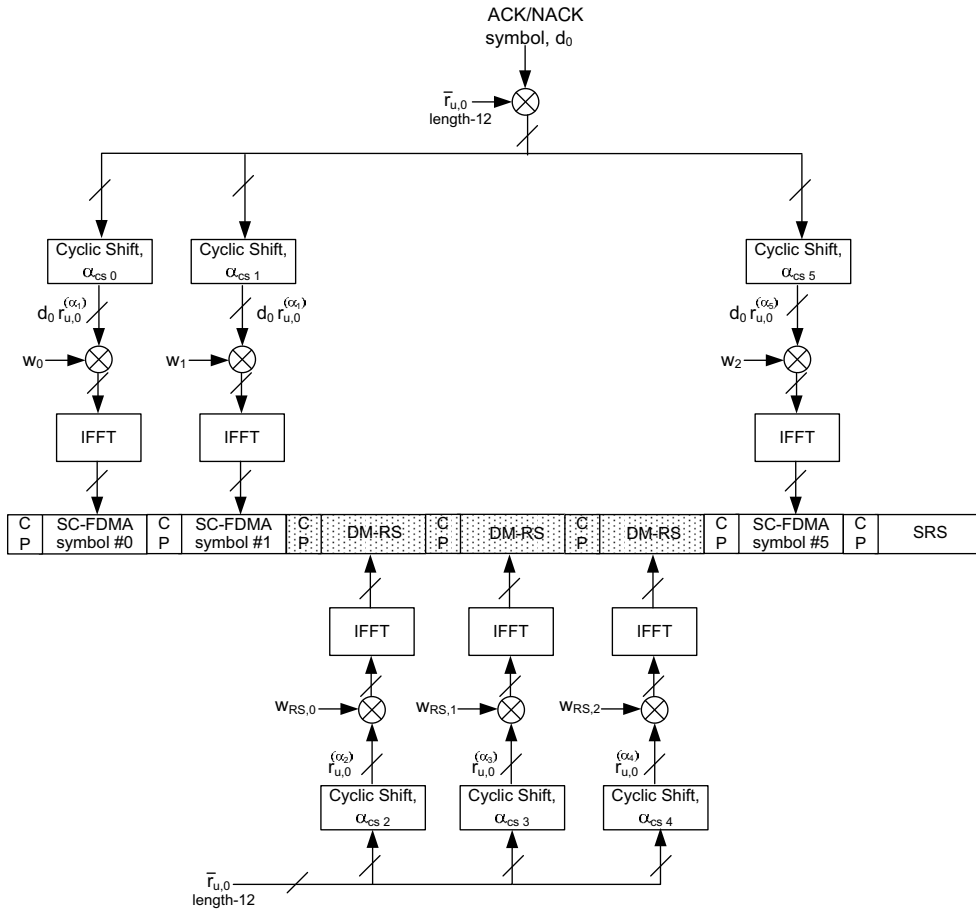


Figure 16.14: Shortened PUCCH ACK/NACK structure when simultaneous SRS and ACK/NACK is enabled in the cell.

The PUCCH resource index (n_{RS}^{PUCCH}, n_{oc}) allocation within a PUCCH RB format 1/1a/1b, is shown in Tables 16.5, 16.6 and 16.7, for $\Delta_{shift}^{PUCCH} \in \{1, 2, 3\}$ with 36, 18, and 12 resource indices respectively for the normal CP case [13]. For the extended CP, with two time-domain orthogonal spreading code sequences, only the first two columns of the orthogonal code sequence index, $n_{oc} = 1, 2$ are used, resulting in 24, 12 and 8 resource indices for $\Delta_{shift}^{PUCCH} \in \{1, 2, 3\}$ respectively.

The PUCCH resources are first indexed in the cyclic time-shift domain, followed by the orthogonal time spreading code domain.

The cyclic time shifts used on *adjacent* orthogonal codes can also be staggered, providing the opportunity to separate the channel estimates prior to de-spreading. As high-Doppler breaks down the orthogonality between the spread blocks, offsetting the cyclic time shift

Table 16.5: PUCCH RB format 1/1a/1b resource index allocation,
 $\Delta_{\text{shift}}^{\text{PUCCH}} = 1$, 36 resource indices, normal CP.

Cyclic shift index, $n_{\text{RS}}^{\text{PUCCH}}$	Orthogonal code sequence index, n_{oc}		
	$n_{\text{oc}} = 0$	$n_{\text{oc}} = 1$	$n_{\text{oc}} = 2$
0	0	12	24
1	1	13	25
2	2	14	26
3	3	15	27
4	4	16	28
5	5	17	29
6	6	18	30
7	7	19	31
8	8	20	32
9	9	21	33
10	10	22	34
11	11	23	35

Table 16.6: PUCCH RB format 1/1a/1b resource index allocation,
 $\Delta_{\text{shift}}^{\text{PUCCH}} = 2$, 18 resource indices, normal CP.

Cyclic shift index, $n_{\text{RS}}^{\text{PUCCH}}$	Orthogonal code sequence index, n_{oc}		
	$n_{\text{oc}} = 0$	$n_{\text{oc}} = 1$	$n_{\text{oc}} = 2$
0	0		12
1		6	
2	1		13
3		7	
4	2		14
5		8	
6	3		15
7		9	
8	4		16
9		10	
10	5		17
11		11	

values within each SC-FDMA symbol can restore orthogonality at moderate delay spreads. This can enhance the tracking of high-Doppler channels [14].

In order to randomize intra-cell interference, PUCCH resource index remapping is used in the second slot [15]. Index remapping includes both cyclic shift remapping and orthogonal block spreading code remapping (similar to the case of CSI – see Section 16.3.3).

The PUCCH resource index remapping function in an odd slot is based on the PUCCH resource index in the even slot of the subframe, as defined in [3, Section 5.4.1].

Table 16.7: PUCCH RB format 1/1a/1b resource index allocation, $\Delta_{\text{shift}}^{\text{PUCCH}} = 3, 12$ resource indices, normal CP.

Cyclic shift index, $n_{\text{RS}}^{\text{PUCCH}}$	Orthogonal code sequence index, n_{oc}		
	$n_{\text{oc}} = 0$	$n_{\text{oc}} = 1$	$n_{\text{oc}} = 2$
0	0		
1		4	
2			7
3	1		
4		5	
5			8
6	2		
7			
8			
9	3		
10		6	
11			9

16.3.6 Multiplexing of CSI and HARQ ACK/NACK in the Same (Mixed) PUCCH RB

The multiplexing of CSI and HARQ ACK/NACK in different PUCCH RBs can in general simplify the system. However, in the case of small system bandwidths, such as 1.4 MHz, the control signalling overhead can become undesirably high with separate CSI and ACK/NACK RB allocations (two out of a total of 6 RBs for control signalling in 1.4 MHz system bandwidths). Therefore, multiplexing of CSI and ACK/NACK from different UEs in the same mixed PUCCH RB is supported in LTE to reduce the total control signalling overhead.

The ZC cyclic time-shift structure facilitates the orthogonal multiplexing of channel quality information and ACK/NACK signals with different numbers of RS symbols. This is achieved by assigning different sets of adjacent cyclic time shifts to CSI and ACK/NACK signals [16] as shown in Table 16.8. As can be seen from this table, $N_{\text{cs}}^{(1)} \in \{0, 1, \dots, 7\}$ cyclic time shifts are used for PUCCH ACK/NACK formats 1/1a/1b in the mixed PUCCH RB case, where $N_{\text{cs}}^{(1)}$ is a cell-specific broadcast parameter (see [3, Section 5.4]) restricted to integer multiples of $\Delta_{\text{shift}}^{\text{PUCCH}}$. A guard cyclic time shift is used between the ACK/NACK and CSI cyclic shift resources to improve orthogonality and channel separation between UEs transmitting CSI and those transmitting ACK/NACK. To avoid mixing of the cyclic time shifts for ACK/NACK and CSI, the cyclic time shift (i.e. the PUCCH resource index remapping function) for the odd slot of the subframe is not used; the same cyclic time shift as in the first slot of the subframe is used.

16.3.7 Scheduling Request (SR) Transmission on PUCCH (Format 1)

The structure of the SR PUCCH format 1 is the same as that of the ACK/NACK PUCCH format 1a/1b explained in Section 16.3.5, where a cyclic time shift of the base RS sequence is modulated with time-domain orthogonal block spreading. The SR uses simple On–Off

Table 16.8: Multiplexing of ACK/NACK (format 1/1a/1b) and CSI (format 2/2a/2b) from different UEs in the same (mixed) PUCCH RB by using different sets of cyclic time shifts.

Cyclic shift index	Cyclic shift index allocation
0	Format 1/1a/1b (HARQ ACK/NACK, SR) cyclic shifts
1	
2	
⋮	
$N_{cs}^{(1)}$	
$N_{cs}^{(1)} + 1$	Guard cyclic shift
$N_{cs}^{(1)} + 2$	Format 2/2a/2b (CSI) cyclic shifts
⋮	
10	
11	

keying, with the UE transmitting an SR using the modulation symbol $d(0) = +1$ (i.e. the same constellation point as is used for DTX/NACK or DTX/(NACK,NACK) for ACK/NACK transmission using PUCCH format 1a/1b – see Figure 16.10) to request a PUSCH resource (positive SR transmission), and transmitting nothing when it does not request to be scheduled (negative SR).

Since the HARQ ACK/NACK structure is reused for the SR, different PUCCH resource indices (i.e. different cyclic time shift/orthogonal code combinations) in the same PUCCH region can be assigned for SR (Format 1) or HARQ ACK/NACK (Format 1a/1b) from different UEs. This results in orthogonal multiplexing of SR and HARQ ACK/NACK in the same PUCCH region. The PUCCH resource index to be used by a UE for SR transmission, $m_{\text{PUCCH,SRI}}^{(1)}$, is configured by UE-specific higher-layer signalling.

If a UE needs to transmit a positive SR in the same subframe as a scheduled CSI transmission, the CSI is dropped and only the SR is transmitted, in order to maintain the low CM of the transmit signal. In the case of SRS coinciding with an ACK/NACK or positive SR in the same subframe, the UE drops its SRS transmission if the parameter ‘ackNackSRS-SimultaneousTransmission’ is set to ‘FALSE’ and transmits the SRS otherwise (see [6, Section 8.2]).

If an SR and ACK/NACK happen to coincide in the same subframe, the UE transmits the ACK/NACK on the assigned SR PUCCH resource for a positive SR and transmits ACK/NACK on its assigned ACK/NACK PUCCH resource in case of a negative SR (see [6, Section 7.3]).

The expected behaviour in some other cases of different types of UCI and SRS coinciding in the same subframe can be found in [17].

16.4 Multiplexing of Control Signalling and UL-SCH Data on PUSCH

When UCI is to be transmitted in a subframe in which the UE has been allocated transmission resources for the PUSCH, the UCI is multiplexed together with the UL-SCH data prior to DFT spreading, in order to preserve the low CM single-carrier property; the PUCCH is never transmitted in the same subframe as the PUSCH in Releases 8 and 9. The multiplexing of CQI/PMI, HARQ ACK/NACK, and RI with the PUSCH data symbols onto uplink resource elements (REs) is shown in Figure 16.15.

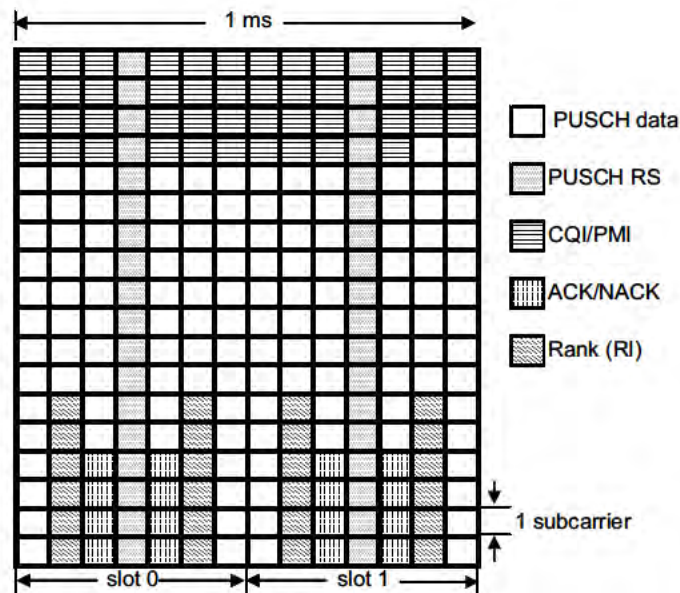


Figure 16.15: Multiplexing of control signalling with UL-SCH data.

The number of REs used for each of CQI/PMI, ACK/NACK and RI is based on the MCS assigned for PUSCH and an offset parameter, $\Delta_{\text{offset}}^{\text{CQI}}$, $\Delta_{\text{offset}}^{\text{HARQ-ACK}}$ or $\Delta_{\text{offset}}^{\text{RI}}$, which is semi-statically configured by higher-layer signalling (see [2, Section 5.2.2.6]). This allows different code rates to be used for the different types of UCI. PUSCH data and UCI are never mapped to the same RE. UCI is mapped in such a way that it is present in both slots of the subframe. Since the eNodeB has prior knowledge of UCI transmission, it can easily demultiplex the UCI and data.

As shown in Figure 16.15, CQI/PMI resources are placed at the beginning of the UL-SCH data resources and mapped sequentially to all SC-FDMA symbols on one subcarrier before continuing on the next subcarrier. The UL-SCH data is rate-matched (see Section 10.3.2.4) around the CQI/PMI data. The same modulation order as UL-SCH data on PUSCH is used for CQI/PMI. For small CQI and/or PMI report sizes up to 11 bits, a (32, *k*) block code,

similar to the one used for PUCCH, is used, with optional circular repetition of encoded data (see [2, Section 5.2.2.6.4]); no Cyclic Redundancy Check (CRC) is applied. For large CSI reports (> 11 bits), an 8-bit CRC is attached and channel coding and rate matching is performed using the tail-biting convolutional code as described in Chapter 10.

The HARQ ACK/NACK resources are mapped to SC-FDMA symbols by puncturing the UL-SCH PUSCH data. Positions next to the RS are used, so as to benefit from the best possible channel estimation. The maximum amount of resource for HARQ ACK/NACK is 4 SC-FDMA symbols.

The coded RI symbols are placed next to the HARQ ACK/NACK symbol positions irrespective of whether ACK/NACK is actually present in a given subframe. The modulation of the 1- or 2-bit ACK/NACK or RI is such that the Euclidean distance of the modulation symbols carrying ACK/NACK and RI is maximized (see [2, Section 5.2.2.6]). The outermost constellation points of the higher-order 16/64-QAM PUSCH modulations are used, resulting in increased transmit power for ACK/NACK/RI relative to the average PUSCH data power.

The coding of the RI and CQI/PMI are separate, with the UL-SCH data being rate-matched around the RI REs similarly to the case of CQI/PMI.

In the case of 1-bit ACK/NACK or RI, repetition coding is used. For the case of 2-bit ACK/NACK/RI, a (3, 2) simplex code is used with optional circular repetition of the encoded data (see [2, Section 5.2.2.6]). The resulting code achieves the theoretical maximum values of the minimum Hamming distance of the output codewords in an efficient way. The (3, 2) simplex codeword mapping is shown in Table 16.9.

Table 16.9: (3, 2) Simplex code for 2-bit ACK/NACK and RI.

2-bit Information Bit Sequence	3-bit Output Codeword
00	000
01	011
10	101
11	110

Control signalling (using QPSK modulation) can also be scheduled to be transmitted on PUSCH without UL-SCH data. The control signalling (CQI/PMI, RI, and/or HARQ ACK/NACK) is multiplexed and scrambled prior to DFT spreading, in order to preserve the low CM single-carrier property. The multiplexing of HARQ ACK/NACK and RI with the CQI/PMI QPSK symbols onto uplink REs is similar to that shown in Figure 16.15. HARQ ACK/NACK is mapped to SC-FDMA symbols next to the RS, by puncturing the CQI data and RI symbols, irrespective of whether ACK/NACK is actually present in a given subframe. The number of REs used for each of ACK/NACK and RI is based on a reference MCS for CQI/PMI and offset parameters, $\Delta_{\text{offset}}^{\text{CQI}}$, $\Delta_{\text{offset}}^{\text{HARQ-ACK}}$ or $\Delta_{\text{offset}}^{\text{RI}}$. The reference CQI/PMI MCS is computed from the CSI payload size and resource allocation. The channel coding and rate matching of the control signalling without UL-SCH data is the same as that of multiplexing control with UL-SCH data as described above.

16.5 ACK/NACK Repetition

It is possible for the network to configure a UE to repeat each ACK/NACK transmission in multiple successive subframes. This is particularly beneficial for ensuring the reliability of the ACK/NACK signalling from power-limited UEs at the cell edge, especially in large cell deployments.

The number of subframes over which each ACK/NACK transmission is repeated is signalled by the UE-specific parameter ‘ackNackRepetition’.

16.6 Multiple-Antenna Techniques

In LTE Releases 8 and 9, simultaneous transmissions from multiple-transmit antennas of a single UE are not supported. Only a single power-amplifier is assumed to be available at the UE. However, LTE does support closed-loop antenna selection transmit diversity in the uplink from UEs which have multiple transmit antennas.

LTE is also designed to support uplink SDMA (or ‘Virtual MU-MIMO’), which is discussed in more detail in Section 16.6.2.

More advanced multiple-antenna techniques, including closed-loop spatial multiplexing for SU-MIMO and transmit diversity for control signalling are introduced in Release 10 for LTE-Advanced, as explained in Section 29.4.

16.6.1 Closed-Loop Switched Antenna Diversity

Uplink closed-loop antenna selection (for up to two transmit antennas) is supported as an optional UE capability in LTE and configured by higher layers (see [18, Section 4.3.4.1]).

If a UE signals that it supports uplink antenna selection, the eNodeB may take this capability into consideration when configuring and scheduling the UE.⁷

16.6.1.1 UE Antenna Selection Indication for PUSCH

When closed-loop antenna selection is enabled, the eNodeB indicates which antenna should be used for the PUSCH by implicitly coding this information in the uplink scheduling grant (Downlink Control Information (DCI) Format 0⁸ – see Section 9.3): the 16 CRC parity bits are scrambled (modulo-2 addition) by one of two antenna selection masks [19], $\langle 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1 \rangle$ for the UE’s first transmit antenna and $\langle 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1 \rangle$ for the second antenna.⁹ The antenna selection mask is applied in addition to the UE-ID masking which indicates for which UE the scheduling grant is intended. This implicit encoding avoids the use of an explicit antenna selection bit which would result in an increased overhead for UEs not supporting (or not configured) for transmit antenna selection.

⁷Alternatively, the eNodeB may permit the UE to use open-loop antenna selection, in which case the UE is free to determine which antenna to transmit from. This may be based on uplink–downlink channel reciprocity, for example in the case of TDD operation (see Section 23.5.2.5.)

⁸Note that closed-loop antenna selection is not supported in conjunction with DCI Format 4 for SU-MIMO transmission in Release 10.

⁹This is applicable only for FDD and half-duplex FDD.

The minimum Hamming distance between the antenna selection masks is only 1 rather than the maximum possible Hamming distance of 16; since the CRC is masked by both the antenna selection indicator and the 16-bit UE-ID, the minimum Hamming distance between the correct UE-ID/antenna selection mask and the nearest erroneous UE-ID/antenna selection mask is 1 for any antenna selection mask. Out of the possible $2^{16} - 1$ incorrect masks, the vast majority ($2^{16} - 2$) result in the misidentification of the UE-ID, such that the performance is similar regardless of the Hamming distance between antenna selection masks. The primary advantage of using the chosen masks is the ease of implementation due to simpler half-space identification, as the eNodeB can allocate UE-IDs with a fixed Most Significant Bit (e.g. MSB set to '0', or, equivalently, UE-IDs from 0 to $2^{15} - 1$). The UE-ID can be detected directly from the 15 least significant bits of the decoded mask without needing to use the transmitted antenna selection mask (bit 16).

The UE behaviour for adaptive/non-adaptive HARQ retransmissions when configured for antenna selection is as follows [19]:

- **Adaptive HARQ.** The antenna indicator (via CRC masking) is always sent in the uplink grant to indicate which antenna to use. For example, for a high-Doppler UE with adaptive HARQ, the eNodeB might instruct the UE to alternate between the transmit antennas or, alternatively, to select the primary antenna. In typical UE implementations, a transmit antenna gain imbalance of 3 to 6 dB between the secondary and primary antennas is not uncommon.
- **Non-adaptive HARQ.** The UE behaviour is unspecified as to which antenna to use. Thus, for low-Doppler conditions, the UE could use the same antenna as that signalled in the uplink grant, while at high-Doppler the UE could hop between antennas or just select the primary antenna. For large numbers of retransmissions with non-adaptive HARQ, the antenna indicated in the uplink grant may not be the best and it is better to let the UE select the antenna to use. If the eNodeB wishes to instruct the UE to use a specific antenna for the retransmissions, it can use adaptive HARQ.

16.6.1.2 Antenna Selection for SRS

If the eNodeB enables a UE's closed-loop antenna selection capability, the SRS transmissions then alternate between the transmit antennas in successive configured SRS transmission subframes, irrespective of whether frequency hopping is enabled or disabled, except when the UE is configured for a single one-shot SRS transmission (see [6, Section 8.2]).

16.6.2 Multi-User 'Virtual' MIMO or SDMA

Uplink MU-MIMO consists of multiple UEs transmitting on the same set of RBs, each using a single transmit antenna. From the point of view of an individual UE, such a mode of operation is hardly visible, being predominantly a matter for the eNodeB to handle in terms of scheduling and uplink reception.

However, in order to support uplink MU-MIMO, LTE specifically provides orthogonal DM-RS using different cyclic time shifts (see Section 15.2.2) to enable the eNodeB to derive independent channel estimates for the uplink from each UE.

A cell can assign up to eight different cyclic time shifts using the 3-bit PUSCH cyclic time shift offset in the uplink scheduling grant. As a maximum of eight cyclic time shifts can be assigned, SDMA of up to eight UEs can be supported in a cell. SDMA between cells (i.e. uplink inter-cell cooperation) is supported in LTE by assigning the same base sequence-groups and/or RS hopping patterns to the different cells as explained in Section 15.3.

16.7 Summary

The main uplink physical channels are the PUSCH for data transmission and the PUCCH for control signalling.

The PUSCH supports resource allocation for both frequency-selective scheduling and frequency-diverse transmissions, the latter being by means of intra- and/or inter-subframe frequency hopping. In Release 10, dual-cluster resource allocations on the PUSCH are also supported, as explained in Section 28.3.6.2.

Control signalling (consisting of ACK/NACK, CQI/PMI and RI) is carried by the PUCCH when no PUSCH resources have been allocated. The PUCCH is deliberately mapped to resource blocks near the edge of the system bandwidth, in order to reduce out-of-band emissions caused by data transmissions on the inner RBs, as well as maximizing flexibility for PUSCH scheduling in the central part of the band.

The control signalling from multiple UEs is multiplexed via orthogonal coding by using cyclic time shifts and/or time-domain block spreading.

Simple multiple antenna techniques in the uplink were already incorporated in LTE Releases 8 and 9, in particular through the support of closed-loop switched antenna diversity and SDMA. These techniques are cost-effective for a UE implementation, as they do not assume simultaneous transmissions from multiple UE antennas. SU-MIMO is introduced in Release 10, as explained in Section 29.4.

References¹⁰

- [1] Motorola, 'R1-072671: Uplink Channel Interleaving', www.3gpp.org, 3GPP TSG RAN WG1, meeting 49bis, Orlando, USA, June 2007.
- [2] 3GPP Technical Specification 36.212, 'Evolved Universal Terrestrial Radio Access (E-UTRA); Multiplexing and Channel Coding', www.3gpp.org.
- [3] 3GPP Technical Specification 36.211, 'Evolved Universal Terrestrial Radio Access (E-UTRA); Physical Channels and Modulation', www.3gpp.org.
- [4] B. Classon, P. Sartori, V. Nangia, X. Zhuang and K. Baum, 'Multi-dimensional Adaptation and Multi-User Scheduling Techniques for Wireless OFDM Systems' in *Proc. IEEE International Conference on Communications*, Anchorage, Alaska, May 2003.
- [5] Motorola, 'R1-071340: Considerations and Recommendations for UL Sounding RS', www.3gpp.org, 3GPP TSG RAN WG1, meeting 48bis, St Julian's, Malta, March 2007.
- [6] 3GPP Technical Specification 36.213, 'Evolved Universal Terrestrial Radio Access (E-UTRA); Physical Layer Procedures', www.3gpp.org.
- [7] Motorola, 'R1-073756: Benefit of Non-Persistent UL Sounding for Frequency Hopping PUSCH', www.3gpp.org, 3GPP TSG RAN WG1, meeting 50, Athens, Greece, August 2007.

¹⁰All web sites confirmed 1st March 2011.

- [8] A. Ghosh, R. Ratasuk, W. Xiao, B. Classon, V. Nangia, R. Love, D. Schwent and D. Wilson, 'Uplink Control Channel Design for 3GPP LTE' in *Proc. IEEE International Symposium on Personal, Indoor and Mobile Radio Communications*, Athens, Greece, September 2007.
- [9] Motorola, 'R1-070084: Coexistence Simulation Results for 5 MHz E-UTRA → UTRA FDD Uplink with Revised Simulation Assumptions', www.3gpp.org, 3GPP TSG RAN WG4, meeting 42, St Louis, USA, February 2007.
- [10] S. Zhou, G. B. Giannakis and Martret C. L., 'Chip-Interleaved Block-Spread Code Division Multiple Access'. *IEEE Trans. on Communications*, Vol. 50, pp. 235–248, February 2002.
- [11] Texas Instruments, 'R1-080190: Embedding ACK/NAK in CQI Reference Signals and Receiver Structures', www.3gpp.org, 3GPP TSG RAN WG1, meeting 51bis, Sevilla, Spain, January 2008.
- [12] Samsung, 'R1-073564: Selection of Orthogonal Cover and Cyclic Shift for High Speed UL ACK Channels', www.3gpp.org, 3GPP TSG RAN WG1, meeting 50, Athens, Greece, August 2007.
- [13] Samsung, Panasonic, Nokia, Nokia Siemens Networks and Texas Instruments, 'R1-080035: Joint Proposal on Uplink ACK/NACK Channelization', www.3gpp.org, 3GPP TSG RAN WG1, meeting 51bis, Sevilla, Spain, January 2008.
- [14] Motorola, 'R1-062072: Uplink Reference Signal Multiplexing Structures for E-UTRA', www.3gpp.org, 3GPP TSG RAN WG1, meeting 46, Tallinn, Estonia, August 2006.
- [15] Panasonic, Samsung and ETRI, 'R1-080983: Way forward on the Cyclic Shift Hopping for PUCCH', www.3gpp.org, 3GPP TSG RAN WG1, meeting 52, Sorrento, Italy, February 2008.
- [16] Nokia Siemens Networks, Nokia and Texas Instruments, 'R1-080931: ACK/NACK channelization for PRBs containing both ACK/NACK and CQI', www.3gpp.org, 3GPP TSG RAN WG1, meeting 52, Sorrento, Italy, February 2008.
- [17] Alcatel-Lucent Shanghai Bell, Alcatel-Lucent, LG Electronics, Panasonic and Qualcomm, 'R1-104135: Confirmation of UE behaviour in case of simultaneously-triggered SR, SRS and CQI', www.3gpp.org, 3GPP TSG RAN WG1, meeting 61bis, Dresden, Germany, June 2010.
- [18] 3GPP Technical Specification 36.306, 'Evolved Universal Terrestrial Radio Access (E-UTRA); User Equipment (UE) Radio Access Capabilities', www.3gpp.org.
- [19] Motorola, Mitsubishi Electric and Nortel, 'R1-081928: Way Forward on Indication of UE Antenna Selection for PUSCH', www.3gpp.org, 3GPP TSG RAN WG1, meeting 53, Kansas City, USA, May 2008.
- [20] Motorola, 'R1-082109: UE Transmit Antenna Selection', www.3gpp.org, 3GPP TSG RAN WG1, meeting 53, Kansas City, USA, May 2008.

17

Random Access

Pierre Bertrand and Jing Jiang

17.1 Introduction

An LTE User Equipment (UE) can only be scheduled for uplink transmission if its uplink transmission timing is synchronized. The LTE Random Access CHannel (RACH) therefore plays a key role as an interface between non-synchronized UEs and the orthogonal transmission scheme of the LTE uplink radio access. In this chapter, the main roles of the LTE RACH are elaborated, together with its differences from the Wideband Code Division Multiple Access (WCDMA) RACH. The rationale for the design of the LTE Physical Random Access CHannel (PRACH) is explained, and some possible implementation options are discussed for both the UE and the eNodeB.

17.2 Random Access Usage and Requirements in LTE

In WCDMA, the RACH is primarily used for initial network access and short message transmission. LTE likewise uses the RACH for initial network access, but in LTE the RACH cannot carry any user data, which is exclusively sent on the Physical Uplink Shared CHannel (PUSCH). Instead, the LTE RACH is used to achieve uplink time synchronization for a UE which either has not yet acquired, or has lost, its uplink synchronization. Once uplink synchronization is achieved for a UE, the eNodeB can schedule orthogonal uplink transmission resources for it. Relevant scenarios in which the RACH is used are therefore:

- (1) A UE in RRC_CONNECTED state, but not uplink-synchronized, needing to send new uplink data or control information (e.g. an event-triggered measurement report);

LTE – The UMTS Long Term Evolution: From Theory to Practice, Second Edition.

Stefania Sesia, Issam Toufik and Matthew Baker.

© 2011 John Wiley & Sons, Ltd. Published 2011 by John Wiley & Sons, Ltd.

- (2) A UE in RRC_CONNECTED state, but not uplink-synchronized, needing to receive new downlink data, and therefore to transmit corresponding ACKnowledgement/Negative ACKnowledgement (ACK/NACK) in the uplink;
- (3) A UE in RRC_CONNECTED state, handing over from its current serving cell to a target cell;
- (4) For positioning purposes in RRC_CONNECTED state, when timing advance is needed for UE positioning (See Section 19.4);
- (5) A transition from RRC_IDLE state to RRC_CONNECTED, for example for initial access or tracking area updates;
- (6) Recovering from radio link failure.

One additional exceptional case is that an uplink-synchronized UE is allowed to use the RACH to send a Scheduling Request (SR) if it does not have any other uplink resource allocated in which to send the SR. These roles require the LTE RACH to be designed for low latency, as well as good detection probability at low Signal-to-Noise Ratio (SNR) (for cell edge UEs undergoing handover) in order to guarantee similar coverage to that of the PUSCH and Physical Uplink Control CHannel (PUCCH).¹

A successful RACH attempt should allow subsequent UE transmissions to be inserted among the scheduled synchronized transmissions of other UEs. This sets the required timing estimation accuracy which must be achievable from the RACH, and hence the required RACH transmission bandwidth: due to the Cyclic Prefix (CP) of the uplink transmissions, the LTE RACH only needs to allow for round-trip delay estimation (instead of the timing of individual channel taps), and this therefore reduces the required RACH bandwidth compared to WCDMA.

This is beneficial in minimizing the overhead of the RACH, which is another key consideration. Unlike in WCDMA, the RACH should be able to be fitted into the orthogonal time-frequency structure of the uplink, so that an eNodeB which wants to avoid interference between the RACH and scheduled PUSCH/PUCCH transmissions can do so. It is also important that the RACH is designed so as to minimize interference generated to adjacent scheduled PUSCH/PUCCH transmissions.

17.3 Random Access Procedure

The LTE random access procedure comes in two forms, allowing access to be either *contention-based* (implying an inherent risk of collision) or *contention-free*.

A UE initiates a contention-based random access procedure for all use-cases listed in Section 17.2. In this procedure, a random access preamble signature is randomly chosen by the UE, which may result in more than one UE simultaneously transmitting the same signature, leading to a need for a subsequent contention resolution process.

For the use-cases (2) (new downlink data) and (3) (handover) the eNodeB has the option of preventing contention occurring by allocating a dedicated signature to a UE, resulting in contention-free access. This is faster than contention-based access – a particularly important factor for the case of handover, which is time-critical.

¹See Chapter 26 for an analysis of the coverage of the PUSCH and PUCCH.

Unlike in WCDMA, a fixed number (64) of preamble signatures is available in each LTE cell, and the operation of the two types of RACH procedure depends on a partitioning of these signatures between those for contention-based access and those reserved for allocation to specific UEs on a contention-free basis.

The two procedures are outlined in the following sections.

17.3.1 Contention-Based Random Access Procedure

The contention-based procedure consists of four-steps as shown in Figure 17.1:

- Step 1: Preamble transmission;
- Step 2: Random access response;
- Step 3: Layer 2 / Layer 3 (L2/L3) message;
- Step 4: Contention resolution message.




Figure 17.1: Contention-based random access procedure. Reproduced by permission of © 3GPP.

17.3.1.1 Step 1: Preamble Transmission

The UE selects one of the $64 - N_{cf}$ available PRACH contention-based signatures, where N_{cf} is the number of signatures reserved by the eNodeB for contention-free RACH. The set of contention-based signatures is further subdivided into two subgroups, so that the choice of signature can carry one bit of information relating to the amount of transmission resource needed to transmit the message at Step 3. The broadcast system information indicates which signatures are in each of the two subgroups (each subgroup corresponding to one value of the one bit of information), as well as the meaning of each subgroup. The UE selects a

signature from the subgroup corresponding to the size of transmission resource needed for the appropriate RACH use case (some use cases require only a few bits to be transmitted at Step 3, so choosing the small message size avoids allocating unnecessary uplink resources). In selecting the appropriate resource size to indicate, the UE takes into account the current downlink path-loss and the required transmission power for the Step 3 message, in order to avoid being granted resources in Step 3 for a message size that would need a transmission power exceeding that which the UE's maximum power would allow. The transmission power required for the Step 3 message is calculated based on some parameters broadcast by the eNodeB, in order that the network has some flexibility to adapt the maximum size for the Step 3 message. The eNodeB can control the number of signatures in each subgroup according to the observed loads in each group.

The initial preamble transmission power setting is based on an open-loop estimation with full compensation for the path-loss. This is designed to ensure that the received power of the preambles is independent of the path-loss. The UE estimates the path-loss by averaging measurements of the downlink Reference Signal Received Power (RSRP). The eNodeB may also configure an additional power offset, depending for example on the desired received Signal to Interference plus Noise Ratio (SINR), the measured uplink interference and noise level in the time-frequency slots allocated to RACH preambles, and possibly also on the preamble format (see Section 17.4.2.2).

17.3.1.2 Step 2: Random Access Response

The Random Access Response (RAR) is sent by the eNodeB on the Physical Downlink Shared CHannel (PDSCH), and addressed with an ID, the Random Access Radio Network Temporary Identifier (RA-RNTI), identifying the time-frequency slot in which the preamble was detected. If multiple UEs had collided by selecting the same signature in the same preamble time-frequency resource, they would each receive the RAR.

The RAR conveys the identity of the detected preamble, a timing alignment instruction to synchronize subsequent uplink transmissions from the UE, an initial uplink resource grant for transmission of the Step 3 message, and an assignment of a temporary Cell Radio Network Temporary Identifier (C-RNTI) (which may or may not be made permanent as a result of the next step – contention resolution). The RAR message can also include a 'backoff indicator' which the eNodeB can set to instruct the UE to back off for a period of time before retrying a random access attempt.

The UE expects to receive the RAR within a time window, of which the start and end are configured by the eNodeB and broadcast as part of the cell-specific system information. The earliest subframe allowed by the specifications occurs 2 ms after the end of the preamble subframe, as illustrated in Figure 17.2. However, a typical delay (measured from the end of the preamble subframe to the beginning of the first subframe of RAR window) is more likely to be 4 ms. Figure 17.2 shows the RAR consisting of the step 2 message (on PDSCH) together with its downlink transmission resource allocation message 'G' (on the Physical Downlink Control CHannel (PDCCH) – see Section 9.3.5).

If the UE does not receive a RAR within the configured time window, it selects a signature again (as in Step 1) and transmits another preamble. The minimum delay for the transmission of another preamble after the end of the RAR window is 3 ms. If the UE receives the PDCCH signalling the downlink resource used for the RAR but cannot satisfactorily decode the RAR




Figure 17.2: Timing of the Random Access Response (RAR) window.

message itself, the minimum delay is increased to 4 ms, to allow for the time taken by the UE in attempting to decode the RAR.

The eNodeB may configure *preamble power ramping* so that the transmission power for each transmitted preamble is increased by a fixed step. However, since the random access preambles in LTE are normally orthogonal to other uplink transmissions, it is less critical than it was in WCDMA to ensure that the initial preamble power is kept low to control interference. Therefore, the proportion of random access attempts which succeed at the first preamble transmission is likely to be higher than in WCDMA, and the need for power ramping is likely to be reduced.

17.3.1.3 Step 3: Layer 2/Layer 3 (L2/L3) Message

This message is the first scheduled uplink transmission on the PUSCH and makes use of Hybrid Automatic Repeat reQuest (HARQ). It conveys the actual random access procedure message, such as an RRC connection request, tracking area update, or scheduling request, but no Non-Access Stratum (NAS) message. It is addressed to the temporary C-RNTI allocated in the RAR at Step 2 and carries either the C-RNTI if the UE already has one (RRC_CONNECTED UEs) or an initial UE identity (the SAE² Temporary Mobile Subscriber Identity (S-TMSI) or a random number). In case of a preamble collision having occurred at Step 1, the colliding UEs will receive the same temporary C-RNTI through the RAR and will also collide in the same uplink time-frequency resources when transmitting their L2/L3 message. This may result in such interference that no colliding UE can be decoded, and the UEs restart the random access procedure after reaching the maximum number of HARQ retransmissions. However, if one UE is successfully decoded, the contention remains unresolved for the other UEs. The following downlink message (in Step 4) allows a quick resolution of this contention.

If the UE successfully receives the RAR, the UE minimum processing delay before message 3 transmission is 5 ms minus the round-trip propagation time. This is shown in Figure 17.3 for the case of the largest supported cell size of 100 km.

17.3.1.4 Step 4: Contention Resolution Message

The contention resolution message uses HARQ. It is addressed to the C-RNTI (if indicated in the L2/L3 message) or to the temporary C-RNTI, and, in the latter case, echoes the UE identity contained in the L2/L3 message. In case of a collision followed by successful decoding of the L2/L3 message, the HARQ feedback is transmitted only by the UE which detects its own UE identity (or C-RNTI); other UEs understand there was a collision, transmit no HARQ feedback, and can quickly exit the current random access procedure and

²System Architecture Evolution.




Figure 17.3: Timing of the message 3 transmission.

start another one. The UE's behaviour upon reception of the contention resolution message therefore has three possibilities:

- The UE correctly decodes the message and detects its own identity: it sends back a positive ACKnowledgement, 'ACK'.
- The UE correctly decodes the message and discovers that it contains another UE's identity (contention resolution): it sends nothing back (Discontinuous Transmission, 'DTX').
- The UE fails to decode the message or misses the DL grant: it sends nothing back ('DTX').

17.3.2 Contention-Free Random Access Procedure

The slightly unpredictable latency of the random access procedure can be circumvented for some use cases where low latency is required, such as handover and resumption of downlink traffic for a UE, by allocating a dedicated signature to the UE on a per-need basis. In this case the procedure is simplified as shown in Figure 17.4. The procedure terminates with the RAR.

17.4 Physical Random Access Channel Design

The random access preamble part of the random access procedure is mapped at the physical layer onto the PRACH. The design of the preamble is crucial to the success of the RACH procedure, and therefore the next section focuses on the details of the PRACH design.

17.4.1 Multiplexing of PRACH with PUSCH and PUCCH

The PRACH is time- and frequency-multiplexed with PUSCH and PUCCH as illustrated in Figure 17.5. PRACH time-frequency resources are semi-statically allocated within the PUSCH region, and repeat periodically. The possibility of scheduling PUSCH transmissions within PRACH slots is left to the eNodeB's discretion.

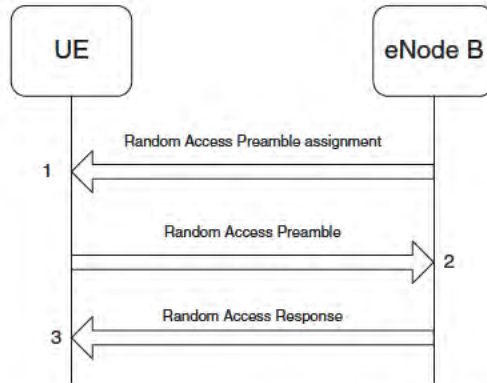


Figure 17.4: Contention-free random access procedure. Reproduced by permission of © 3GPP.

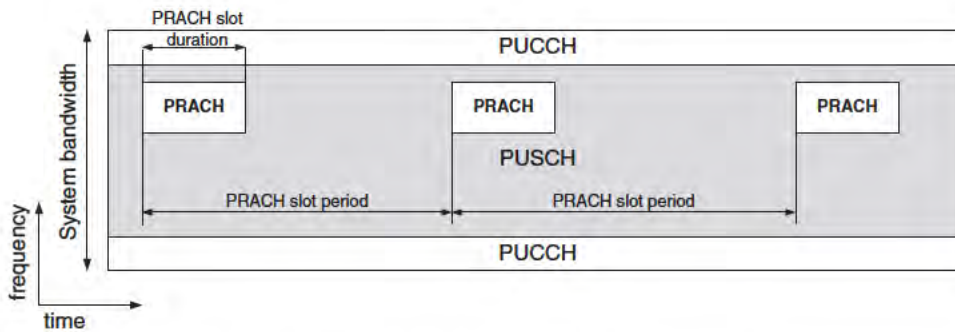


Figure 17.5: PRACH multiplexing with PUSCH and PUCCH.

17.4.2 The PRACH Structure

17.4.2.1 DFT-S-OFDM PRACH Preamble Symbol

Similarly to WCDMA, the LTE PRACH preamble consists of a complex sequence. However, it differs from the WCDMA preamble in that it is also an OFDM symbol following the DFT-S-OFDM structure of the LTE uplink, built with a CP, thus allowing for an efficient frequency-domain receiver at the eNodeB. As shown in Figure 17.6, the end of the sequence is appended at the start of the preamble, thus allowing a periodic correlation at the PRACH receiver (as opposed to a less efficient aperiodic correlation) [1].

The UE aligns the start of the random access preamble with the start of the corresponding uplink subframe at the UE assuming a timing advance of zero (see Section 18.2), and the preamble length is shorter than the PRACH slot in order to provide room for a Guard Time (GT) to absorb the propagation delay. Figure 17.6 shows two preambles at the eNodeB received with different timings depending on the propagation delay: as for a conventional

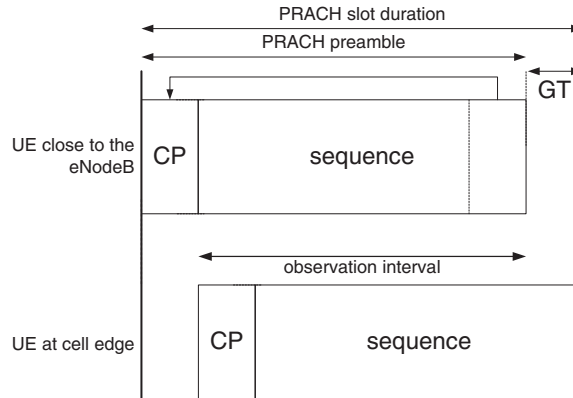


Figure 17.6: PRACH preamble received at the eNodeB.

OFDM symbol, a single observation interval can be used regardless of the UE's delay, within which periodic correlation is possible.

As further elaborated in Section 17.4.3, the LTE PRACH preamble sequence is optimized with respect to its periodic autocorrelation property. The dimensioning of the CP and GT is addressed in Section 17.4.2.4.

17.4.2.2 PRACH Formats

Four Random Access preamble formats are defined for Frequency Division Duplex (FDD) operation [2]. Each format is defined by the durations of the sequence and its CP, as listed in Table 17.1. The format configured in a cell is broadcast in the System Information.

Table 17.1: Random access preamble formats. Reproduced by permission of © 3GPP.

Preamble format	T_{CP} (μ s)	T_{SEQ} (μ s)	Typical usage
0	103.13	800	Normal 1 ms random access burst with 800 μ s preamble sequence, for small to medium cells (up to ~14 km)
1	684.38	800	2 ms random access burst with 800 μ s preamble sequence, for large cells (up to ~77 km) without a link budget problem
2	203.13	1600	2 ms random access burst with 1600 μ s preamble sequence, for medium cells (up to ~29 km) supporting low data rates
3	684.38	1600	3 ms random access burst with 1600 μ s preamble sequence, for very large cells (up to ~100 km)

The rationale behind these choices of sequence duration, CP and GT, are discussed below.

17.4.2.3 Sequence Duration

The sequence duration, T_{SEQ} , is driven by the following factors:

- Trade-off between sequence length and overhead: a single sequence must be as long as possible to maximize the number of orthogonal preambles (see Section 7.2.1), while still fitting within a single subframe in order to keep the PRACH overhead small in most deployments;
- Compatibility with the maximum expected round-trip delay;
- Compatibility between PRACH and PUSCH subcarrier spacings;
- Coverage performance.

Maximum round-trip time. The lower bound for T_{SEQ} must allow for unambiguous round-trip time estimation for a UE located at the edge of the largest expected cell (i.e. 100 km radius), including the maximum delay spread expected in such large cells, namely 16.67 μs . Hence

$$T_{\text{SEQ}} \geq \frac{200 \cdot 10^3}{3 \cdot 10^8} + 16.67 \cdot 10^{-6} = 683.33 \mu\text{s} \quad (17.1)$$

Subcarrier spacing compatibility. Further constraints on T_{SEQ} are given by the Single-Carrier Frequency Division Multiple Access (SC-FDMA – see Chapter 14) signal generation principle (see Section 17.5), such that the size of the DFT and IDFT³, N_{DFT} , must be an integer number:

$$N_{\text{DFT}} = f_s T_{\text{SEQ}} = k, \quad k \in \mathbb{N} \quad (17.2)$$

where f_s is the system sampling rate (e.g. 30.72 MHz). Additionally, it is desirable to minimize the orthogonality loss in the frequency domain between the preamble subcarriers and the subcarriers of the surrounding uplink data transmissions. This is achieved if the PUSCH data symbol subcarrier spacing Δf is an integer multiple of the PRACH subcarrier spacing Δf_{RA} :

$$\Delta f_{\text{RA}} = \frac{f_s}{N_{\text{DFT}}} = \frac{1}{T_{\text{SEQ}}} = \frac{1}{k T_{\text{SYM}}} = \frac{1}{k} \Delta f, \quad k \in \mathbb{N} \quad (17.3)$$

where $T_{\text{SYM}} = 66.67 \mu\text{s}$ is the uplink subframe symbol duration. In other words, the preamble duration must be an integer multiple of the uplink subframe symbol duration:

$$T_{\text{SEQ}} = k T_{\text{SYM}} = \frac{k}{\Delta f}, \quad k \in \mathbb{N} \quad (17.4)$$

An additional benefit of this property is the possibility to reuse the FFT/IFFT⁴ components from the SC-FDMA signal processing for the PUSCH to implement the large DFT/IDFT blocks involved in the PRACH transmitter and receiver (see Section 17.5).

³Discrete Fourier Transform and Inverse Discrete Fourier Transform.

⁴Fast Fourier Transform / Inverse Fast Fourier Transform.

Coverage performance. In general a longer sequence gives better coverage, but better coverage requires a longer CP and GT in order to absorb the corresponding round-trip delay (Figure 17.6). The required CP and GT lengths for PRACH format 0, for example, can therefore be estimated from the maximum round-trip delay achievable by a preamble sequence which can fit into a 1 ms subframe.

Under a noise-limited scenario, as is typical of a low density, medium to large suburban or rural cell, coverage performance can be estimated from a link budget calculation. Under the assumption of the Okumura-Hata empirical model of distance-dependent path-loss $L(r)$ [3,4] (where r is the cell radius in km), the PRACH signal power P_{RA} received at the eNodeB baseband input can be computed as follows:

$$P_{RA}(r) = P_{\max} + G_a - L(r) - LF - P_L \text{ (dB)} \quad (17.5)$$

where the parameters are listed in Table 17.2 (mainly from [5]).

Table 17.2: Link budget parameters for analysis of PRACH preamble coverage.

Parameter	Value
Carrier frequency (f)	2000 MHz
eNodeB antenna height (h_b)	30 m / 60 m
UE antenna height (h_m)	1.5 m
UE transmitter EIRP ^a (P_{\max})	24 dBm (250 mW)
eNodeB Receiver Antenna Gain (including cable loss) (G_a)	14 dBi
Receiver noise figure (N_f)	5.0 dB
Thermal noise density (N_0)	-174 dBm/Hz
Percentage of the area covered by buildings (α)	10%
Required E_p/N_0 (eNodeB with 2 Rx antenna diversity)	18 dB (six-path Typical Urban channel model)
Penetration loss (P_L)	0 dB
Log-normal fade margin (LF)	0 dB

^aEquivalent Isotropic Radiated Power.

The required PRACH preamble sequence duration T_{SEQ} is then derived from the required preamble sequence energy to thermal noise ratio E_p/N_0 to meet a target missed detection and false alarm probability, as follows:

$$T_{SEQ} = \frac{N_0 N_f E_p}{P_{RA}(r) N_0} \quad (17.6)$$

where N_0 is the thermal noise power density (in mW/Hz) and N_f is the receiver noise figure (in linear scale).

Assuming that $E_p/N_0 = 18$ dB is required to meet missed detection and false alarm probabilities of 10^{-2} and 10^{-3} respectively (see Section 17.4.3.3), Figure 17.7 plots the coverage performance of the PRACH sequence as a function of the sequence length T_{SEQ} .

It can be observed from Figure 17.7 that the potential coverage performance of a 1 ms PRACH preamble is in the region of 14 km. As a consequence, the required CP and GT




Figure 17.7: PRACH coverage performance versus sequence duration.

lengths are approximately $(2 \cdot 14000)/(3 \cdot 10^8) = 93.3 \mu s$, so that the upper bound for T_{SEQ} is given by

$$T_{SEQ} \leq 1000 - 2 \cdot 93.33 = 813 \mu s \quad (17.7)$$

Therefore, the longest sequence simultaneously satisfying Equations (17.1), (17.4) and (17.7) is $T_{SEQ} = 800 \mu s$, as used for preamble formats 0 and 1. The resulting PRACH sub-carrier spacing is $\Delta f_{RA} = 1/T_{SEQ} = 1.25$ kHz.

The 1600 μs preamble sequence of formats 2 and 3 is implemented by repeating the baseline 800 μs preamble sequence. These formats can provide up to 3 dB link budget improvement, which is useful in large cells and/or to balance PUSCH/PUCCH and PRACH coverage at low data rates.

17.4.2.4 CP and GT Duration

Having chosen T_{SEQ} , the CP and GT dimensioning can be specified more precisely.

For formats 0 and 2, the CP is dimensioned to maximize the coverage, given a maximum delay spread d : $T_{CP} = (1000 - 800)/2 + d/2 \mu s$, with $d \approx 5.2 \mu s$ (corresponding to the longest normal CP of a PUSCH SC-FDMA symbol). The maximum delay spread is used as a guard period at the end of CP, thus providing protection against multipath interference even for the cell-edge UEs.

In addition, for a cell-edge UE, the delay spread energy at the end of the preamble is replicated at the end of the CP (see Figure 17.8) and is therefore within the observation interval. Consequently, there is no need to include the maximum delay spread in the GT dimensioning. Hence, instead of locating the sequence in the centre of the PRACH slot, it is shifted later by half the maximum delay spread, allowing the maximum Round-Trip Delay

(RTD) to be increased by the same amount. Note that, as for a regular OFDM symbol, the residual delay spread at the end of the preamble from a cell-edge UE spills over into the next subframe, but this is taken care of by the CP at the start of the next subframe to avoid any Inter-Symbol Interference (ISI).




Figure 17.8: PRACH CP/GT dimensioning for formats 0 and 2.

For formats 1 and 3, the CP is dimensioned to address the maximum cell range in LTE, 100 km, with a maximum delay spread of $d \approx 16.67 \mu\text{s}$. In practice, format 1 is expected to be used with a 3-subframe PRACH slot; the available GT in 2 subframes can only address a 77 km cell range. It was chosen to use the same CP length for both format 1 and format 3 for implementation simplicity. Of course, handling larger cell sizes than 100 km with suboptimal CP dimensioning is still possible and is left to implementation.

Table 17.3 shows the resulting cell radius and delay spread ranges associated with the four PRACH preamble formats. The CP lengths are designed to be integer multiples of the assumed system sampling period for LTE, $T_S = 1/30.72 \mu\text{s}$.

Table 17.3: Field durations and achievable cell radius of the PRACH preamble formats.

Preamble format	Number of allocated subframes	CP duration		GT duration		Max. delay spread (μs)	Max. cell radius (km)
		in μs	as multiple of T_S	in μs	as multiple of T_S		
0	1	103.13	3168	96.88	2976	6.25	14.53
1	2	684.38	21024	515.63	15840	16.67	77.34
2	2	203.13	6240	196.88	6048	6.25	29.53
3	3	684.38	21024	715.63	21984	16.67	100.16




Figure 17.9: PRACH preamble formats and cell size dimensioning.

These are also illustrated in Figure 17.9.

17.4.2.5 PRACH Resource Configurations

The PRACH slots shown in Figure 17.5 can be configured to occur in up to 16 different layouts, or *resource configurations*. Depending on the RACH load, one or more PRACH resources may need to be allocated per PRACH slot period. The eNodeB has to process the PRACH very quickly so that message 2 of the RACH procedure can be sent within the required window, as shown in Figure 17.2. In case of more than one PRACH resource per PRACH period it is generally preferable to multiplex the PRACH resources in time (Figure 17.10 – right) rather than in frequency (Figure 17.10 – left). This helps to avoid processing peaks at the eNodeB.




Figure 17.10: Processing peaks of the random access preamble receiver.

Extending this principle, the available slot configurations are designed to facilitate a PRACH receiver which may be used for multiple cells of an eNodeB, assuming a periodic pattern with period 10 ms or 20 ms.

Assuming an operating collision probability per UE, $p_{coll}^{UE} = 1\%$, one PRACH time-frequency resource (with 64 signatures) per 10 ms per 5 MHz can handle an offered load $G = -64 \ln(1 - p_{coll}^{UE}) = 0.6432$ average PRACH attempts, which translates into 128 attempts per second in 10 MHz. This is expected to be a typical PRACH load in LTE.

Assuming a typical PRACH load, example usages and cell allocations of the 16 available resource configurations are shown in Figure 17.11 for different system bandwidths and different numbers of cells per eNodeB. Resource configurations 0–2 and 15 use a 20 ms PRACH period, which can be desirable for small bandwidths (e.g. 1.4 MHz) in order to reduce the PRACH overhead at the price of higher waiting times.




Figure 17.11: Random Access resource configurations.

As can be observed from Figure 17.11, in a three-cell scenario, time collision of PRACH resources can always be avoided except for the 20 MHz case, where collisions can be minimized to two PRACH resources occurring in the same subframe. It should also be noted that in a scenario with six cells per eNodeB, at most two PRACH resources will occur in the same subframe for bandwidths below 20 MHz.

The variety of configurations provided therefore enables efficient dimensioning of a multicell PRACH receiver.

17.4.3 Preamble Sequence Theory and Design

As noted above, 64 PRACH signatures are available in LTE, compared to only 16 in WCDMA. This can not only reduce the collision probability, but also allow for 1 bit of information to be carried by the preamble and some signatures to be reserved for contention-free access (see Section 17.3.2). Therefore, the LTE PRACH preamble called for an improved sequence design with respect to WCDMA. While Pseudo-Noise (PN) based sequences were used in WCDMA, in LTE prime-length Zadoff–Chu (ZC) [6, 7] sequences have been chosen (see Chapter 7 for an overview of the properties of ZC sequences). These sequences enable improved PRACH preamble detection performance. In particular:

- The power delay profile is built from periodic instead of aperiodic correlation;
- The intra-cell interference between different preambles received in the same PRACH resource is reduced;
- Intra-cell interference is optimized with respect to cell size: the smaller the cell size, the larger the number of orthogonal signatures and the better the detection performance;
- The eNodeB complexity is reduced;
- The support for high-speed UEs is improved.

The 800 μs LTE PRACH sequence is built by cyclicly-shifting a ZC sequence of prime-length N_{ZC} , defined as

$$x_u(n) = \exp\left[-j\frac{\pi un(n+1)}{N_{\text{ZC}}}\right], \quad 0 \leq n \leq N_{\text{ZC}} - 1 \quad (17.8)$$

where u is the ZC sequence index and the sequence length $N_{\text{ZC}} = 839$ for FDD.

The reasons that led to this design choice are elaborated in the next sections.

17.4.3.1 Preamble Bandwidth

In order to ease the frequency multiplexing of the PRACH and the PUSCH resource allocations, a PRACH slot must be allocated a bandwidth BW_{PRACH} equal to an integer multiple of Resource Blocks (RBs), i.e. an integer multiple of 180 kHz.

For simplicity, BW_{PRACH} in LTE is constant for all system bandwidths; it is chosen to optimize both the detection performance and the timing estimation accuracy. The latter drives the lower bound of the PRACH bandwidth. Indeed, a minimum bandwidth of ~ 1 MHz is necessary to provide a one-shot accuracy of about $\pm 0.5 \mu\text{s}$, which is an acceptable timing accuracy for PUCCH/PUSCH transmissions.

Regarding the detection performance, one would intuitively expect that the higher the bandwidth, the better the detection performance, due to the diversity gain. However, it is important to make the comparison using a constant signal energy to noise ratio, E_p/N_0 , resulting from accumulation (or despreading) over the same preamble duration, and the same false alarm probability, p_{fa} , for all bandwidths. The latter requires the detection threshold

to be adjusted with respect to the search window size, which increases with the bandwidth. Indeed, it will become clear from the discussion in Section 17.5.2.3 that the larger the search window, the higher p_{fa} . In other words, the larger the bandwidth the higher the threshold relative to the noise floor, given a false alarm target p_{fa_target} and cell size L ; equivalently, the larger the cell size, the higher the threshold relative to the noise floor, given a target p_{fa_target} and bandwidth. As a result, under the above conditions, a smaller bandwidth will perform better than a large bandwidth in a single-path static AWGN channel, given that no diversity improvement is to be expected from such a channel.

Figure 17.12 shows simulation results for the TU-6 fading channel,⁵ comparing detection performance of preamble bandwidths BW_{PRACH} of 6, 12, 25 and 50 RBs. For each bandwidth, the sequence length is set to the largest prime number smaller than $1/BW_{PRACH}$, the false alarm rate is set to $p_{fa_target} = 0.1\%$, the cell radius is 0.7 km and the receiver searches for 64 signatures constructed from 64 cyclic shifts of one root ZC sequence.




Figure 17.12: PRACH missed detection performance comparison for different BW_{PRACH} of 6, 12, 25 and 50 RBs.

We can observe that the best detection performance is achieved by preambles of 6 RBs and 12 RBs for low and high SNRs respectively. The 25-RB preamble has the overall best performance considering the whole SNR range. Thus the diversity gain of large bandwidths only compensates the increased detection threshold in the high SNR region corresponding to misdetection performances in the range of 10^{-3} and below. At a typical 10^{-2} detection probability target, the 6-RB allocation only has 0.5 dB degradation with respect to the best case.

Therefore, a PRACH allocation of 6 RBs provides a good trade-off between PRACH overhead, detection performance and timing estimation accuracy. Note that for the smallest

⁵The six-path Typical Urban (TU) channel model [8].

system bandwidth (1.4 MHz, 6 RBs) the PRACH overlaps with the PUCCH; it is left to the eNodeB implementation whether to implement scheduling restrictions during PRACH slots to avoid collisions, or to let PRACH collide with PUCCH and handle the resulting interference.

Finally, the exact preamble transmission bandwidth is adjusted to isolate PRACH slots from surrounding PUSCH/PUCCH allocations through guard bands, as elaborated in the following section.

17.4.3.2 Sequence Length

The sequence length design should address the following requirements:

- Maximize the number of ZC sequences with optimal cross-correlation properties;
- Minimize the interference to/from the surrounding scheduled data on the PUSCH.

The former requirement is guaranteed by choosing a prime-length sequence. For the latter, since data and preamble OFDM symbols are neither aligned nor have the same durations, strict orthogonality cannot be achieved. At least, fixing the preamble duration to an integer multiple of the PUSCH symbol provides some compatibility between preamble and PUSCH subcarriers. However, with the 800 μ s duration, the corresponding sequence length would be 864, which does not meet the prime number requirement. Therefore, shortening the preamble to a prime length slightly increases the interference between PUSCH and PRACH by slightly decreasing the preamble sampling rate.

The interference from PUSCH to PRACH is further amplified by the fact that the operating E_s/N_0 of PUSCH (where E_s is the PUSCH symbol energy) is much greater than that of the PRACH (typically as much as 24 dB greater if we assume 13 dB E_s/N_0 for 16QAM PUSCH, while the equivalent ratio for the PRACH would be -11 dB assuming $E_p/N_0 = 18$ dB and adjusting by $-10 \log_{10}(864)$ to account for the sequence length). This is illustrated in Figure 17.13 showing the missed detection rate (P_m) with and without data interference adjacent to the PRACH. The simulations assume a TU-6 channel, two receive antennas at the eNodeB and 15 km/h UE speed. The PRACH shows about 1 dB performance loss at $P_m = 1\%$.

The PRACH uses *guard bands* to avoid the data interference at preamble edges. A cautious design of preamble sequence length not only retains a high inherent processing gain, but also allows avoidance of strong data interference. In addition, the loss of spectral efficiency (by reservation of guard subcarriers) can also be well controlled at a fine granularity ($\Delta f_{RA} = 1.25$ kHz). Figure 17.13 shows the missed detection rate for a cell radius of 0.68 km, for various preamble sequence lengths with and without 16QAM data interference.

In the absence of interference, there is no significant performance difference between sequences of similar prime length. In the presence of interference, it can be seen that reducing the sequence length below 839 gives no further improvement in detection rate. No effect is observed on the false alarm rate.

Therefore the sequence length of 839 is selected for LTE PRACH, corresponding to 69.91 PUSCH subcarriers in each SC-FDMA symbol, and offers $72 - 69.91 = 2.09$ PUSCH subcarriers protection, which is very close to one PUSCH subcarrier protection on each side of the preamble. This is illustrated in Figure 17.14; note that the preamble is positioned




Figure 17.13: Missed detection rates of PRACH preamble with and without 16QAM interferer for different sequence lengths (cell radius of 0.68 km).

centrally in the block of 864 available PRACH subcarriers, with 12.5 null subcarriers on each side.

Finally, the PRACH preamble signal $s(t)$ can therefore be defined as follows [2]:

$$s(t) = \beta_{\text{PRACH}} \sum_{k=0}^{N_{\text{ZC}}-1} \sum_{n=0}^{N_{\text{ZC}}-1} x_{u,v}(n) \cdot \exp\left[-j \frac{2\pi nk}{N_{\text{ZC}}}\right] \times \exp[j2\pi[k + \varphi + K(k_0 + \frac{1}{2})]\Delta f_{\text{RA}}(t - T_{\text{CP}})] \quad (17.9)$$

where $0 \leq t < T_{\text{SEQ}} + T_{\text{CP}}$, β_{PRACH} is an amplitude scaling factor and $k_0 = n_{\text{PRB}}^{\text{RA}} N_{\text{SC}}^{\text{RB}} - N_{\text{RB}}^{\text{UL}} N_{\text{SC}}^{\text{RB}} / 2$. The location in the frequency domain is controlled by the parameter $n_{\text{PRB}}^{\text{RA}}$, expressed as an RB number configured by higher layers and fulfilling $0 \leq n_{\text{PRB}}^{\text{RA}} \leq N_{\text{RB}}^{\text{UL}} - 6$. The factor $K = \Delta f / \Delta f_{\text{RA}}$ accounts for the ratio of subcarrier spacings between the PUSCH and PRACH. The variable φ (equal to 7 for LTE FDD) defines a fixed offset determining the frequency-domain location of the random access preamble within the RBs. $N_{\text{RB}}^{\text{UL}}$ is the uplink system bandwidth (in RBs) and $N_{\text{SC}}^{\text{RB}}$ is the number of subcarriers per RB, i.e. 12.

17.4.3.3 Cyclic Shift Dimensioning (N_{CS}) for Normal Cells

Sequences obtained from cyclic shifts of *different* ZC sequences are not orthogonal (see Section 7.2.1). Therefore, orthogonal sequences obtained by cyclically shifting a single root sequence should be favoured over non-orthogonal sequences; additional ZC root sequences should be used only when the required number of sequences (64) cannot be generated by cyclic shifts of a single root sequence. The cyclic shift dimensioning is therefore very important in the RACH design.




Figure 17.14: PRACH preamble mapping onto allocated subcarriers.

The cyclic shift offset N_{CS} is dimensioned so that the Zero Correlation Zone (ZCZ) (see Section 7.2.1) of the sequences guarantees the orthogonality of the PRACH sequences regardless of the delay spread and time uncertainty of the UEs. The minimum value of N_{CS} should therefore be the smallest integer number of sequence sample periods that is greater than the maximum delay spread and time uncertainty of an uplink non-synchronized UE, plus some additional guard samples provisioned for the spill-over of the pulse shaping filter envelope present in the PRACH receiver (Figure 17.15).




Figure 17.15: Cyclic shift dimensioning.

Table 17.4: Cell scenarios with different cyclic shift increments.

Cell scenario	Number of cyclic shifts per ZC sequence	Number of ZC root sequences	Cyclic shift size N_{CS} (samples)	Cell radius (km)
1	64	1	13	0.7
2	32	2	26	2.5
3	18	4	46	5
4	9	8	93	12

The resulting lower bound for cyclic shift N_{CS} can be written as

$$N_{CS} \geq \left\lceil \left(\frac{20}{3} r - \tau_{ds} \right) \frac{N_{ZC}}{T_{SEQ}} \right\rceil + n_g \quad (17.10)$$

where r is the cell size (km), τ_{ds} is the maximum delay spread, $N_{ZC} = 839$ and T_{SEQ} are the PRACH sequence length and duration (measured in μs) respectively, and n_g is the number of additional guard samples due to the receiver pulse shaping filter.

The delay spread can generally be assumed to be constant for a given environment. However, the larger the cell, the larger the cyclic shift required to generate orthogonal sequences, and consequently, the larger the number of ZC root sequences necessary to provide the 64 required preambles.

The relationship between cell size and the required number of ZC root sequences allows for some system optimization. In general, the eNodeB should configure N_{CS} independently in each cell, because the expected inter-cell interference and load (user density) increases as cell size decreases; therefore smaller cells need more protection from co-preamble interference than larger cells.

Some practical examples of this optimization are given in Table 17.4, showing four cell scenarios resulting from different N_{CS} values configured by the eNodeB. For each scenario, the total number of sequences is 64, but resulting from different combinations of the number of root sequences and cyclic shifts.

N_{CS} set design. Given the sequence length of 839, allowing full flexibility in signalling N_{CS} would lead to broadcasting a 10-bit parameter, which is over-dimensioning. As a result, in LTE the allowed values of N_{CS} are quantized to a predefined set of just 16 configurations. The 16 allowed values of N_{CS} were chosen so that the number of orthogonal preambles is as close as possible to what could be obtained if there were no restrictions on the value of N_{CS} [9]. This is illustrated in Figure 17.16 (left), where the cell radii are derived assuming a delay spread of $5.2 \mu s$ and 2 guard samples n_g for the pulse shaping filter.

The effect of the quantization is shown in Figure 17.16 (right), which plots the probability p_2 that two UEs randomly select two preambles on the same root sequence, as a function of the cell radius, for both the quantized N_{CS} set and an ideal unquantized set. The larger p_2 , the better the detection performance. The figure also shows an ideal unquantized set. It can be seen that the performance loss due to the quantization is negligible.

Figure 17.17 illustrates the range of N_{CS} values and their usage with the various preamble formats. Note that this set of N_{CS} values is designed for use in low-speed cells. LTE also




Figure 17.16: Number of orthogonal preambles (left) and probability that two UEs select two orthogonal preambles (right).

provides a second N_{CS} set specially designed for high-speed cells, as elaborated in the following sections.




Figure 17.17: N_{CS} values and usage with the various preamble formats (low speed cells).

17.4.3.4 Cyclic Shift (N_{CS}) Restriction for High-Speed Cells

The support of 64 RACH preambles as described above assumes little or no frequency shifting due to Doppler spread, in the presence of which ZC sequences lose their zero auto-correlation property. In the presence of a frequency offset δf , it can be shown that the PRACH

ZC sequence in Equation (17.8) is distorted as follows:

$$\begin{aligned} x_u(n, \delta f) &= \exp \left[-j\pi \frac{u(n - 1/u)(n - 1/u + 1)}{N_{ZC}} \right] \\ &\quad \times \exp \left[j2\pi \frac{n}{N_{ZC}} (\delta f T_{SEQ} - 1) \right] \exp \left[-\frac{j\pi}{N_{ZC}} \frac{u - 1}{u} \right] \\ &= x_u(n - 1/u) \exp \left[j2\pi \frac{n}{N_{ZC}} (\delta f T_{SEQ} - 1) \right] e^{j\Phi_u} \end{aligned} \quad (17.11)$$

A similar expression can be written for the opposite frequency offset.

As can be observed, frequency offsets as large as one PRACH subcarrier ($\delta f = \pm \Delta f_{RA} = \pm 1/T_{SEQ} = \pm 1.25$ kHz) result in cyclic shifts $d_u = (\pm 1/u) \bmod N_{ZC}$ on the ZC sequence $x_u(n)$. (Note that $u \cdot d_u \bmod N_{ZC} = \pm 1$.) This frequency offset δf can be due to the accumulated frequency uncertainties at both UE transmitter and eNodeB receiver and the Doppler shift resulting from the UE motion in a Line of Sight (LOS) radio propagation condition.

Figure 17.18 illustrates the impact of the cyclic shift distortion on the received Power Delay Profile (PDP): it creates false alarm peaks whose relative amplitude to the correct peak depends on the frequency offset. The solution adopted in LTE to address this issue is referred to as ‘cyclic shift restriction’ and consists of ‘masking’ some cyclic shift positions in the ZC root sequence. This makes it possible to retain an acceptable false alarm rate, while also combining the PDPs of the three uncertainty windows, thus also maintaining a high detection performance even for very high-speed UEs.




Figure 17.18: Side peaks in PDP due to frequency offset.

It should be noted that at $|\delta f| = \Delta f_{RA}$, the preamble peak completely disappears at the desired location (as per Equation (17.11)). However, the false image peak begins to appear even with $|\delta f| < \Delta f_{RA}$. Another impact of the side peaks is that they restrict the possible cyclic shift range so as to prevent from side peaks from falling into the cyclic shift region (see Figure 17.19). This restriction on N_{CS} is captured by Equations (17.12) and (17.13) and is important for the design of the high-speed N_{CS} set (explained in Section 17.4.3.5) and the

order in which the ZC sequences are used (explained in Section 17.4.3.6):

$$N_{CS} \leq d < (N_{ZC} - N_{CS})/2 \tag{17.12}$$

where

$$d = \begin{cases} d_u, & 0 \leq d_u < N_{ZC}/2 \\ N_{ZC} - d_u, & d_u \geq N_{ZC}/2 \end{cases} \tag{17.13}$$




Figure 17.19: Side peaks within the signature search window.

We use C_{-1} and C_{+1} to denote the two wrong cyclic shift windows arising from the frequency offset, while C_0 denotes the correct cyclic shift window (Figure 17.18). The cyclic restriction rule must be such that the two wrong cyclic shift windows C_{-1} and C_{+1} of a cyclicly-shifted ZC sequence overlap none of C_0 , C_{-1} or C_{+1} of other cyclicly-shifted ZC sequences, nor the correct cyclic shift window C_0 of the same cyclicly-shifted ZC sequence, nor each other [10]. Finally, the restricted set of cyclic shifts is obtained such that the minimum difference between two cyclic shifts is still N_{CS} but the cyclic shifts are not necessarily multiples of N_{CS} .

It is interesting to check the speed limit beyond which it is worth considering a cell to be a high-speed cell. This is done by assessing the performance degradation of the PRACH at the system-level as a function of the UE speed when no cyclic shift restriction is applied.

For this analysis, we model the RACH access attempts of multiple concurrent UEs with a Poisson arrival rate. A preamble detection is considered to be correct if the timing estimation is within $2 \mu s$. A target E_p/N_0 of 18 dB is used for the first preamble transmission, with a power ramping step of 1 dB for subsequent retransmissions. The cell radius is random between 0.5 and 12 km, with either AWGN or a six-path TU channel, and a 2 GHz carrier frequency. eNodeB and UE frequency errors are modelled randomly within ± 0.05 ppm. The access failure rate is the measure of the number of times a UE unsuccessfully re-tries access attempts (up to a maximum of three retransmissions), weighted by the total number of new access attempts.

Figure 17.20 shows the access failure rate performance for both channel types as a function of the UE speed, for various offered loads G . It can be observed that under fading conditions, the RACH failure rates experience some degradation with the UE speed (which translates into Doppler spread), but remains within acceptable performance even at 350 km/h. For the AWGN channel (where the UE speed translates into Doppler shift) the RACH failure rate stays below 10^{-2} up to UE speeds in the range 150 to 200 km/h. However, at 250 km/h and above, the throughput collapses. Without the cyclic shift restrictions the upper bound for useful performance is around 150–200 km/h.




Figure 17.20: Random access failure rate as a function of UE speed.

17.4.3.5 Cyclic Shift Configuration for High-Speed Cells

The cyclic shift dimensioning for high-speed cells in general follows the same principle as for normal cells, namely maximizing the sequence reuse when group quantization is applied to cyclic shift values. However, for high-speed cells, the cyclic shift restriction needs to be considered when deriving the sequence reuse factor with a specific cyclic shift value. Note that there is no extra signalling cost to support an additional set of cyclic shift configurations for high-speed cells since the one signalling bit which indicates a ‘high-speed cell configuration’ serves this purpose.

The N_{CS} values for high-speed cells are shown in Figure 17.21 for the number of available and used preambles, with both consecutive and non-consecutive (quantized) cyclic shift values. The number of available preambles assumes no cyclic shift restriction at all, as in low-speed cells. It should be noted that with the cyclic shift restriction above, the largest usable high-speed cyclic shift value among all root sequences is 279 (from Equation (17.12)). As is further elaborated in the next section, only the preambles with Cubic Metric (CM) (see Section 21.3.3) below 1.2 dB are considered in Figure 17.21.

Since for small N_{CS} values the sequence usage is not so tight with a generally high sequence reuse factor, a way to simplify design, while still achieving a high reuse factor, is to reuse the small N_{CS} values for normal cells. In Figure 17.21, N_{CS} values up to 46 are from the normal cyclic shift values, corresponding to a cell radius up to 5.8 km. At the high end, the value of 237 rather than 242 is chosen to support a minimum of two high-speed cells when all the 838 sequences are used. The maximum supportable high-speed cell radius is approximately 33 km, providing sufficient coverage for preamble formats 0 and 2.

17.4.3.6 Sequence Ordering

A UE using the contention-based random access procedure described in Section 17.3.1 needs to know which sequences are available to select from. As explained in Section 17.4.3.3, the full set of 64 sequences may require the use of several ZC root sequences, the identity of each of which must be broadcast in the cell. Given the existence of 838 root sequences, in LTE the




Figure 17.21: Number of available and used preambles in low CM group.

signalling is streamlined by broadcasting only the index of the *first* root sequence in a cell, and the UE derives the other preamble signatures from it given a predefined ordering of all the sequences.

Two factors are taken into account for the root sequence ordering, namely the CM [11, 12] of the sequence, and the maximum supportable cell size for high-speed cells (or equivalently the maximum supported cyclic shift). Since CM has a direct impact on cell coverage, the first step in ordering the root sequences is to divide the 838 sequences into a low CM group and a high CM group, using the CM of Quadrature Phase Shift Keying (QPSK) (1.2 dB) as a threshold. The low CM group would be used first in sequence planning (and also for high-speed cells) since it is more favourable for coverage.

Then, within each CM group, the root sequences are classified into subgroups based on their maximum supportable cell radius, to facilitate sequence planning including high-speed cells. Specifically, a sequence subgroup g is the set of all root sequences with their maximum allowed cyclic shifts (S_{max}) derived from Equation (17.12) lying between two consecutive high-speed N_{CS} values according to

$$\begin{aligned}
 N_{CS}(g) \leq S_{max} < N_{CS}(g + 1), \text{ for } g = 0, 1, \dots, G - 2, \text{ and} \\
 S_{max} \geq N_{CS}(G - 1)
 \end{aligned}
 \tag{17.14}$$

for G cyclic shift values, with the set of N_{CS} values being those for high-speed cells. Sequences in each subgroup are ordered according to their CM values.

Figure 17.22 shows the CM and maximum allowed cyclic shift values at high speed for the resulting root sequence index ordering used in LTE. This ordering arrangement reflects a continuous CM transition across subgroups and groups, which ensures that consecutive sequences always have close CM values when allocated to a cell. Thus, consistent cell coverage and preamble detection can be achieved in one cell.




Figure 17.22: CM and maximum allowed cyclic shift of reordered group sequences.

The LTE specifications define the mapping from root sequence index u to a reordered index in a table [2], an enhanced extract of which is given in Figure 17.23. The first 16 subgroups are from the low CM group, and the last 16 from the high CM group. Figure 17.23 shows the corresponding high-speed N_{CS} value for each subgroup. Note that sequences with S_{max} less than 15 cannot be used by any high-speed cells, but they can be used by any normal cell which requires no more than 24 root sequences from this group for a total of 64 preambles. Note also that the ordering of the physical ZC root sequence indices is pairwise, since root sequence indices u and $N_{ZC} - u$ have the same CM and S_{max} values. This helps to simplify the PRACH receiver, as elaborated in Section 17.5.2.2.

17.5 PRACH Implementation

This section provides some general principles for practical implementation of the PRACH function.

CM group	Sub-group no.	N_{CS} (High-Speed)	Logical index (i.e. re-ordered)	Physical root sequence index u (in increasing order of the corresponding logical index number)
Low	0	-	0~23	129, 710, 140, 699, 120, 719, 210, 629, 168, 671, 84, 755, 105, 734, 93, 746, 70, 769, 60, 779, 2, 837, 1, 838
	1	15	24~29	56, 783, 112, 727, 148, 691
	2	18	30~35	80, 759, 42, 797, 40, 799
	3	22	36~41	35, 804, 73, 766, 146, 693
	4	26	42~51	31, 808, 28, 811, 30, 809, 27, 812, 29, 810

High	15	23 7	384~455	3, 836, 19, 820, 22, 817, 41, 798, 38, 801, 44, 795, 52, 787, 45, 794, 63, 776, 67, 772, 72, 767, 76, 763, 94, 745, 102, 737, 90, 749, 109, 730, 165, 674, 111, 728, 209, 630, 204, 635, 117, 722, 188, 651, 159, 680, 198, 641, 113, 726, 183, 656, 180, 659, 177, 662, 196, 643, 155, 684, 214, 625, 126, 713, 131, 708, 219, 620, 222, 617, 226, 613
	16	23 7	456~513	230, 609, 232, 607, 262, 577, 252, 587, 418, 421, 416, 423, 413, 426, 411, 428, 376, 463, 395, 444, 283, 556, 285, 554, 379, 460, 390, 449, 363, 476, 384, 455, 388, 451, 386, 453, 361, 478, 387, 452, 360, 479, 310, 529, 354, 485, 328, 511, 315, 524, 337, 502, 349, 490, 335, 504, 324, 515

	29	18	810~815	309, 530, 265, 574, 233, 606
	30	15	816~819	367, 472, 296, 543
31	-	820~837	336, 503, 305, 534, 373, 466, 280, 559, 279, 560, 419, 420, 240, 599, 258, 581, 229, 610	

Figure 17.23: Example of mapping from logical index to physical root sequence index.

17.5.1 UE Transmitter

The PRACH preamble can be generated at the system sampling rate, by means of a large IDFT as illustrated in Figure 17.24.




Figure 17.24: Functional structure of PRACH preamble transmitter.

Note that the DFT block in Figure 17.24 is optional as the sequence can be mapped directly in the frequency domain at the IDFT input, as explained in Section 7.2.1. The cyclic shift can be implemented either in the time domain after the IDFT, or in the frequency domain before

the IDFT through a phase shift. For all possible system sampling rates, both CP and sequence durations correspond to an integer number of samples.

The method of Figure 17.24 does not require any time-domain filtering at baseband, but leads to large IDFT sizes (up to 24 576 for a 20 MHz spectrum allocation), which are cumbersome to implement in practice.

Therefore, another option for generating the preamble consists of using a smaller IDFT, actually an IFFT, and shifting the preamble to the required frequency location through time-domain upsampling and filtering (hybrid frequency/time-domain generation, shown in Figure 17.25). Given that the preamble sequence length is 839, the smallest IFFT size that can be used is 1024, resulting in a sampling frequency $f_{\text{IFFT}} = 1.28$ Msps. Both the CP and sequence durations have been designed to provide an integer number of samples at this sampling rate. The CP can be inserted before the upsampling and time-domain frequency shift, so as to minimize the intermediate storage requirements.




Figure 17.25: Hybrid frequency/time domain PRACH generation.

17.5.2 eNodeB PRACH Receiver

17.5.2.1 Front-End

In the same way as for the preamble transmitter, a choice can be made for the PRACH receiver at the eNodeB between full frequency-domain and hybrid time/frequency-domain approaches. As illustrated in Figure 17.26, the common parts to both approaches are the CP removal, which always occurs at the front-end at the system sampling rate f_s , the PDP computation and signature detection. The approaches differ only in the computation of the frequency tones carrying the PRACH signal(s).

The full frequency-domain method computes, from the 800 μs worth of received input samples during the observation interval (Figure 17.6), the full range of frequency tones used for UL transmission given the system bandwidth. As a result, the PRACH tones are directly extracted from the set of UL tones, which does not require any frequency shift or time-domain filtering but involves a large DFT computation. Note that even though $N_{\text{DFT}} = n \cdot 2^m$, thus allowing fast and efficient DFT computation algorithms inherited from the building-block construction approach [13], the DFT computation cannot start until the complete sequence is stored in memory, which increases delay.

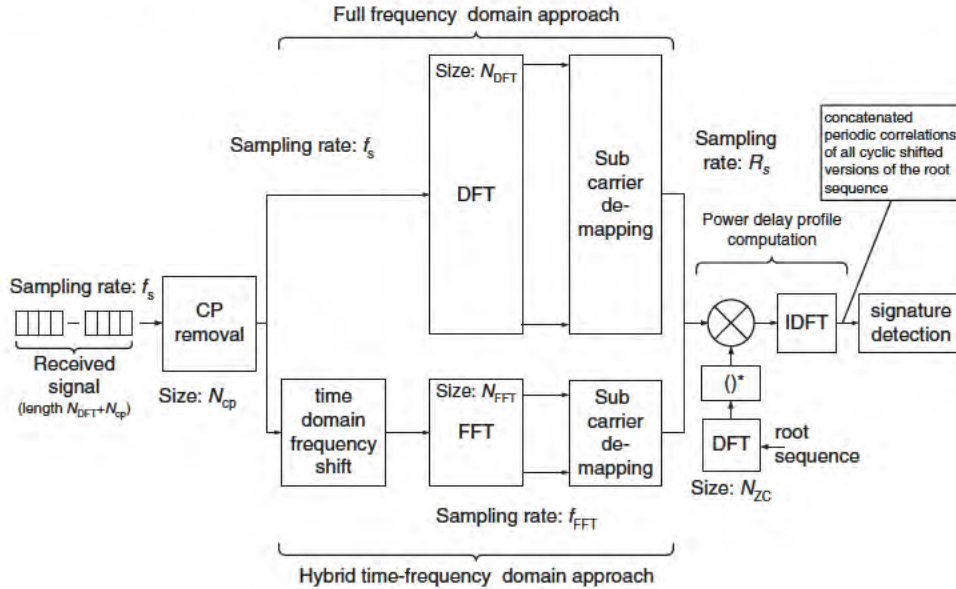


Figure 17.26: PRACH receiver options.

On the other hand, the hybrid time/frequency-domain method first extracts the relevant PRACH signal through a time-domain frequency shift with down-sampling and anti-aliasing filtering. There follows a small-size DFT (preferably an FFT), computing the set of frequency tones centered on the PRACH tones, which can then be extracted. The down-sampling ratio and corresponding anti-aliasing filter are chosen to deliver a number of PRACH time samples suitable for an FFT or simple DFT computation at a sampling rate which is an integer fraction of the system sampling rate. Unlike the full frequency-domain approach, the hybrid time/frequency-domain computation can start as soon as the first samples have been received, which helps to reduce latency.

17.5.2.2 Power Delay Profile (PDP) Computation

The LTE PRACH receiver can benefit from the PRACH format and Constant Amplitude Zero AutoCorrelation (CAZAC) properties as described in the earlier sections by computing the PRACH power delay profile through a frequency-domain periodic correlation. The PDP of the received sequence is given by

$$PDP(l) = |z_u(l)|^2 = \left| \sum_{n=0}^{N_{ZC}-1} y(n) x_u^*[(n + l)_{N_{ZC}}] \right|^2 \tag{17.15}$$

where $z_u(l)$ is the discrete periodic correlation function at lag l of the received sequence $y(n)$ and the reference searched ZC sequence $x_u(n)$ of length N_{ZC} , and where $(\cdot)^*$ denotes the complex conjugate. Given the periodic convolution of the complex sequences $y(n)$ and $x_u(n)$

defined as

$$[y(n) * x_u(n)](l) = \sum_{n=0}^{N_{ZC}-1} y(n)x_u[(l-n)_{N_{ZC}}], \quad (17.16)$$

$z_u(l)$ can be expressed as follows:

$$z_u(l) = [y(n) * x_u^*(-n)](l) \quad (17.17)$$

Let $X_u(k) = R_{X_u}(k) + jI_{X_u}(k)$, $Y_u(k) = R_Y(k) + jI_Y(k)$ and $Z_u(k)$ be the DFT coefficients of the time-domain ZC sequence $x_u(n)$, the received baseband samples $y(n)$, and the discrete periodic correlation function $z_u(n)$ respectively. Using the properties of the DFT, $z_u(n)$ can be efficiently computed in the frequency domain as

$$\begin{cases} Z_u(k) = Y(k)X_u^*(k) & \text{for } k = 0, \dots, N_{ZC} - 1 \\ z_u(n) = \text{IDFT}\{Z_u(k)\}_n & \text{for } n = 0, \dots, N_{ZC} - 1 \end{cases} \quad (17.18)$$

The PDP computation is illustrated in Figure 17.27.




Figure 17.27: PDP computation per root sequence.

The zero padding aims at providing the desired oversampling factor and/or adjusting the resulting number of samples to a convenient IFFT size. Note that for high-speed cells, additional non-coherent combining over three timing uncertainty windows can be performed for each receive antenna, as shown in Figure 17.18.

In addition, the pairwise sequence indexing (Section 17.4.3.6) allows further efficient ‘paired’ matched filtering.

Let $W_u(k)$ be defined as

$$W_u(k) = Y(k)X_u(k) \quad \text{for } k = 0, \dots, N_{ZC} - 1 \quad (17.19)$$

The element-wise multiplications $Z_u(k)$ and $W_u(k)$ can be computed jointly [14], where the partial products $R_Y(k)R_{X_u}(k)$, $I_Y(k)I_{X_u}(k)$, $I_Y(k)R_{X_u}(k)$ and $I_{X_u}(k)R_Y(k)$ only need to be computed once for both $Z_u(k)$ and $W_u(k)$ frequency-domain matched filters.

Now let $w_u(l)$ be defined as

$$w_u(l) = \text{IDFT}\{W_u(k)\}_l \quad \text{for } l = 0, \dots, N_{ZC} - 1 \quad (17.20)$$

By substituting Equation (17.19) into Equation (17.20) and using the property $\text{DFT}\{x_{N_{ZC}-u}(n-1)\}_k = X_u^*(k)$, we can obtain [15]

$$w_u(l) = \sum_{n=0}^{N_{ZC}-1} y(n)x_{N_{ZC}-u}(l+n-1) = [y(n) * x_{N_{ZC}-u}(-n)](l-1) \quad (17.21)$$

As a result, the IDFTs of the joint computation of the frequency-domain matched filters $Z_u(k)$ and $W_u(k)$ provide the periodic correlations of time-domain ZC sequences $x_u(n)$ and $x_{N_{ZC}-u}(n)$, the latter being shifted by one sequence sample.

17.5.2.3 Signature Detection

The fact that different PRACH signatures are generated from cyclic shifts of a common root sequence means that the frequency-domain computation of the PDP of a root sequence provides in one shot the concatenated PDPs of all signatures derived from the same root sequence.

Therefore, the signature detection process consists of searching, within each ZCZ defined by each cyclic shift, the PDP peaks above a detection threshold over a search window corresponding to the cell size. Figure 17.28 shows the basic functions of the signature detector.




Figure 17.28: Signature detection per root sequence.

Detection threshold setting. The target false alarm probability $p_{fa}(T_{det})$ drives the setting of the detection threshold T_{det} .

Under the assumption that the L samples in the uncertainty window are uncorrelated⁶ Gaussian noise with variance σ_n^2 in the absence of preamble transmission, the complex sample sequence $z_a^m(\tau)$ received from antenna a (delayed to reflect a targeted time offset τ of the search window, and despread over a coherent accumulation length (in samples) N_{ca} against the reference code sequence) is a complex Gaussian random variable with variance $\sigma_{n,ca}^2 = N_{ca}\sigma_n^2$. In practice, N_{ca} is the size of the IFFT in Figure 17.27. The non-coherent accumulation $z_{nca}(\tau)$ is modelled as follows:

$$z_{nca}(\tau) = \sum_{a=1}^{N_a} \sum_{m=0}^{N_{nca}-1} |z_a^m(\tau)|^2 \tag{17.22}$$

where N_a is the number of antennas and N_{nca} is the number of additional non-coherent accumulations (e.g. in case of sequence repetition).

$z_{nca}(\tau)$ follows a central chi-square distribution with $2N = 2N_a \cdot N_{nca}$ degrees of freedom, with mean (defining the noise floor) $\gamma_n = N\sigma_{n,ca}^2$ and Cumulative Distribution Function (CDF) $F(T_{det}) = 1 - p_{fa}(T_{det})^L$. It is worth noticing that instead of the absolute threshold we can

⁶The assumption of no correlation between samples holds true in practice up to an oversampling factor of 2.

consider the threshold T_r relative to the noise floor γ_n as follows:

$$T_r = \frac{T_{\text{det}}}{\gamma_n} = \frac{T_{\text{det}}}{N_a \cdot N_{\text{nca}} N_{\text{ca}} \sigma_n^2} \quad (17.23)$$

This removes the dependency of $F(T_r)$ on the noise variance [16]:

$$F(T_r) = 1 - e^{-N_a \cdot N_{\text{nca}} T_r} \sum_{k=0}^{N_a \cdot N_{\text{nca}} - 1} \frac{1}{k!} (N_a \cdot N_{\text{nca}} T_r)^k \quad (17.24)$$

As a result, the relative detection threshold can be precomputed and stored.

Noise floor estimation. For the PDP arising from the transmissions of each root sequence, the noise floor can be estimated as follows:

$$\gamma_n = \frac{1}{N_s} \sum_{i=0, z_{\text{nca}}(\tau_i) < T_{\text{det_ini}}}^{L-1} z_{\text{nca}}(\tau_i) \quad (17.25)$$

where the summation is over all samples less than the absolute noise floor threshold $T_{\text{det_ini}}$ and N_s is the number of such samples. In a real system implementation, the number of additions can be made a power of two by repeating some additions if needed. The initial absolute threshold $T_{\text{det_ini}}$ is computed using an initial noise floor estimated by averaging across all search window samples.

Collision detection. In any cell, the eNodeB can be made aware of the maximum expected delay spread. As a result, whenever the cell size is more than twice the distance corresponding to the maximum delay spread, the eNodeB may in some circumstances be able to differentiate the PRACH transmissions of two UEs if they appear distinctly apart in the PDP. This is illustrated in Figure 17.29, where the upper PDP reflects a small cell, where collision detection is never possible, while the lower PDP represents a larger cell where it may sometimes be possible to detect two distinct preambles within the same ZCZ. If an eNodeB detects a collision, it would not send any random access response, and the colliding UEs would each randomly reselect their signatures and retransmit.

17.5.2.4 Timing Estimation

The primary role of the PRACH preamble is to enable the eNodeB to estimate a UE's transmission timing. Figure 17.30 shows the CDF of the typical performance achievable for the timing estimation. From Figure 17.30 one can observe that the timing of 95% of UEs can be estimated to within 0.5 μs , and more than 98% within 1 μs . These results are obtained assuming a very simple timing advance estimation algorithm, using only the earliest detected peak of a detected signature. No collision detection algorithm is implemented here. The IFFT size is 2048 and the system sampling rate 7.68 MHz, giving an oversampling rate of 2.44.

17.5.2.5 Channel Quality Estimation

For each detected signature, the relative frequency-domain channel quality of the transmitting UE can be estimated from the received preamble. This allows the eNodeB to schedule the L2/L3 message (message 3) in a frequency-selective manner within the PRACH bandwidth.




Figure 17.29: Collision detection.




Figure 17.30: CDF of timing estimation error from the PRACH preamble.

Figure 17.31 shows the BLock Error Rate (BLER) performance of the L2/L3 message of the RACH procedure when frequency-selectively scheduled or randomly scheduled, assuming a typical 10 ms delay between the PRACH preamble and the L2/L3 message. A Least Squares (LS) filter is used for the frequency-domain interpolation, and a single RB is assumed for the size of the L2/L3 message.




Figure 17.31: BLER performance of post-preamble scheduled data.

It can be seen that the performance of a frequency-selectively scheduled L2/L3 message at 10% BLER can be more than 2 dB better than ‘blind scheduling’ at 3 km/h, and 0.5 dB at 10 km/h.

17.6 Time Division Duplex (TDD) PRACH

As discussed in Chapter 23, one design principle of LTE is to maximize the commonality of FDD and TDD transmission modes. With this in mind, the random access preamble formats 0 to 3 are supported in both FDD and TDD operation. In addition, a short preamble format, ‘format 4’, is supported for TDD operation. Format 4 is designed to fit into the short uplink special field known as UpPTS⁷ for small cells (see Section 23.4). Unlike in FDD where there can be at most one PRACH opportunity in one subframe, TDD allows frequency multiplexing of up to six PRACH opportunities in one subframe, to compensate for the smaller number of uplink subframes.

17.6.1 Preamble Format 4

A ZC sequence of length 139 is used for preamble format 4. The preamble starts 157 μ s before the end of the UpPTS field at the UE; it has a T_{CP} of 14.6 μ s and T_{SEQ} is 133 μ s.

Unlike preamble formats 0 to 3, a restricted preamble set for high-speed cells is not necessary for preamble format 4, which uses a 7.5 kHz subcarrier spacing. With a random access duration of two OFDM symbols (157 μ s), the preamble format 4 is mainly used for

⁷Uplink Pilot TimeSlot.

small cells with a cell radius less than 1.5 km, and where cyclic shift restrictions for high UE velocities (see Section 17.4.3.4) are not needed. Therefore, considering that Layer 2 always sees 64 preambles and a sequence length of 139, a smaller set of cyclic shift configurations can be used, as shown in Table 17.5.

Table 17.5: Cyclic shift configuration for preamble format 4. Reproduced by permission of © 3GPP.

N_{CS} configuration	N_{CS} value	Required number of ZC root sequences per cell
0	2	1
1	4	2
2	6	3
3	8	4
4	10	5
5	12	6
6	15	8

Unlike for preamble formats 0 to 3, the root ZC sequence index for preamble format 4 follows the natural pairwise ordering of the physical ZC sequences, with no special restrictions related to the CM or high-speed scenarios. The sequence mapping for preamble format 4 can be formulated as

$$u = \left((-1)^v \left\lfloor \frac{v}{2} + 1 \right\rfloor \right) \bmod N_{ZC} \quad (17.26)$$

from logical sequence index v to physical sequence index u , therefore avoiding the need to prestore a mapping table.

17.7 Concluding Remarks

In this chapter, the detailed design choices of the LTE PRACH have been explained based on theoretical derivations and performance evaluations. In particular, it can be seen how the PRACH preamble addresses the high performance targets of LTE, such as high user density, very large cells, very high speed, low latency and a plurality of use cases, while fitting with minimum overhead within the uplink SC-FDMA transmission scheme.

Many of these aspects benefit from the choice of ZC sequences for the PRACH preamble sequences in place of the PN sequences used in earlier systems. The properties of these sequences enable substantial numbers of orthogonal preambles to be transmitted simultaneously.

Considerable flexibility exists in the selection of the PRACH slot formats and cyclic shifts of the ZC sequences to enable the LTE PRACH to be dimensioned appropriately for different cell radii and loadings.

Some options for the implementation are available, by which the complexity of the PRACH transmitter and receiver can be minimized without sacrificing the performance.

References⁸

- [1] A. V. Oppenheim and R. W. Schaffer, *Discrete-Time Signal Processing*. Englewood Cliffs, NJ: Prentice Hall, 1999.
- [2] 3GPP Technical Specification 36.211, 'Evolved Universal Terrestrial Radio Access (E-UTRA); Physical Channels and Modulation', www.3gpp.org.
- [3] M. Shafi, S. Ogoose and T. Hattori, *Wireless Communications in the 21st Century*. New York: Wiley Inter-Science, 2002.
- [4] Panasonic and NTT DoCoMo, 'R1-062175: Random Access Burst Design for E-UTRA', www.3gpp.org, 3GPP TSG RAN WG1, meeting 46, Tallinn, Estonia, August 2006.
- [5] 3GPP Technical Report 25.814 'Physical layer aspects for E-UTRA (Release 7)', www.3gpp.org.
- [6] R. L. Frank, S. A. Zadoff and R. Heimiller, 'Phase Shift Pulse Codes with Good Periodic Correlation Properties'. *IRE IEEE Trans. on Information Theory*, Vol. 7, pp. 254–257, October 1961.
- [7] D. C. Chu, 'Polyphase Codes with Good Periodic Correlation Properties'. *IEEE Trans. on Information Theory*, Vol. 18, pp. 531–532, July 1972.
- [8] ETSI EN 300 910, 'Radio Transmission and Reception (Release 1999)', www.etsi.org.
- [9] Huawei, 'R1-072325: Multiple Values of Cyclic Shift Increment NCS', www.3gpp.org, 3GPP TSG RAN WG1, meeting 49, Kobe, Japan, May 2007.
- [10] Panasonic and NTT DoCoMo, 'R1-073624: Limitation of RACH Sequence Allocation for High Mobility Cell', www.3gpp.org, 3GPP TSG RAN WG1, meeting 50, Athens, Greece, August 2007.
- [11] Motorola, 'R1-040642: Comparison of PAR and Cubic Metric for Power De-rating', www.3gpp.org, 3GPP TSG RAN WG1, meeting 37, Montreal, Canada, May 2004.
- [12] Motorola, 'R1-060023: Cubic Metric in 3GPP-LTE', www.3gpp.org, 3GPP TSG RAN WG1, LTE adhoc meeting, Helsinki, Finland, January 2006.
- [13] W. W. Smith and J. M. Smith, *Handbook of Real-time Fast Fourier Transforms*. New York: Wiley Inter-Science, 1995.
- [14] Panasonic, 'R1-071517: RACH Sequence Allocation for Efficient Matched Filter Implementation', www.3gpp.org, 3GPP TSG RAN WG1, meeting 48bis, St Julians, Malta, March 2007.
- [15] Huawei, 'R1-071409: Efficient Matched Filters for Paired Root Zadoff–Chu Sequences', www.3gpp.org, 3GPP TSG RAN WG1, meeting 48bis, St Julians, Malta, March 2007.
- [16] J. G. Proakis, *Digital Communications*. New York: McGraw–Hill, 1995.

⁸All web sites confirmed 1st March 2011.

18

Uplink Transmission Procedures

Matthew Baker

18.1 Introduction

In this chapter two procedures are explained which are fundamental to the efficient operation of the LTE uplink: *Timing control* is essential for the orthogonal uplink intra-cell multiple access scheme, while *power control* is important for maintaining Quality-of-Service (QoS), ensuring acceptable User Equipment (UE) battery life and controlling inter-cell interference.

18.2 Uplink Timing Control

18.2.1 Overview

As explained in Chapters 14 and 16, a key feature of the uplink transmission scheme in LTE is that it is designed for orthogonal multiple access in time and frequency between the different UEs. This is fundamentally different from WCDMA, in which the uplink transmissions are non-orthogonal; in WCDMA, from the point of view of the multiple access, there is therefore no need to arrange for the uplink signals from different UEs to be received with any particular timing at the NodeB. The dominant consideration for the uplink transmission timing in WCDMA is the operation of the power control loop, which was designed (in most cases) for a loop delay of just one timeslot (0.666 ms). This is achieved by setting the uplink transmission timing as close as possible to a fixed offset relative to the received downlink timing, without taking into account any propagation delays. Propagation delays in uplink and downlink are absorbed at the NodeB, by means of reducing the time spent measuring the Signal-to-Interference Ratio (SIR) to derive the next power control command.

LTE – The UMTS Long Term Evolution: From Theory to Practice, Second Edition.

Stefania Sesia, Issam Toufik and Matthew Baker.

© 2011 John Wiley & Sons, Ltd. Published 2011 by John Wiley & Sons, Ltd.

For LTE, uplink orthogonality is maintained by ensuring that the transmissions from different UEs in a cell are time-aligned at the receiver of the eNodeB. This avoids intra-cell interference occurring, both between UEs assigned to transmit in consecutive subframes and between UEs transmitting on adjacent subcarriers.

Time alignment of the uplink transmissions is achieved by applying a *timing advance* at the UE transmitter, relative to the received downlink timing. The main role of this is to counteract differing propagation delays between different UEs, as shown in Figure 18.1. A similar approach is used in GSM.

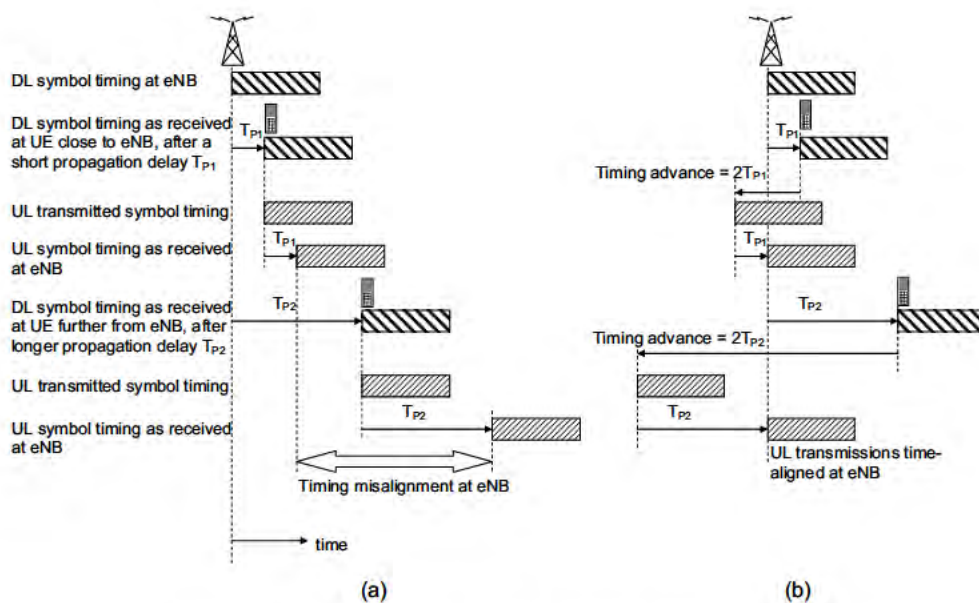


Figure 18.1: Time alignment of uplink transmissions:
(a) without timing advance; (b) with timing advance.

18.2.2 Timing Advance Procedure

18.2.2.1 Initial Timing Advance

After a UE has first synchronized its receiver to the downlink transmissions received from the eNodeB (see Section 7.2), the initial timing advance is set by means of the random access procedure described in Section 17.3. This involves the UE transmitting a random access preamble from which the eNodeB estimates the uplink timing and responds with an 11-bit initial timing advance command contained within the Random Access Response (RAR) message. This allows the timing advance to be configured by the eNodeB with a granularity of $0.52 \mu\text{s}$ from 0 up to a maximum of 0.67 ms ,¹ corresponding to a cell radius

¹This is equal to $16T_s$, where $T_s = 1/30.72 \mu\text{s}$ is the basic unit of time introduced in Section 5.4.1.

of 100 km.² The timing advance was limited to this range in order to avoid further restricting the processing time available at the UE between receiving the downlink signal and having to make a corresponding uplink transmission (see Figures 10.13 and 10.14). In any case, a cell range of 100 km is sufficient for most practical scenarios, and is far beyond what could be achieved with the early versions of GSM, in which the range of the timing advance restricted the cell range to about 35 km. Support of cell sizes even larger than 100 km in LTE is left to the eNodeB implementation to handle.

The granularity of 0.52 μs enables the uplink transmission timing to be set with an accuracy well within the length of the uplink CP (the smallest value of which is 4.7 μs). This granularity is also significantly finer than the length of a cyclic shift of the uplink reference signals (see Chapter 15). Simulations have shown [1, 2] that timing misalignment of up to at least 1 μs does not cause significant degradation in system performance due to increased interference. Thus the granularity of 0.52 μs is sufficiently fine to allow for additional timing errors arising from the uplink timing estimation in the eNodeB and the accuracy with which the UE sets its initial transmission timing – the latter being required to be better than 0.39 μs ³ (except if the downlink system bandwidth is 1.4 MHz, in which case the UE initial transmission timing accuracy is relaxed to 0.79 μs) (see [3, Section 7.1]).

18.2.2.2 Timing Advance Updates

After the initial timing advance has been established for each UE, it will then need to be updated from time to time to counteract changes in the arrival time of the uplink signals at the eNodeB. Such changes may arise from:

- The movement of a UE, causing the propagation delay to change at a rate dependent on the velocity of the UE relative to the eNodeB; at 500 km/h (the highest speed considered for LTE), the round-trip propagation delay would change by a maximum of 0.93 $\mu\text{s/s}$;
- Abrupt changes in propagation delay due to existing propagation paths disappearing and new ones arising; such changes typically occur most frequently in dense urban environments as the UEs move around the corners of buildings;
- Oscillator drift in the UE, where the accumulation of small frequency errors over time may result in timing errors; the frequency accuracy of the oscillator in an LTE UE is required to be better than 0.1 ppm (see [4, Section 6.5.1]), which would result in a maximum accumulated timing error of 0.1 $\mu\text{s/s}$;
- Doppler shift arising from the movement of the UE, especially in Line-Of-Sight (LOS) propagation conditions, resulting in an additional frequency offset of the uplink signals received at the eNodeB.⁴

The updates of the timing advance to counteract these effects are performed by a closed-loop mechanism whereby the eNodeB measures the received uplink timing and issues timing

²In theory it would also be possible to support small negative timing advances (i.e. a timing delay) up to the duration of the Cyclic Prefix (CP), for UEs very close to the eNodeB, without causing loss of uplink time-domain orthogonality. However, this is not supported in LTE.

³This is equal to $12T_s$.

⁴In non-LOS conditions, this becomes a Doppler spread, where the error is typically a zero-mean random variable.

advance update commands to instruct the UE to adjust its transmission timing accordingly, relative to its current transmission timing.⁵

In deriving the timing advance update commands, the eNodeB may measure any uplink signal which is useful. This may include the Sounding Reference Signals (SRSs), DeModulation Reference Signals (DM-RS), Channel Quality Indicator (CQI), ACKnowledgements/Negative ACKnowledgements (ACK/NACKs) sent in response to downlink data, or the uplink data transmissions themselves. In general, wider-bandwidth uplink signals enable a more accurate timing estimate to be made, although this is not, in itself, likely to be a sufficient reason to configure all UEs to transmit wideband SRS very frequently. The benefit of highly accurate timing estimation has to be traded off against the uplink overhead from such signals. In addition, cell-edge UEs are power-limited and therefore also bandwidth-limited for a given uplink Signal-to-Interference-plus-Noise Ratio (SINR); in such cases, the timing estimation accuracy of narrower-bandwidth uplink signals can be improved through averaging multiple measurements over time and interpolating the resulting power delay profile. The details of the uplink timing measurements at the eNodeB are not specified but left to the implementation.

A timing advance update command received at the UE is applied at the beginning of the uplink subframe which begins 4–5 ms after the end of the downlink subframe in which the command is received (depending on the propagation delay), as shown in Figure 18.2. For a TDD or half-duplex FDD system configuration, the new uplink transmission timing takes effect at the start of the first uplink transmission after this point. In the case of an increase in the timing advance relative to the previous transmission, the first part of the subframe in which the new timing is applied is skipped.




Figure 18.2: Application of timing advance update commands.

⁵Transmission timing adjustments arising from timing advance update commands are always made relative to the latest uplink timing. The initial transmission timing is set relative to the received downlink timing; further, if no timing advance update commands are received, the UE is required autonomously to adjust its uplink transmission timing to track changes in the received downlink timing, as specified in [3, Section 7.1.2].

The timing advance update commands are generated at the Medium Access Control (MAC) layer in the eNodeB and transmitted to the UE as MAC control elements which may be multiplexed together with data on the Physical Downlink Shared CHannel (PDSCH). Like the initial timing advance command in the response to the Random Access CHannel (RACH) preamble, the update commands have a granularity of 0.52 μs ; in the case of the update commands, the UE is required to implement them with an accuracy of $\pm 0.13 \mu\text{s}$. The range of the update commands is $\pm 16.6 \mu\text{s}$, allowing a step change in uplink timing equivalent to the length of the extended CP (i.e. 16.67 μs). They would typically not be sent more frequently than about 2 Hz. In practice, fast updates are unlikely to be necessary, as even for a UE moving at 500 km/h the change in round-trip time is not more than 0.93 $\mu\text{s/s}$. The eNodeB must balance the overhead of sending regular timing advance update commands to all the UEs in the cell against a UE's ability to transmit quickly when data arrives in its transmit buffer. The eNodeB therefore configures a timer for each UE, which the UE restarts each time a timing advance update command is received; if the UE does not receive another timing advance update command before the timer expires, it must then consider its uplink to have lost synchronization.⁶ In such a case, in order to avoid the risk of generating interference to uplink transmissions from other UEs, the UE is not permitted to make another uplink transmission of any sort without first transmitting a random access preamble to reinitialize the uplink timing.

One further use of timing advance is to create a switching time between uplink reception at the eNodeB and downlink transmission for TDD and half-duplex FDD operation. This switching time can be generated by applying an additional timing advance offset to the uplink transmissions, to increase the amount of timing advance beyond what is required to compensate for the round-trip propagation delay. Typically a switching time of up to 20 μs may be needed. This is discussed in more detail in Section 23.4.1.

18.3 Power Control

18.3.1 Overview

Uplink transmitter power control in a mobile communication system serves an important purpose: it balances the need for sufficient transmitted energy per bit to maintain the link quality corresponding to the required Quality-of-Service (QoS), against the needs to minimize interference to other users of the system and to maximize the battery life of the mobile terminal.

In achieving this purpose, the power control has to adapt to the characteristics of the radio propagation channel, including path-loss, shadowing and fast fading, as well as overcoming interference from other users – both within the same cell and in neighbouring cells.

The requirements for uplink interference management in LTE are quite different from those for WCDMA. In WCDMA, the uplink is basically non-orthogonal,⁷ and the primary source of interference which has to be managed is intra-cell interference between different users in the same cell. Uplink users in WCDMA share the same time-frequency resources,

⁶Note that loss of uplink synchronization is possible without leaving RRC_CONNECTED state (see Chapter 3).

⁷The later releases of WCDMA do, however, introduce a greater element of orthogonality into the uplink transmissions, by means of lower spreading factors and greater use of time-division multiplexing of different users in HSDPA and HSUPA.

and they generate an interference rise above thermal noise at the NodeB receiver; this is known as ‘Rise over Thermal’ (RoT), and it has to be carefully controlled and shared between users. The primary mechanism for increasing the uplink data rate for a given user in WCDMA is to reduce the spreading factor and increase the transmission power accordingly, consuming a larger proportion of the total available RoT in the cell.

By contrast, in LTE the uplink is basically orthogonal by design, and intra-cell interference management is consequently less critical than in WCDMA. The primary mechanisms for varying the uplink data rate in LTE are varying the transmitted bandwidth and varying the Modulation and Coding Scheme (MCS), while the transmitted Power Spectral Density (PSD) could typically remain approximately constant for a given MCS.

Moreover, in WCDMA the power control [5] was primarily designed with continuous transmission in mind for circuit-switched services, while in LTE fast scheduling of different UEs is applied at 1 ms intervals. This is reflected in the fact that power control in WCDMA is periodic with a loop delay of 0.67 ms and a normal power step of ± 1 dB, while LTE allows for larger power steps (which do not have to be periodic), with a minimum loop delay of about 5 ms (see Figure 18.5).

With these considerations in mind, the power control scheme provided in LTE employs a combination of open-loop and closed-loop control. This in theory requires less feedback than a purely closed-loop scheme, as the closed-loop feedback is only needed to compensate for cases when the UE’s own estimate of the required power setting is not satisfactory.

A typical mode of operation for power control in LTE involves setting a coarse operating point for the transmission PSD⁸ by open-loop means, based on path-loss estimation. This would give a suitable PSD for a reference MCS in the prevailing path-loss and shadowing conditions. Faster adaptation can then be applied around the open-loop operating point by closed-loop power control. This can control interference and fine-tune the power setting to suit the channel conditions (including fast fading). However, due to the orthogonal nature of the LTE uplink, the LTE closed-loop power control does not need to be as fast as in WCDMA – in LTE it would typically be expected to operate at no more than a few hundred Hertz.

Meanwhile, the fastest and most frequent adaptation of the uplink transmissions is by means of the uplink scheduling grants, which vary the transmitted bandwidth (and accordingly the total transmitted power), together with setting the MCS, in order to reach the desired transmitted data rate.

With this combination of mechanisms, the power control scheme in LTE in practice provides support for more than one mode of operation. It can be seen as a ‘toolkit’ from which different power control strategies can be selected and used depending, for example, on the deployment scenario or system loading.

18.3.2 Detailed Power Control Behaviour

Detailed power control formulae are specified in LTE for the Physical Uplink Shared CHannel (PUSCH), Physical Uplink Control CHannel (PUCCH) and the Sounding Reference Signals (SRSs) [6]. The formula for each of these uplink signals follows the same basic principles; though they appear complex, in all cases they can be considered as a summation of two main terms: a basic open-loop operating point derived from static or semi-static parameters

⁸In LTE, the PSD is set as a power per Resource Block (RB); if multiple RBs are transmitted by a UE in a subframe the power per RB is the same for all RBs.

signalled by the eNodeB, and a dynamic offset updated from subframe to subframe:

$$\text{Power per resource block} = \text{basic open-loop operating point} + \text{dynamic offset} .$$

18.3.2.1 Basic Open-Loop Operating Point

The basic open-loop operating point for the transmit power per Resource Block (RB) depends on a number of factors including the inter-cell interference and cell load. It can be further broken down into two components:

- a semi-static base level, P_0 , comprising a nominal power level that is common for all UEs in the cell (measured in dBm per RB) and a UE-specific offset;
- an open-loop path-loss compensation component.

Different base levels can be configured for PUSCH data transmissions depending on the scheduling mode: those which are dynamically scheduled (i.e. using Physical Downlink Control CHannel (PDCCH) signalling) and those which use Semi-Persistent Scheduling (SPS) (see Section 4.4.2.1). This in principle allows different BLER (BBlock Error Rate) operating points to be used for dynamically scheduled and SPS transmissions. One possible use for different BLER operating points is to achieve a lower probability of retransmission for SPS transmissions, hence avoiding the PDCCH signalling overhead associated with dynamically scheduled retransmissions; this is consistent with using SPS for delivery of services such as VoIP with minimal signalling overhead.

The UE-specific offset component of the base level P_0 enables the eNodeB to correct for systematic offsets in a UE's transmission power setting, for example arising from errors in path-loss estimation or in absolute output power setting.

The path-loss compensation component is based on the UE's estimate of the downlink path-loss, which can be derived from the UE's measurement of Reference Signal Received Power (RSRP) (see Section 22.3.1.1) and the known transmission power of the downlink reference signals, which is broadcast by the eNodeB. In order to obtain a reasonable indication of the uplink path-loss, the UE should filter the downlink path-loss estimate with a suitable time-window to remove the effect of fast fading but not shadowing. Typical filter lengths are between 100 and 500 ms for effective operation.

For the PUSCH and SRS, the degree to which the uplink PSD is adapted to compensate for the path-loss can be set by the eNodeB, on a scale from 'no compensation' to 'full compensation'. This is an important feature of power control in LTE and is known as *fractional power control*; it is configured by means of a fractional path-loss compensation factor, referred to as α .

In principle, the combination of the base level P_0 and the path-loss compensation component together allow the eNodeB to configure the degree to which the UE responds to the path-loss. At one extreme, the eNodeB could configure the base level to the lowest level (−126 dBm) and rely entirely on the UE's path-loss measurement to raise the power towards the cell edge; alternatively, the eNodeB can set the base level to a higher value, possibly in conjunction with only partial path-loss compensation.

Disregarding the UE-specific offset, the range of the base level P_0 for the PUSCH (−126 dBm to +24 dBm per RB) is designed to cover the full range of target SINR values for different degrees of path-loss compensation, transmission bandwidths and interference levels. For example, the highest value of P_0 , +24 dBm per RB, corresponds to the maximum likely

transmission power of an LTE UE and would typically only be used if path-loss compensation was not being used at all. The lowest value of P_0 for the PUSCH, -126 dBm, is relevant to a case when full path-loss compensation is used and the uplink transmission and reception conditions are optimal: for example, taking a single RB transmission, with a target SINR at the eNodeB of -5 dB (around the lowest useful SINR), interference-free reception and a 0 dB noise figure for the eNodeB receiver (see Section 21.4.4.2), then the required value of P_0 is the thermal noise level in one RB (180 kHz) minus 5 dB, which gives $P_0 = -126$ dBm.

In general, the maximum path-loss that can be compensated (either by P_0 or by the path-loss compensation component) depends on the required SINR and the transmission bandwidth. Some examples are shown in Figure 18.3, for typical ranges of SINR from -5 dB to $+30$ dB, interference rise above thermal noise from 0 dB to $+30$ dB, and transmission bandwidth from one RB to the maximum LTE system bandwidth of 110 RBs (19.8 MHz). Note that this assumes full path-loss compensation and ignores the dynamic offset.




Figure 18.3: Maximum path-loss compensation in typical scenarios for a 23 dBm UE.

The fractional path-loss compensation factor α can be seen as a tool to trade off the fairness of the uplink scheduling against the total cell capacity. Full path-loss compensation maximizes fairness for cell-edge UEs. However, when considering multiple cells together as a system, the use of partial path-loss compensation can increase the total system capacity in the uplink, as less resources are spent ensuring the success of transmissions from cell-edge UEs and less inter-cell interference is caused to neighbouring cells. Path-loss compensation factors around 0.7 – 0.8 typically give a close-to-maximal uplink system capacity (typically around 15 – 25% greater than can be achieved with full path-loss compensation) without causing significant degradation to the cell-edge data rate.

The principle of fractional power control is illustrated in Figure 18.4. The target received PSD for a given MCS is reduced as the path-loss increases, so that cell-edge UEs cause less inter-cell interference.

Inter-cell interference is of particular concern for UEs located near the edge of a cell, as they may disrupt the uplink transmissions in neighbouring cells. LTE consequently provides an interference coordination mechanism whereby a frequency-dependent ‘Overload Indicator’ (OI) may be signalled directly between eNodeBs to warn a neighbouring eNodeB of high uplink interference levels in specific RBs. In response to this, the neighbouring eNodeB may reduce the permitted power per RB of the UEs which are scheduled in the corresponding RBs




Figure 18.4: (a) Fractional and (b) non-fractional power control.

in its cell(s). It is also possible for eNodeBs to cooperate to avoid scheduling cell-edge UEs in neighbouring cells to transmit in the same RBs. This is discussed further in Section 12.5.

In summary, the basic operating point for the transmit power per RB can be expressed as:

$$\text{Basic operating point} = P_0 + \alpha \cdot \text{PL} \tag{18.1}$$

where α is the fractional path-loss compensation factor which allows the trade-off between total uplink capacity and cell-edge data rate, and PL is the downlink path-loss estimate.

For the low-rate PUCCH (carrying ACK/NACK and CQI signalling), full path-loss compensation is always applied, as the PUCCH transmissions from different users are code-division-multiplexed. Full path-loss compensation facilitates good control of the interference between the different users, and hence helps to maximize the number of users which can be accommodated simultaneously on the PUCCH. A different base level P_0 is also provided for the PUCCH compared to that used for the PUSCH, assuming a reference PUCCH format.

For the SRS, an additional semi-static offset relative to the PUSCH power operating point may be configured by RRC signalling.

18.3.2.2 Dynamic Offset

The dynamic offset part of the power per RB can also be broken down into two components:

- A component dependent on the MCS;
- Explicit Transmitter Power Control (TPC) commands.

MCS-dependent component. The MCS-dependent component (referred to in the LTE specifications as Δ_{TF} , where TF stands for ‘Transport Format’) allows the transmitted power per RB to be adapted according to the transmitted information data rate. Ideally, the transmission power required for a given information data rate should follow the fundamental capacity limit, such that

$$R_N = \log_2(1 + \text{SNR}) \tag{18.2}$$

where R_N is the normalized information data rate per unit bandwidth and can be calculated as the number of information bits per Resource Element (RE) in the RB, denoted here BPRE (Bits Per RE), and SNR is the Signal-to-Noise Ratio. Practical limitations of the system and

receiver can be modelled with a scaling factor $k (> 1)$:

$$\text{BPRE} = \frac{1}{k} \log_2(1 + \text{SNR}) \quad (18.3)$$

It follows that the transmission power required per RB is proportional to $2^{k \cdot \text{BPRE}} - 1$. A suitable value for k is taken as 1.25 for the MCS-dependent power offset when enabled [7].

The MCS-dependent component of the transmit power setting can act like a power control command, as the MCS is under the direct control of the eNodeB scheduler: by changing the MCS which the UE has permission to transmit, the eNodeB can quickly apply an indirect adjustment to the UE's transmit PSD via the MCS-dependent component of the transmit power setting. This may be done to take into account the instantaneous buffer status, available power headroom and QoS requirements of the UE.

The MCS-dependent component can also be used to allow an element of frequency-dependent power control, for example in cases where explicit power control commands (discussed in more detail below) are not transmitted frequently and are therefore following only the wideband fading characteristics; for example, by scheduling a low-rate MCS when the UE is granted permission to transmit in a particular part of the band, the eNodeB can dictate a low transmission power in those RBs.

Another use for the MCS-dependent component is in cases where the number of uplink RBs allocated to a UE in a subframe is not matched to the desired data rate and SIR. One example is to enable the transmit power to be reduced if the amount of data to be transmitted is less than the rate supported by the radio channel in a single RB.

The MCS-dependent component for the PUSCH can be set to zero if it is not needed, for example if fast Adaptive Modulation and Coding (AMC) is used instead.⁹

Transport-format-dependent power control is also particularly relevant for the PUCCH, as the PUCCH bandwidth for a UE does not vary depending on the amount of information to be transmitted in a given subframe (ranging from a single bit for a scheduling request or ACK/NACK, to 22 bits for combined dual-codeword ACK/NACK and CQI together – see Section 16.3). For the PUCCH, the magnitude of the power offset for each combination of control information can be adjusted semi-statically by the eNodeB, in order to set a suitable error-rate operating point for each PUCCH format (see Table 16.2). This is analogous to the different power offsets which may be set in HSDPA¹⁰ for ACK/NACK and CQI signalling according to the error rate desired by the network.

UE-specific power control commands. The other component of the dynamic offset is the UE-specific TPC commands. These can operate in two different modes: *accumulative* TPC commands (available for PUSCH, PUCCH and SRS) and *absolute* TPC commands (available for PUSCH only). For the PUSCH, the switch between these two modes is configured semi-statically for each UE by RRC signalling – i.e. the mode cannot be changed dynamically.

With the accumulative TPC commands, each TPC command signals a power step relative to the previous level. This is the default mode and is particularly well-suited to fine-tuning of the transmission power, and to situations where a UE receives power control commands in groups of successive subframes. This mode is similar to the closed-loop power control

⁹In Release 10, the MCS-dependent component cannot be used in conjunction with PUSCH transmission mode 2 for multiple codeword uplink Single User MIMO transmission (see Section 29.4.1).

¹⁰High Speed Downlink Packet Access.

operation in WCDMA, except that the exact values of the power steps are different: in LTE, two sets of power step values are provided: $\{-1, +1\}$ dB and $\{-1, 0, +1, +3\}$ dB (compared to the sets $\{-1, +1\}$ dB and $\{-2, +2\}$ dB in WCDMA). Which of these two sets of power steps is used is determined by the format of the TPC commands and RRC configuration. The maximum size of power step that can be made using accumulative TPC commands is therefore $+3/-1$ dB, but the range over which the power can be adjusted relative to the semi-static operating point is unlimited (except for the maximum and minimum power limits according to the UE power class – see Section 21.3.1.2). Larger power steps can be achieved by combining an accumulative TPC command with an MCS-dependent power step, by changing the MCS. The provision of one set of power step values containing a 0 dB step size enables the transmit power to be kept constant if needed (i.e. without necessarily having to change the transmission power every time a scheduling grant is sent). This is useful, for example, in scenarios where the interference is not expected to vary significantly over time.

By contrast, the transmit power setting that results from an absolute TPC command is independent of the sequence of TPC commands that may have been received previously; the transmit power setting depends only on the most recently received absolute TPC command, which independently signals a power offset relative to the semi-static operating point.¹¹ The set of offsets which can be signalled by absolute TPC commands is $\{-4, -1, +1, +4\}$ dB. Thus the absolute power control mode can only control the power within a range of ± 4 dB from the semi-static operating point, but a relatively large power step can be triggered by a single command (up to ± 8 dB). This mode is therefore suited to scenarios where the scheduling of the UE’s uplink transmissions may be intermittent; an absolute TPC command enables the UE’s transmission power to be adjusted to a suitable level in a single step after each transmission gap. Absolute TPC commands can also be useful for dynamic frequency-domain inter-cell interference coordination.

The timing of the closed-loop TPC commands is illustrated in Figure 18.5. The transmission power change resulting from a TPC command usually takes effect at the fourth uplink subframe after the TPC command is received.¹² Transmission power changes have to be completed within 20 μ s of the relevant subframe boundary (see [4, Section 6.3.4]).

18.3.2.3 Total Transmit Power Setting

Finally, for the PUSCH and SRS, the total transmit power of the UE in each subframe is scaled up linearly from the power level derived from the semi-static operating point and dynamic offset, according to the number of RBs actually scheduled for transmission from the UE in the subframe.

Thus the overall power control equation is as follows:

$$\text{UE transmit power} = \underbrace{P_0 + \alpha \cdot \text{PL}}_{\text{basic open-loop operating point}} + \underbrace{\Delta_{\text{TF}} + f(\Delta_{\text{TPC}})}_{\text{dynamic offset}} + \underbrace{10 \log_{10} M}_{\text{bandwidth factor}}$$

where Δ_{TPC} denotes a TPC command, $f(\cdot)$ represents accumulation in the case of accumulative TPC commands, and M is the number of allocated RBs.

¹¹The absolute TPC mode can be seen as a low-overhead way to adjust the UE-specific offset in the base level component of the semi-static operating point.

¹²For TDD, the execution of the power changes may occur later, depending on the availability of uplink subframes; full details can be found in [6, Section 5.1].

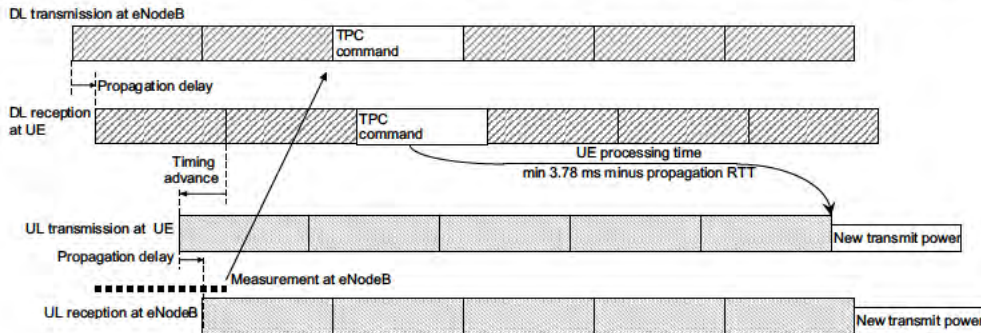


Figure 18.5: Timing of the uplink power control loop.

This overall power control formula allows the UE's transmit power to be controlled with a granularity of 1 dB within the range set -40 dBm to $+23$ dBm (corresponding to a maximum transmission power of 0.2 W). The maximum transmission power of a UE may, however, be subject to additional restrictions:

- The maximum transmission power of the UEs in a cell may be restricted to a lower level by RRC signalling, for example in a hospital scenario;
- In some configurations, reductions in maximum output power may be applied in order to satisfy emissions requirements (see Section 21.3.1.2);
- For UEs which support aggregation of multiple uplink carriers (from Release 10 onwards), the transmission power of an individual carrier or uplink physical channel may have to be scaled down according to defined rules in order to satisfy the total output power constraints (see Section 28.3.5).

The accuracy with which a UE is required to be able to set its total transmission power depends on the length of time since the last uplink transmission (greater or less than 20 ms) and the size of the required change in transmission power; details can be found in Section 21.3.1.2 and in [4, Section 6.3.5].

18.3.2.4 Transmission of TPC Commands

TPC commands for the dynamic offset part of the power control are sent to the UE in messages on the PDCCH. The UE is required to check for a TPC command in every subframe unless it is specifically configured in Discontinuous Reception (DRX – see Section 4.4.1.1). However, unlike in WCDMA, the TPC commands in LTE are not necessarily periodic.

One method by which TPC commands are transmitted to the UEs is in the uplink resource scheduling grant messages for each specific UE. This is logical as it results in all the applicable information for an uplink transmission (set of RBs, transport format, and power setting) being included in a single message.

Additionally, individual accumulative TPC commands for multiple UEs can be jointly coded into a special PDCCH message dedicated to power control (PDCCH Formats 3 and 3A – see Section 9.3.5.1). Such grouped TPC commands may be useful for controlling the

power of uplink SPS transmissions (see Section 4.4.2.1), SRSs, or non-adaptive PUSCH retransmissions (see Section 4.4.1.1). Furthermore, for the PUCCH only, TPC commands can be sent in downlink resource assignment messages on the PDCCH. Both these methods for TPC command transmission enable the power control loop to track changes in channel conditions even when the UE is not scheduled for uplink data transmission, and they can therefore be seen as an alternative to the use of absolute TPC commands. The LTE specifications do not allow jointly coded TPC commands on the PDCCH to be used if the UE is configured in the absolute power control mode.

Due to the structure of the PDCCH signalling (see Section 9.3.5), in all cases the TPC commands are protected by a CRC (Cyclic Redundancy Check); this means that they should be considerably more reliable than in WCDMA. The only likely source of error in LTE would be the UE's failure to detect a PDCCH message, which should typically have a probability around 1% (compared to a typical power control error rate of 4–10% in WCDMA).

The eNodeB can use a number of techniques to determine how to command each UE to adjust its transmit power. One method will be the received SIR, based, for example, on measurements of the SRS and uplink demodulation RSs; in addition the BLER experienced on the decoding of uplink data packets may be used.

The eNodeB may also take into account interference coordination with neighbouring cells, for example if it has received an OI indicating that interference from a UE is causing a problem in a neighbouring cell. Note, however, that although eNodeBs may signal OIs to each other, an eNodeB receiving an OI cannot know for certain whether the overload situation is caused by a UE in its cell or not; it can only infer that, if the received OI relates to a group of RBs where it has scheduled a cell-edge UE, then it is possible that the interference arises from its cell and it should therefore react. Further details of interference coordination are explained in Section 12.5.

18.3.3 UE Power Headroom Reporting

In order to assist the eNodeB to schedule the uplink transmission resources to different UEs in an appropriate way, it is important that the UE can report its available power headroom to the eNodeB. The eNodeB can use the Power Headroom Reports (PHRs) to determine how much more uplink bandwidth per subframe a UE is capable of using. This can help to avoid allocating uplink transmission resources to UEs which are unable to use them; as the uplink is basically orthogonal in LTE, no other UE would be able to use such resources, so system capacity would be wasted.

The range of the PHR is from +40 dB to –23 dB. The negative part of the range enables the UE to signal to the eNodeB the extent to which it has received an uplink resource grant which would require more transmission power than the UE has available. This would enable the eNodeB to reduce the size (i.e. the number of RBs in the frequency domain) of a subsequent grant, thus freeing up transmission resources to allocate to other UEs.

A PHR can only be sent in subframes in which a UE has an uplink transmission grant; the report relates to the subframe in which it is sent. The PHR is therefore a prediction rather than a direct measurement; the UE cannot directly measure its actual transmission power headroom for the subframe in which the report is transmitted. It, therefore, relies on reasonably accurate calibration of the UE's power amplifier output, especially at high output powers when reliable knowledge of the headroom is more critical to system performance.

A number of criteria are defined to trigger a PHR. These include:

- A significant change in estimated path-loss since the last PHR;
- More than a configured time elapse since the previous PHR (controlled by the ‘PHR prohibit timer’);
- More than a configured number of TPC commands implemented by the UE.

The eNodeB can configure parameters to control each of these triggers depending on, for example, the system loading and the requirements of its scheduling algorithm.

In Release 10, some additional aspects are included in the PHRs; these are explained in Section 28.3.5.

18.3.4 Summary of Uplink Power Control Strategies

In summary, a variety of degrees of freedom are available for uplink power control in LTE. Not every parameter will be actively used in every network deployment, but each deployment will select a mode of power control appropriate to the scenario or scheduling strategy. The use of fractional power control facilitates an appropriate trade-off between fairness and system capacity.

One typical mode of operation would be to set the semi-static operating point (via P_0 and the fractional path-loss compensation factor α) to achieve at least the required SINR at the eNodeB for the required QoS for each UE, compensating for path-loss and wideband shadowing. Further control for interference management and rate adaptation can be exercised by means of frequency-domain scheduling and bandwidth adaptation – these being degrees of freedom for power management which were not available in WCDMA. Bandwidth adaptation may also be used in conjunction with changing the MCS to set different BLER operating points for different HARQ processes.

Finally, dynamic transmission power offsets can be used to give a finer degree of control, by means of the MCS-dependent offsets and the closed-loop corrections using the explicit TPC commands.

References¹³

- [1] Nokia, ‘R1-063377: UL Timing Control Accuracy and Update Rate’, www.3gpp.org, 3GPP TSG RAN WG1, meeting 47, Riga, Latvia, November 2006.
- [2] Texas Instruments, ‘R1-072841: Simulation of Uplink Timing Error Impact on PUSCH’, www.3gpp.org, 3GPP TSG RAN WG1, meeting 49bis, Orlando, USA, June 2007.
- [3] 3GPP Technical Specification 36.133, ‘Evolved Universal Terrestrial Radio Access (E-UTRA); Requirements for Support of Radio Resource Management’, www.3gpp.org.
- [4] 3GPP Technical Specification 36.101, ‘Evolved Universal Terrestrial Radio Access (E-UTRA); User Equipment (UE) radio transmission and reception’, www.3gpp.org.
- [5] M. P. J. Baker and T. J. Moulosley, ‘Power Control in UMTS Release 99’ in *Proc. First IEE Int. Conf. on 3G Communications*, March 2000.
- [6] 3GPP Technical Specification 36.213, ‘Evolved Universal Terrestrial Radio Access (E-UTRA); Physical Layer Procedures (Release 8)’, www.3gpp.org.
- [7] Ericsson, ‘R1-080881: Range and Representation of Delta_MCS’, www.3gpp.org, 3GPP TSG RAN WG1, meeting 52, Sorrento, Italy, February 2008.

¹³All web sites confirmed 1st March 2011.

Part IV

Practical Deployment Aspects

19

User Equipment Positioning

Karri Ranta-aho and Zukang Shen

19.1 Introduction

The ability both to locate an object and to communicate with it is a combination that enables a wide range of location-based services – from navigator-like map services to location-based advertising to tracking children, cars or even convicted criminals. This provides a natural motivation for mobile phones to have positioning capabilities. Another strong motivation is a requirement from the Federal Communications Commission (FCC) of the USA that emergency calls, whether fixed or mobile, can be located with a high degree of accuracy.

The most straightforward method by which a mobile phone can know its location is to use a GPS¹ receiver operating completely independently of the cellular network. Alternatively, the cellular network itself can establish the approximate location of a connected mobile phone using knowledge of the coverage area of each cell.

The first version of LTE (Release 8) does not provide any direct protocol support for locating the User Equipment (UE). However, a Release 8 LTE UE can nonetheless be located by means of Assisted Global Navigation Satellite System (A-GNSS) and Enhanced Cell-ID-based techniques (see Sections 19.2 and 19.4 respectively) in conjunction with a general-purpose positioning protocol known as Secure User Plane Location (SUPL), defined by the Open Mobile Alliance (OMA). SUPL operates as a service in the application layer and requires only a normal User Plane (UP) connection between a server in the network (known in OMA as a SUPL Location Platform (SLP) and in LTE as an Evolved Serving Mobile Location Centre (E-SMLC)) and the SUPL client application in the UE. These techniques can provide sufficient average positioning accuracy for many commercial applications, but even when used together they may not meet the emergency call positioning requirements of the FCC with sufficiently high probability in all deployment scenarios.

¹Global Positioning System: www.gps.gov.

LTE – The UMTS Long Term Evolution: From Theory to Practice, Second Edition.

Stefania Sesia, Issam Toufik and Matthew Baker.

© 2011 John Wiley & Sons, Ltd. Published 2011 by John Wiley & Sons, Ltd.

The FCC requires all mobile network operators in the USA to comply with the ‘E911 Phase II’ requirements by January 2013 [1]. For UE-based techniques such as GPS:

- 67% of emergency calls need to be located within 50 m;
- 80% of emergency calls need to be located within 150 m (rising to 90% within 6 years).

GPS can easily meet the accuracy aspect of the requirements, but it cannot provide the required availability due to the satellite signals being blocked in indoor and urban environments. Therefore other techniques are needed in addition.

In GERAN² and UTRAN³ deployments, cell-ID-based location techniques have been successfully used in conjunction with GPS. However, cell-ID-based techniques are limited in their accuracy by the density of deployed cell sites, and therefore the second release of the LTE specifications (Release 9) introduced support for Observed Time Difference Of Arrival (OTDOA) positioning. OTDOA was already defined for UTRAN and a similar technique has been successfully used in CDMA2000 systems.

In addition to support for UP-connection-based positioning (which is transparent to the Radio Access Network (RAN)), support for control plane techniques was also introduced in LTE Release 9 to provide a solution that is more robust against network congestion in emergency scenarios, providing a direct link between the location server and the network nodes.

Table 19.1 summarizes the UE positioning techniques supported in LTE Release 9 [2]. The only positioning technique that can operate in a purely UE-based mode is A-GNSS positioning; for all other techniques the location determination takes place in a location server in the network. eNodeB involvement is not needed for the A-GNSS or OTDOA methods, with the positioning procedure operating directly between the UE and the location server. For Enhanced Cell-ID, the location server may request additional measurement reports from the eNodeB, such as the Angle of Arrival (AoA) of the received signal. These techniques are explained in detail in the following sections.

Table 19.1: UE positioning techniques supported in LTE Release 9.

	UE-based	UE-assisted, E-SMLC-based	eNodeB-assisted
A-GNSS	Yes	Yes	No
OTDOA	No	Yes	No
Enhanced Cell-ID	No	Yes	Yes

²GSM/EDGE Radio Access Network.

³Universal Terrestrial Radio Access Network.

19.2 Assisted Global Navigation Satellite System (A-GNSS) Positioning

A-GNSS-based positioning refers in general to any satellite-based positioning system, such as GPS, Galileo⁴ or GLONASS,⁵ in conjunction with assistance provided over a terrestrial network to improve the sensitivity and/or speed of detection of the satellites, as illustrated in Figure 19.1.




Figure 19.1: Basic principle of A-GNSS.

Basic GNSS-based positioning relies on accurate knowledge of the locations of the satellites and the transmission times of their signals. With some simplification it can be said that a GNSS receiver measures the exact time at which it receives the signal of each satellite it can detect. Using this information, it is possible to calculate the location of the UE. Since GNSS receivers typically do not have high-accuracy atomic clocks, the problem to be solved is four-dimensional (x , y , z and time), and at least four satellites have to be detected for a position estimate.

A GNSS receiver typically has a 'Time To First Fix' (TTFF)⁶ that is too long for most practical uses of mobile terminal positioning. First the receiver must detect the signals transmitted by at least four satellites and decode a message modulated on each satellite's signal providing details of its position in orbit. Only then does the GNSS receiver have sufficient information to perform the measurements and calculate its position. Stand-alone GPS navigation devices typically solve this problem by remaining powered on continually, so that although the first fix may take some time, subsequently the position is always known provided that the satellite signals remain sufficiently strong. However, this would be impractical in a mobile phone as the GNSS receiver would quickly drain the battery even if the location information was rarely needed.

⁴The GNSS of the European Union: www.esa.int/esaNA/galileo.html.

⁵GLOBAL'naya NAVigatsionnaya Sputnikovaya Sistema, the Russian GNSS: www.federal-space.ru.

⁶TTFF describes the time required for a GPS device to acquire the necessary satellite signals and calculate an accurate position.

In order to reduce the TTFF, the terrestrial cellular network can provide assistance data to UEs equipped with GNSS receivers. The simplest such assistance data comprises the messages the satellites broadcast to enable the receiver to calculate their orbital positions at any given time. Receiving this information over the fast cellular network rather than the slow satellite link (50 bps for GPS) directly reduces the TTFF. In addition, the GNSS receiver sensitivity and/or detection speed can be improved by reducing the size of the required satellite signal search window by providing fine time assistance and reference location data to the UE. The more accurate the UE's knowledge of location and time, the smaller the search window it needs for the satellites' signals.

Note that the assistance data serves only to reduce the time taken to determine the position and the probability that the position determination succeeds; it does not improve the accuracy of the location estimate.

A-GNSS can also operate in a UE-assisted mode, where rather than determining its location itself the UE makes measurements of the timing of the received satellite signals and reports them to the E-SMLC. The E-SMLC may then take additional information into account in calculating the UE's position.

More information on A-GNSS can be found in [3].

19.3 Observed Time Difference Of Arrival (OTDOA) Positioning

The principle behind OTDOA positioning is similar to GPS, but for OTDOA the signals measured by the UE are the terrestrial downlink transmissions from the eNodeBs. The location determination is typically distributed between the UE and the network. Unlike GPS positioning, the UE does not acquire an accurate reference time, but the position estimate is based on the received time difference of at least two pairs of cells.

OTDOA in LTE is based on the UE measuring the time difference observed (by the UE receiver) between the Reference Signals (RS – see Section 8.2) of a neighbour cell and those of the serving cell; this is known as a Reference Signal Time Difference (RSTD) measurement (see Figure 19.2). In order to be able to calculate the UE's location, the network needs to know accurately the locations of the eNodeB transmit antennas and the transmission timing of each cell. The fact that the eNodeBs are not orbiting the earth at a high velocity makes the location determination mathematics somewhat simpler than for GPS. An efficient mathematical method for calculating the UE location based on OTDOA measurements is described in [4].

OTDOA positioning does not suffer from long TTFF like GPS, but the requirement that the eNodeB transmission timings are accurately known is not necessarily straightforward if the cellular network is asynchronous. Assuming, however, that these basic requirements can be fulfilled, the key factor governing the achievable performance of a cellular OTDOA system is whether the signals to be measured can be detected by the UE sufficiently fast and with sufficiently high probability.

The requirement that the UE can detect cells of at least three different eNodeBs with a high probability does not sit well with the fact that LTE was designed primarily to provide high-speed data services with good spectral efficiency, for which reason a frequency reuse factor of one is typically used. This can render the Release 8 synchronization signals and RSs of




Figure 19.2: Basic principle of OTDOA in LTE.

distant neighbour cells undetectable to the UE;⁷ on many occasions this would prevent the RSTD measurements required by OTDOA from being made if the available signals were only those defined in Release 8 of the LTE specifications.

Figure 19.3 shows that in an interference-limited hexagonal cell deployment the probability that the SINR of the cell-specific RS of the third-best cell (i.e. the second-best neighbour cell) is above -6 dB is about 70%, which clearly does not give a sufficient level of availability. Figure 19.3 also shows that, even in a 500 m inter-site distance deployment with no interference from user data traffic, the 5-percentile for the SINR of the cell-specific RS of the third best cell is about -13 dB (i.e. for the 95% availability requirement of the FCC E911 Phase II requirements, an SINR of only -13 dB can be assured). Thus it is obvious that even if the UE were enabled to skip the detection of the synchronization signals by providing network assistance information (indicating, for example, the potential neighbour cells for which to search), OTDOA would not work sufficiently well if it were based on the Release 8 cell-specific RSs alone, especially if an interference-free environment could not be guaranteed.

For this purpose, special Positioning Reference Signals (PRSs) are introduced in Release 9 of the LTE specifications, as detailed in the following section.

19.3.1 Positioning Reference Signals (PRS)

In order to increase the probability that the UE can detect sufficient neighbour cells, and therefore achieve good OTDOA positioning reliability, special ‘positioning subframes’ are designated in LTE Release 9. Positioning subframes are designed to aid the ‘hearability’ of neighbour cells by reducing the interference and increasing the RS energy: typically they do not carry any Physical Downlink Shared CHannel (PDSCH) data, but provide PRSs in addition to the Release 8 cell-specific RSs.

⁷A neighbour cell is only considered detectable if the Signal-to-Interference-plus-Noise Ratio (SINR) of the synchronization signals is at least -6 dB.




Figure 19.3: SINR of the Cell-specific RS (CRS) for the best 5 cells, inter-site distance = 500 m (unloaded system) [5]. Reproduced by permission of © Motorola.

Figure 19.4 shows the arrangement of the PRSs in a Resource Block (RB).⁸ The PRS pattern is designed so that it never overlaps with the Physical Downlink Control Channel (PDCCH), nor with the cell-specific RS of any other antenna port. Overlap between PRS patterns in neighbouring cells can be avoided by means of a cell-specific frequency shift of a number of subcarriers given by the Physical Cell ID (PCI) modulo 6, allowing six different non-overlapping frequency shifts. The PRS sequence is constructed in the same way as the cell-specific RS (see Section 8.2.1).

The PRSs are designed to provide more RS energy and a larger reuse factor (namely 6) than is available with the Release 8 cell-specific RS. In synchronous network deployments, even if the interference from data transmissions is eliminated by providing subframes free of PDSCH transmission, the cell-specific RS of cells with two or more antenna ports have a reuse factor of only 3, with the result that they interfere with each other, resulting in reduced hearability of the neighbour cells.

The positioning subframes occur in groups of consecutive downlink subframes known as ‘positioning occasions’, consisting of between one and six positioning subframes (see [6, Section 6.10.4.3]). In TDD deployments, uplink subframes and special subframes cannot contain PRSs.

In order for the UEs to benefit fully from the reduced inter-cell interference towards the PRS of other cells during positioning subframes, the positioning subframes should overlap (at least partially) between the serving cell and all the neighbour cells to be detected, and they need to occur sufficiently frequently for accurate RSTD measurements to be made. The

⁸The eNodeB antenna port used for transmission of the PRS is defined as antenna port number six and is used for PRS transmission only.




Figure 19.4: PRS arrangement (R_6) in a cell with two cell-specific transmit antenna ports. Reproduced by permission of © 3GPP.

best overlap can be achieved in a fully synchronized network, but even partial overlap can be shown to be beneficial.

In some cases a network may further reduce the inter-cell interference by muting the PRS transmissions themselves from time to time, which can be beneficial if the PRS patterns in neighbouring cells overlap. If this is the case, signalling is provided to inform the UEs as to the subframes in which the PRSs are muted. For simplicity, muting can only occur in all or none of the subframes in a given positioning occasion. The signalled pattern of muted positioning occasions is given by a bitmap and repeats periodically; it is applicable in both the serving cell and the neighbour cells. If PRS muting is not applied in a positioning occasion, the UEs can assume that the PRS power is the same in all PRS subframes in the positioning occasion.

In any case, the network should assist the UE in its neighbour-cell search by providing the relative transmission time differences of each neighbour cell relative to the serving cell (in a synchronous network these would always be zero) and the search window length in terms of the maximum propagation delay difference that occurs in the deployment environment. Such assistance can significantly reduce the UE's search space and increase the detection sensitivity.

It should also be noted that OTDOA measurements based on neighbour cell PRS require the UE to have a priori knowledge of the PRS sequences for which to search. As the location servers know the exact geographical locations of the transmit antennas, they can construct a list of candidate neighbour cells for each cell without traditional neighbour list planning. The network can then provide this information to the UEs.

The requirements for the time within which a UE must be capable of reporting an RSTD measurement after receiving OTDOA assistance information can be found in [7].

19.3.2 OTDOA Performance and Practical Considerations

In practical network and UE implementations, a number of sources of error can have an impact on the final accuracy of the location estimates. This is illustrated in Figure 19.5, where the main sources of error are modelled as follows:

- eNodeB synchronization error is caused by imperfect signal transmission timing from the antenna compared to a perfect reference clock. This is modelled by a Gaussian distribution with 100 ns standard deviation. The eNodeB transmit antenna locations are assumed to be perfectly known.
- The UE RSTD measurement quantization error is caused by the UE receiver sampling rate which is assumed to be 32 ns.
- Multipath propagation introduces error to the desired Line-Of-Sight (LOS) propagation delay estimation. In Figure 19.5, the UE takes the centre of gravity of the received multipath signal as the time of reception, and the error is the difference between that estimate and the actual LOS propagation delay.
- Timing offset estimation error and UE frequency instability directly affect the error in the RSTD measurements themselves. The magnitudes of these errors are a function of the SINR.




Figure 19.5: OTDOA positioning accuracy with 4 detected neighbouring cells, inter-site distance = 1732 m, Extended Typical Urban channel model [8].

19.4 Cell-ID-based Positioning

Positioning based on Cell-ID (CID) uses geographical knowledge of a UE’s serving cell. To improve the accuracy, measurements made by the UE and/or the eNodeB can be utilized in addition.

19.4.1 Basic CID Positioning

Basic CID positioning estimates the location of a UE using only the coordinates of its serving eNodeB or cell. Knowledge of the serving eNodeB or cell can be obtained by paging, tracking area updates or other methods. Typically, basic CID positioning provides only coarse estimation of the UE location, with accuracy of roughly the same order as the cell radius.

19.4.2 Enhanced CID Positioning using Round Trip Time and UE Receive Level Measurements

Enhanced CID positioning uses additional information beyond the identity of the eNodeB or cell that is serving the UE. The distance of a UE from its serving eNodeB or cell can be estimated from the Round Trip Time (RTT).

When timing advance is operating (see Section 18.2), Figure 19.6 illustrates the reception and transmission timings at the eNodeB and UE. Two measurements are defined in LTE Release 9 by which an eNodeB can indicate the RTT to the E-SMLC. These are referred to as the ‘Timing Advance Type 1’ and ‘Timing Advance Type 2’ measurements [9]. Alternatively, the E-SMLC or SLP may obtain the timing advance setting directly from the UE.

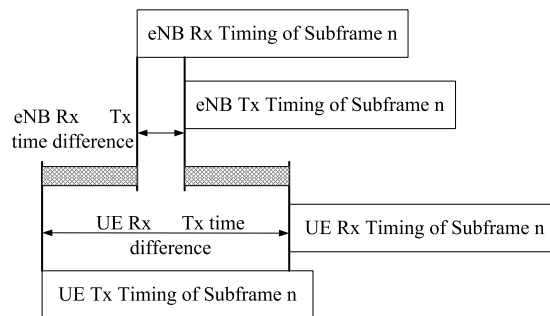


Figure 19.6: RTT estimation in the presence of timing advance.

The Type 1 measurement is the time difference defined as the sum of the receive–transmit time difference at the eNodeB (either a positive or a negative value) and the receive–transmit time difference at the UE (a positive value, due to the operation of timing advance). The

result is represented by the shaded duration in Figure 19.6, corresponding to the RTT. The UE reports its receive–transmit time difference to its serving eNodeB, and the eNodeB calculates its own receive–transmit time difference. The Type 1 method can only be used for UEs which support the reporting of this information – i.e. UEs conforming to Release 9 of the LTE specifications or later.

The Type 2 measurement is simply defined as the receive–transmit time difference at the eNodeB, measured in a radio frame containing a Physical Random Access CHannel (PRACH – see Section 17.4) transmission from the UE whose location is to be ascertained. For the PRACH, the transmit timing of the UE should be aligned with the UE’s receive timing. The Type 2 measurement is equally applicable to Release 8 UEs as to later UEs.

The distance of the UE from the eNodeB or cell is estimated as $d = c \times \text{RTT}/2$, where c is the speed of light. The accuracy of the Type 1 measurement is mainly determined by the receive timing accuracy of the eNodeB and UE, which is typically of the order of $0.3 \mu\text{s}$, corresponding to a distance of around 45 m. The accuracy of the Type 2 measurement depends on the PRACH detection accuracy, which is typically of the order of $1\text{--}2 \mu\text{s}$ depending on the multipath characteristics of the radio channel (see [10]). Thus the Type 1 measurement can potentially provide better positioning accuracy.

In addition to the receive–transmit time difference measurement, the UE may be configured to perform Reference Signal Received Power (RSRP) and Reference Signal Received Quality (RSRQ) (see Sections 22.3.1.1 and 22.3.1.2 respectively) measurements for the serving cell and for neighbouring cells. These measurements can be used to help enhance CID-based positioning by enabling an RSRP/RSRQ map to be compiled throughout the network coverage area. The estimate of a UE’s location can then be enhanced by comparing its reported RSRP/RSRQ with the RSRP/RSRQ map. The usefulness of this information depends on the accuracy of the UE’s RSRP/RSRQ reports and on the accuracy of the RSRP/RSRQ map.

19.4.3 Enhanced CID Positioning using Round Trip Time and Angle of Arrival

Although the RTT can be used to estimate distance, it does not provide directional information for the UE location. Angle of Arrival (AoA) [9] is defined as the estimated angle of a UE with respect to a reference direction which is defined as geographical North; the value of AoA is positive in an anticlockwise direction.

The eNodeB estimates the AoA using uplink transmissions from the UE. The eNodeB antenna configuration is a key factor for AoA estimation. With a linear array of equally spaced antenna elements, the received signals at any two adjacent elements are phase-rotated by a fixed amount θ . The value of θ is a function of the AoA as well as the antenna element spacing and carrier frequency. The latter parameters are known by the eNodeB, and therefore the AoA can be estimated. In general, any uplink signal from the UE can be used to estimate the AoA, but typically a known signal such as the Sounding Reference Signals (SRSs) or DeModulation Reference Signals (DM-RS) would be used (see Sections 15.6 and 15.5 respectively). With knowledge of both the AoA and the RTT, a UE’s position can be estimated [11].

Figure 19.7 shows an example of estimated UE positions with AoA+RTT and GPS for a cellular deployment in Guangzhou, China, where positioning based on AoA and RTT

estimation is used in a TD-SCDMA network. Typically, AoA+RTT based positioning can provide an accuracy of 150 m and 50 m in Non-LOS and LOS environments respectively [11, 12].




Figure 19.7: AoA+RTT positioning results compared to GPS in a TD-SCDMA deployment in Guangzhou, China.

19.5 LTE Positioning Protocols

Figure 19.8 illustrates the signalling connections for the location servers for LTE UE positioning. For the Control Plane (CP), the E-SMLC is shown, while for the User Plane (UP) the SUPL Location Platform (SLP) is shown. The E-SMLC–UE communication is tunnelled through the MME and eNodeB in the CP protocols and the E-SMLC–eNodeB communication is achieved by terminating the tunnel in the eNodeB. The SLP interacts directly with the UE client using normal UP data transport and the same protocol as is used for the E-SMLC–UE communication. It should be noted that the E-SMLC and SLP are logical entities and may therefore be implemented in a single physical server.

Figure 19.9 shows the protocol stack used to carry the LTE Positioning Protocol (LPP) [13] between the UE and the E-SMLC. LPP is a point-to-point protocol, and naturally the E-SMLC can have a number of parallel connections to different UEs. The LPP messages can be divided into four functional groups:

- UE positioning capability information transfer to the E-SMLC;
- Positioning assistance data delivery from the E-SMLC to the UE;
- Location information transfer;
- Session management (for procedure abortion and error handling).

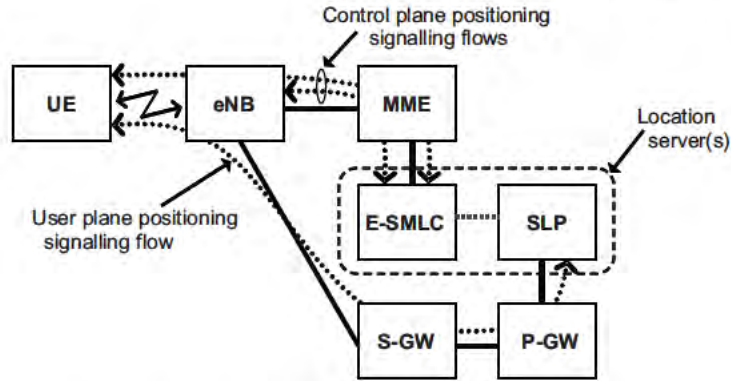


Figure 19.8: CP and UP architecture for positioning.




Figure 19.9: The LPP stack between the E-SMLC and the UE.

Reproduced by permission of © 3GPP.

The underlying protocols and the nodes between the eNodeB and the MME act solely as transport and do not take part in the LPP procedures. This means that LPP is a stand-alone protocol that can use any UE-server delivery mechanism. In particular, SUPL can use the LPP as is and can use normal UP data transport to carry LPP instead of the CP protocols. This is especially beneficial in deployments where both CP and UP positioning solutions are deployed together. The same server can support both architectures and directly use the same protocol for positioning. Similarly, UEs supporting both solutions need to implement only one positioning protocol.

In addition to supporting the 3GPP-defined positioning mechanisms for LTE, LPP includes an 'External Protocol Data Unit (EPDU)', which provides a transparent message exchange mechanism between the UE and the server for the support of additional positioning techniques that may be defined in the future.

In addition to the LPP protocol used in E-SMLC-UE communication, the LTE Positioning Protocol A (LPPa) [14] is defined for E-SMLC-eNodeB communication. Figure 19.10 shows the LPPa protocol stack. LPPa is used in Enhanced CID positioning if eNodeB measurements

such as AoA are used to aid the location determination. LPPa can also be used for the E-SMLC to request details of the eNodeB configuration for OTDOA, which is needed when constructing the OTDOA assistance data to be delivered to the UE. Examples of these parameters are the cell IDs and transmission frequencies, the PRS configuration and the cell timings. Like LPP, LPPa is also transparent to the MME and is tunnelled through it.




Figure 19.10: The LPPa protocol stack between the E-SMLC and the eNodeB. Reproduced by permission of © 3GPP.

19.6 Summary and Future Techniques

Whatever the motivation for deploying the capability to locate a UE, be it regulatory or commercial, the LTE Release 9 specifications provide suitable methods. Different techniques inherently have varying degrees of accuracy, availability and cost of deployment, so a careful case-by-case assesment is needed to identify which techniques should be deployed in a given environment.

A-GNSS is assumed to be the primary positioning technique for LTE UEs, providing very good positioning accuracy and reasonable availability. However, satellite-based techniques alone are not sufficient to provide a location estimate in all cases, due to satellite hearability problems in dense urban and indoor situations.

OTDOA provides a useful alternative positioning technique to A-GNSS in environments where the satellite signals cannot be received but a sufficiently dense deployment of eNodeBs exists to provide a good probability of OTDOA yielding a location estimate. The probability of obtaining a location estimate by OTDOA can be increased by means of Positioning Reference Signals. OTDOA accuracy depends on the deployment density of the eNodeBs and the accuracy with which the antenna locations and their transmission timings are known by the network.

Cell-ID-based techniques give a reasonably good estimate of the UE location, especially in urban environments with dense eNodeB deployments, and can further benefit from UE measurements of the radio environment as well as eNodeB AoA measurements if suitable antenna arrays are deployed. The smaller the cells, the more accurate the positioning estimate. A key benefit of Cell-ID-based techniques is the fact that if the UE is able to connect to a network then the technique is guaranteed to provide a location estimate.

The UE-to-server positioning communications can take place either via the UP (as with other normal client–server application-layer communications), or via the CP in a similar fashion to RRC, NAS or S1AP signalling.

Further development of LTE in Release 11 is expected to add additional positioning techniques to the range already available, namely Uplink Time Difference of Arrival (UTDOA) and RF pattern matching. UTDOA uses eNodeB measurements of uplink signal reception timings at different cells to estimate the location of a UE without any specific involvement of the UE, while RF pattern matching aims to locate a UE by comparing any available radio channel measurements with a database.

References⁹

- [1] Federal Communications Commission Rules, Section 20.18(h), www.fcc.gov/oet/info/rules.
- [2] 3GPP Technical Specification 36.305, ‘Evolved Universal Terrestrial Radio Access Network (E-UTRAN); Stage 2 functional specification of User Equipment (UE) positioning in E-UTRAN (Release 9)’, www.3gpp.org.
- [3] A. Kupper, *Location-Based Services: Fundamentals and Operation*. John Wiley & Sons Ltd/Inc., 2005.
- [4] Y. Chan and K. Ho, ‘A simple and efficient estimator for hyperbolic location’. *IEEE Transactions on Signal Processing*, Vol. 42, pp. 1905–1915, August 1994.
- [5] Motorola, ‘R1-091336: Study on hearability of reference signals in LTE positioning support’, www.3gpp.org, 3GPP TSG RAN WG1, meeting 56bis, Seoul, South Korea, March 2009.
- [6] 3GPP Technical Specification 36.211, ‘Evolved Universal Terrestrial Radio Access (E-UTRA); Physical Channels and Modulation (Release 9)’, www.3gpp.org.
- [7] 3GPP Technical Specification 36.133, ‘Evolved Universal Terrestrial Radio Access (E-UTRA); Requirements for support of radio resource management (Release 9)’, www.3gpp.org.
- [8] Alcatel-Lucent, ‘R1-092307: Analysis of UE Subframe Timing Offset Measurement Sensitivity to OTDOA Performance’, www.3gpp.org, 3GPP TSG RAN WG1, meeting 57bis, Los Angeles, USA, June 2009.
- [9] 3GPP Technical Specification 36.214, ‘Evolved Universal Terrestrial Radio Access (E-UTRA); Physical layer - Measurements (Release 9)’, www.3gpp.org.
- [10] 3GPP Technical Specification 36.104, ‘Evolved Universal Terrestrial Radio Access Network (E-UTRAN); Base Station (BS) radio transmission and reception (Release 9)’, www.3gpp.org.
- [11] CATT, ‘R1-090936: UE Positioning Based on AoA+TA for LTE Rel-9’, www.3gpp.org, 3GPP TSG RAN WG1, meeting 56, Athens, Greece, February 2009.
- [12] RITT, CATT, ‘R1-091979: Further performance evaluation of UE positioning based on AoA+TA’, www.3gpp.org, 3GPP TSG RAN WG1, meeting 57, San Francisco, USA, May 2009.
- [13] 3GPP Technical Specification 36.355, ‘Evolved Universal Terrestrial Radio Access (E-UTRA); LTE Positioning Protocol (LPP) (Release 9)’, www.3gpp.org.
- [14] 3GPP Technical Specification 36.455, ‘Evolved Universal Terrestrial Radio Access (E-UTRA); LTE Positioning Protocol A (LPPa) (Release 9)’, www.3gpp.org.

⁹All web sites confirmed 1st March 2011.

The Radio Propagation Environment

Juha Ylitalo and Tommi Jämsä

20.1 Introduction

Realistic modelling of propagation characteristics is essential for two main reasons in relation to the LTE system. Firstly, the link- and system-level performance of LTE can be evaluated accurately only when the radio channel models are realistic. In particular, the spatial characteristics of the channel models have a significant effect on the performance of multi-antenna systems. Secondly, the model used for radio propagation plays an important role in the network planning phase of LTE deployment.

Different environments such as rural, suburban, urban, and indoor set different requirements for network planning, antenna configuration, and the preferred spatial transmission mode. Moreover, the actual deployment scenarios for LTE cover numerous special cases ranging from mountainous rural surroundings to dense urban and outdoor-to-indoor situations. As an example of the range of different propagation environments of relevance to high-bandwidth Multiple-Input Multiple-Output (MIMO) systems such as LTE, the IST-WINNER¹ project [1] developed wideband radio channel models for 12 propagation environments.

The propagation characteristics are in a large part affected by the carrier frequency. In November 2007 the International Telecommunication Union (ITU) World Radiocommunication Conference (WRC-2007) allocated new frequency bands for International Mobile Telecommunications (IMT) radio access systems between 450 MHz and 3.6 GHz. Thus depending on the deployed carrier frequency, the LTE radio channel characteristics will vary

¹Information Society Technologies – Wireless world INitiative NEw Radio.

LTE – The UMTS Long Term Evolution: From Theory to Practice, Second Edition.

Stefania Sesia, Issam Toufik and Matthew Baker.

© 2011 John Wiley & Sons, Ltd. Published 2011 by John Wiley & Sons, Ltd.

significantly even within a particular type of propagation environment. The carrier frequency used also has an impact on the deployment of MIMO for LTE, due to the fact that the size of the antenna arrangement depends strongly on the carrier wavelength.

As discussed in Chapter 11, the ability to use a variety of MIMO techniques is an important feature of LTE. While conventional wireless systems are designed to counteract multipath fading, MIMO approaches such as those used in LTE are able to take advantage of multipath scattering to increase capacity [2, 3]. Therefore, it is important to create realistic standardized MIMO radio channel models for the evaluation of the performance of LTE and its future enhancements.

It is especially important to model accurately the correlation between the signals of different antenna branches, since this dictates, to a large extent, the preferred spatial transmission mode and its performance in practice. The spatial correlation properties of the MIMO radio channel model define the ultimate limit of the theoretical channel capacity. The applied channel model has to reflect all the instantaneous space-time-frequency characteristics which affect the configuration of diversity, beamforming, and spatial multiplexing techniques. The model also has to support all the antenna designs that are of practical interest.

In the following sections we discuss first the 3GPP Single-Input Single-Output (SISO) and Single-Input Multiple-Output (SIMO) channel models which include only a single antenna at the transmitter and one or two antennas at the receiver respectively. This is followed by a detailed analysis and comparison of the characteristics and modelling principles of a MIMO radio channel, focusing in particular on the models used in 3GPP, up to and including the latest models defined by ITU-R for IMT-Advanced systems.

20.2 SISO and SIMO Channel Models

In practice, realistic modelling of the radio channel propagation characteristics requires extensive measurement campaigns with appropriate carrier frequencies and bandwidths, in the radio environments which are identified to be the most relevant for deployment. The research projects COST207, COST231, COST259, CODIT (UMTS Code Division Testbed), and ATDMA (Advanced TDMA Mobile Access) created extensive wideband measurement datasets for SISO and SIMO channel modelling in the 1980s and 1990s [4–8]. The corresponding channel models form a basis for the ITU models which were largely applied in the development of the third generation wireless communication systems.

Figure 20.1 illustrates a multipath propagation scenario in which the UE has an approximately omnidirectional antenna. The transmitted signal traverses three paths with different delays. As explained in Chapter 5, in wideband communications the delay spread of the propagation paths is larger than the symbol period and the receiver observes the multipath components separately. In the frequency domain this corresponds to frequency-selective fading since the coherence bandwidth of the radio channel is smaller than the signal bandwidth. The coherence bandwidth is proportional to the inverse of the root mean square (r.m.s.) delay spread which can be calculated from the Power Delay Profile (PDP) of the radio channel (see Section 8.2.1). The PDP also defines the *maximum excess delay*, which is a measure of the largest delay difference between the propagation paths. This measure is important when considering the length of the Cyclic Prefix (CP) in LTE, as discussed in Section 5.4.




Figure 20.1: Multipath propagation and PDP.

Three main phenomena can be identified which affect the received signal properties:

- Propagation path-loss;
- Shadow (slow, or ‘large-scale’) fading;
- Multipath (fast, or ‘small-scale’) fading.

The propagation loss determines the average signal level at the receiver. Typical radio access systems can cope with a path-loss up to around 150 dB. The received signal level varies relatively slowly due to shadowing effects from large obstructions such as buildings and hills. Finally, the multipath characteristics define the frequency, time, and space correlations. As several multipath signal components with approximately the same propagation delay add up incoherently at the receiver antenna, the composite signal level can vary by up to 30–40 dB as the User Equipment (UE) moves. This variation can be very fast since the distance between two consecutive fades can be of the order of half wavelength.

20.2.1 ITU Channel Model

ITU channel models were used in developing the Third Generation ‘IMT-2000’ family² of radio access systems [9]. The main user scenarios covered indoor office, outdoor-to-indoor, pedestrian and vehicular radio environments. The key parameters to describe each propagation model include time delay spread and its statistical variability, path-loss and shadow fading characteristics, multipath fading characteristics, and operating radio frequency. Each environment is defined for two cases which have different probabilities of occurrence: a smaller delay spread case and a larger delay spread case. The numbers of paths are as for the 3GPP Channel Models introduced in the following section.

²The initial ‘IMT-2000’ systems of the IMT family were proposed with the year 2000 in view.

20.2.2 3GPP Channel Model

The 3GPP specifications for UMTS made use of propagation conditions largely based on the ITU models. As an example, Table 20.1 shows the propagation conditions used for performance evaluation in different multipath fading environments [10]. All taps use the classical Doppler spectrum [11, 12] (see Section 8.3.1), which assumes that the propagation paths at a receiver antenna are uniformly distributed over 360° in azimuth. The fading of the signals at different antennas is assumed to be independent, which is only appropriate for SISO and some antenna diversity cases. MIMO was not introduced in UMTS until Release 7.

Table 20.1: Propagation conditions for multipath fading in the ITU models.

Tap number	ITU Pedestrian A		ITU Pedestrian B		ITU Vehicular A	
	Relative delay (ns)	Relative mean power (dB)	Relative delay (ns)	Relative mean power (dB)	Relative delay (ns)	Relative mean power (dB)
1	0	0	0	0	0	0
2	110	-9.7	200	-0.9	310	-1.0
3	190	-19.2	800	-4.9	710	-9.0
4	410	-22.8	1200	-8.0	1090	-10.0
5			2300	-7.8	1730	-15.0
6			3700	-23.9	2510	-20.0

20.2.3 Extended ITU Models

The evaluation of LTE techniques demands channel models with increased bandwidth compared to UMTS models, to reflect the fact that the characteristics of the radio channel frequency response are connected to the delay resolution of the receiver. In 3GPP the 20 MHz LTE channel models were based on a synthesis of existing models such as the ITU and 3GPP models. Specifically the six ITU models covering an excess (maximum) delay spread from 35 ns to 4000 ns were chosen as a starting point, together with the Typical Urban (TU) model from GSM³ which has a maximum excess delay of 5000 ns. In this way, extended wideband models with low, medium, and large delay spread values could be identified. The low delay spread gives an Extended Pedestrian A (EPA) model which is employed in an urban environment with fairly small cell sizes (or even up to about 2 km in suburban environments with low delay spread), while the medium and large delay spreads give an Extended Vehicular A (EVA) model and Extended TU (ETU) model respectively. The ETU model has a large maximum excess delay of 5000 ns, which in fact is not very typical in urban environments but it applies to some extreme urban, suburban, and rural cases which occur seldom but which are important in evaluating LTE performance in the most challenging environments. Table 20.2 shows the r.m.s. delay spread values for the three extended models [13].

It was also decided that the extended channel models are applied with low, medium, and high Doppler shifts, namely 5 Hz, 70 Hz and 300 Hz, which at a 2.5 GHz carrier frequency

³Global System for Mobile Communications.

Table 20.2: r.m.s. delay spread for the extended ITU models.

Category	Channel model	Acronym	r.m.s. delay spread (ns)
Low delay spread	Extended Pedestrian A	EPA	43
Medium delay spread	Extended Vehicular A model	EVA	357
High delay spread	Extended Typical Urban model	ETU	991

correspond roughly to mobile velocities of 2, 30 and 130 km/h respectively. Combinations which are likely to be used are EPA 5 Hz, EVA 5 Hz, EVA 70 Hz and ETU 70 Hz. The classical Doppler spectrum is again assumed. The tapped delay line propagation conditions for the LTE performance evaluation are hence summarized in Table 20.3.

Table 20.3: Power delay profiles of extended ITU models.

Tap number	EPA model		EVA model		ETU model	
	Excess tap delay (ns)	Relative power (dB)	Excess tap delay (ns)	Relative power (dB)	Excess tap delay (ns)	Relative power (dB)
1	0	0.0	[0	0.0	0	-1.0
2	30	-1.0	30	-1.5	50	-1.0
3	70	-2.0	150	-1.4	120	-1.0
4	80	-3.0	310	-3.6	200	0.0
5	110	-8.0	370	-0.6	230	0.0
6	190	-17.2	710	-9.1	500	0.0
7	410	-20.8	1090	-7.0	1600	-3.0
8			1730	-12.0	2300	-5.0
9			2510	-16.9	5000	-7.0

LTE conformance tests also include propagation conditions for two high-speed train scenarios. These are based on 300 and 350 km/h, with different Doppler shift trajectories, in non-fading propagation channels [14].

20.3 MIMO Channel Models

As already mentioned, the performance of MIMO systems depends strongly on the underlying propagation conditions. First of all, the Signal-to-Interference-plus-Noise Ratio (SINR) at the MIMO receiver determines the ultimate gains of spatial multiplexing. The SINR is a direct function of the path-loss, so that a terminal far from the base station tends to have a relatively small SINR compared to a terminal close to the base station. A second important factor is the correlation between signals at different antennas. Fading decorrelation is facilitated by the presence of a large number of multipath components at both the transmitter and the receiver, as typically experienced in a Non-Line-Of-Sight (NLOS) situation. The

antenna separation also has a strong impact on the spatial correlation. The largest MIMO gains are obtained in scenarios with large SINR and low spatial correlation (where the channel matrix has a high rank), which in practice may rarely occur. For a detailed analysis of the effect of spatial correlation on MIMO system performance, the interested reader is referred to [2, 3] and references therein.

There are two commonly applied types of MIMO channel model:

- Correlation matrix based channel models;
- Geometry-based channel models.

Models of both types are used for LTE. The 3GPP Spatial Channel Model (SCM) and Spatial Channel Model – Extension (SCME) (see Sections 20.3.1 and 20.3.2 respectively) are geometry-based stochastic models, whereas the extended ITU models (EPA, EVA, ETU) are correlation matrix based models. In Section 20.3.3 we then discuss the geometry-based IST-WINNER channel model to give an insight into the state of the art in channel modelling. It should be noted that it is possible to create a correlation matrix based channel model from a geometry-based one, but not vice versa.

20.3.1 SCM Channel Model

In order to model MIMO channel characteristics realistically, originally for the evaluation of different MIMO schemes for High Speed Downlink Packet Access (HSDPA), 3GPP developed jointly with 3GPP2 a geometry-based SCM [15]. The SCM includes simple tapped-delay line models for calibration purposes and a geometry-based stochastic model for system-level simulations. The time-delay properties of the SCM calibration model are the same as in the 3GPP SISO/SIMO models discussed in Section 20.2.2. The spatial characteristics of the MIMO channel are defined by the Angular Spread (AS) and directional distribution of the propagation paths at both the base station and the UE.

System simulations typically consist of multiple cells/sectors, multiple base stations and multiple UEs. Performance metrics such as throughput and delay are gathered over a large number of simulation runs, called ‘drops’, which may consist of a predefined number of radio frames. During a drop, the large-scale parameters are fixed but the channel undergoes fast fading according to the motion of the terminals. The UEs may feed back channel state information about the instantaneous radio channel conditions and the base station can schedule its transmissions accordingly. Figure 20.2 defines the geometrical framework for the spatial parameters.

The SCM includes three main propagation environments: urban microcell, urban macrocell, and suburban macrocell. Additionally, it is possible to modify the basic scenarios by applying a LOS component in an urban micro case, and adding a far-scattering cluster or an urban canyon in an urban macro case.

The stochastic geometry-based model of the SCM enables realistic modelling of spatial correlation at both the base station and the UE. It is based on the ITU SISO models described in Section 20.2.1, so the number of propagation paths in each environment is six. However, the delay and angular spreads, as well as the directions of the scattering clusters, are random variables with normal or log-normal distribution. The elevation (i.e. vertical) domain is not modelled. Table 20.4 shows the main SCM parameters.




Figure 20.2: Geometry for generating spatial parameters for base station and UE.

Table 20.4: Main parameters for SCM.

Parameter	Suburban macro	Urban macro	Urban micro
Number of paths	6	6	6
Number of subpaths per path	20	20	20
Mean AS at base station	5°	8°, 15°	19°
Per-path AS at base station (fixed)	2°	2°	5°
Mean AS at UE	68°	68°	68°
Per-path AS at UE (fixed)	35°	35°	35°
Mean total r.m.s. delay spread	0.17 μs	0.65 μs	0.25 μs
Standard deviation for log-normal shadowing	8 dB	8 dB	NLOS: 10 dB; LOS: 4 dB
Path-loss model (dB, d = distance in metres from base station to UE)	$31.5 + 35 \log_{10}(d)$	$34.5 + 35 \log_{10}(d)$	NLOS: $34.53 + 38 \log_{10}(d)$; LOS: $30.18 + 26 \log_{10}(d)$

The SCM model is designed so that different antenna constellations can be applied. For example, the base station antenna spacing and beam pattern can be varied, including beam patterns for 3-sector and 6-sector cells. Typical values for the antenna spacing at the base

station vary from 0.5 to 10 carrier wavelengths. At the base station the per-path angular spread is rather small ($<5^\circ$) in all three environments, which is typical of antennas well above rooftop level. The mean angular spread is also relatively small, the largest value being 19° in the urban micro scenario, which is small compared to the beamwidth of a 4-antenna array. At the UE, the per-path and mean angular spreads are large (35°), which reflects the fact that the UE antennas are embedded in the scattering environment. The angular spread is composed of 6×20 subpaths. The per-path angular spread at both the base station and the UE follows a Laplacian distribution, which is generated by giving the 20 subpaths an equal power and fixed azimuth directions with respect to the nominal direction of the corresponding path.

A Matlab[®] implementation of the SCM model can be found in [1].

20.3.2 SCM-Extension Channel Model

An extension to the 3GPP/3GPP2 SCM model was developed in the IST-WINNER project. Known as the SCME (SCM-Extension) model, it is described in [16]. The extension was designed to remain backward-compatible, simple and consistent with the initial SCM concept. The main development is a broadening of the channel model bandwidth from 5 MHz to 100 MHz.

The bandwidth extension is carried out by introducing so-called midpaths which define the intra-cluster delay spread (i.e. the delay spread within a cluster of paths in a similar direction). The midpaths have fixed delay and power offsets in order to keep the SCME model backward-compatible with SCM. Therefore, the low-pass-filtered SCME impulse response corresponds closely to the respective SCM impulse response. As a result of the bandwidth extension the number of delay taps increases from six in SCM to 18 or 24, depending on the scenario. Table 20.5 shows the midpath delays and powers for the SCME propagation scenarios. Frequency-dependent factors are also added to the path-loss formulae, to extend the frequency range from 2 GHz to 5 GHz.

Table 20.5: SCME midpath powers and delays.

Scenario	Urban macro, suburban macro		Urban micro	
	No. of midpaths per path			
	Midpath power	Relative delay (ns)	Midpath power	Relative delay (ns)
Midpath 1	10/20	0	6/20	0
Midpath 2	6/20	7	6/20	5.8
Midpath 3	4/20	26.5	4/20	13.5
Midpath 4	—	—	4/20	27.6

Another important contribution of the SCME model is the introduction of fixed spatio-temporal Tapped Delay Line (TDL) models, referred to as ‘Clustered Delay Line (CDL) models’. The model parameters include the arrival and departure angles in addition to the traditional power and delay, thus covering all the MIMO propagation channel parameters

except the polarization information. The TDL models are intended for the calibration of simulators. The spatio-temporal TDL models of the SCM extension resemble closely the SCM system model; they can actually be viewed as a realization of the system model which is optimized for low correlation in the frequency domain.

The SCME model has a number of optional features which can be applied depending on simulation purposes. In addition to the SCM ‘drop’ concept, the SCME model introduces optional drifting of delays and arrival/departure angles of the propagation paths. In the original SCM all the propagation parameters remain fixed during a drop and the only variation is the fast fading caused by the Doppler effects. Time evolution of the delay taps may be used for example for evaluating receiver synchronization algorithms. Drifting of the path arrival/departure angles is targeted to the testing of beamforming algorithms. Time evolution of shadow fading is also an optional feature; since the correlation of the shadow fading at two locations is related to the distance between them, this is modelled by an exponentially-shaped spatial auto-correlation function applied to the shadow fading such that the correlation of the shadow fading reduces exponentially with distance. Finally, the SCME model includes parameterization of the LOS condition for all the macrocellular scenarios.

A Matlab® implementation of the SCM model can be found in [1].

20.3.3 WINNER Model

The WINNER radio channel models were developed in the two phases of the IST-WINNER project [1]. Since the initial SCME model was not adequate for more advanced simulations, new measurement-based models were developed.

The models are ray-based double-directional multilink models which are antenna-independent, scalable and capable of creating arbitrary geometry-based MIMO channel models. Statistical distributions and channel parameters extracted from measurements in any propagation scenarios can be fitted to the models.

The latest model, known as the Phase II model, forms the basis for the ITU IMT-Advanced models introduced in Section 20.3.6.1. It extends the frequency range to cover frequencies from 2 to 6 GHz and covers 13 propagation scenarios including indoor, outdoor-to-indoor (and vice versa), urban micro- and macrocell and corresponding difficult urban scenarios, suburban and rural macrocell, feeder links, and moving networks. It is possible to vary the number of antennas, the antenna configurations, the geometry and the antenna beam pattern without changing the basic propagation model. This method enables the use of the same channel data in different link- and system-level simulations. The details of the model are described in [17]. The models also include CDL models for system calibration.

Path-losses have been specified for all the scenarios, divided into two subscenarios according to propagation conditions (LOS/NLOS) when applicable. The general structure of the path-loss is given by

$$PL = A \log_{10}(d) + B + C \log_{10}\left(\frac{f_c}{5}\right) + X \quad (20.1)$$

where d is the distance in metres and f_c is the carrier frequency in GHz. For each (sub)scenario, a set of constants A , B and C are defined in [17], as is the additional term X , for special cases. X is an environment-dependent factor, for example modelling additional propagation attenuation due to walls or floors (e.g. if the transmitter and receiver are located on different floors). A Matlab® implementation of the WINNER model can be found in [1].

20.3.4 LTE Evaluation Model

Channel models have been used in the design of LTE for many purposes, including evaluation of UE and eNodeB performance requirements, Radio Resource Management (RRM) requirements (see Chapter 22), system concepts and RF system scenarios (see Chapters 21 and 23). Conformance tests of the UE and eNodeB are discussed in Section 20.4.1.

Fixed TDL models of the SCME were proposed for both link- and system-level simulations in [18]. Two simplification steps from the SCM approach were taken. Firstly, the fixed TDL models, which were originally defined for calibration simulations only, were planned to be applied also in system-level simulations. The stochastic drop concept of SCM was discarded. As a slight modification to the SCME TDL models, the delays were quantized to the resolution which corresponds to fourfold over-sampling of the LTE bandwidth. Secondly, the MIMO characteristics were described by correlation matrices instead of directional propagation parameters such as the angles of arrival/departure and the polarization characteristics. For this purpose, antenna array characteristics were defined for the eNodeB and UE for different scenarios. The LTE channel model scenarios and antenna configurations are listed in Table 20.6 where λ is the carrier wavelength.

Table 20.6: Scenarios of LTE performance evaluation model.

Name	Propagation scenario	eNodeB arrangement	UE arrangement
SCM-A	Suburban Macro	3-sector, 0.5λ spacing	Handset, talk position
SCM-B	Urban Macro (low spread)	6-sector, 0.5λ spacing	Handset, data position
SCM-C	Urban Macro (high spread)	3-sector, 4λ spacing	Laptop
SCM-D	Urban Micro	6-sector, 4λ spacing	Laptop

In the eNodeB, a cross-polarized antenna configuration is used with four antenna elements. Two dual-polarized slanted $+45^\circ$ and -45° elements are spatially separated by a distance d , as shown in Figure 20.3. The polarization of each element is for simplicity assumed to be unchanged over all departure angles. The antenna element patterns are identical to those proposed for the link calibration model of SCM and can be either 3-sector or 6-sector. The radiation pattern as a function of angle, $A(\theta)$, is as follows:

$$A(\theta) = -\min\left[12\left(\frac{\theta}{\theta_{3\text{ dB}}}\right)^2, A_m\right] \quad \text{where } -180^\circ \leq \theta \leq 180^\circ \quad (20.2)$$

where, for a 3-sector arrangement, the 3 dB beamwidth $\theta_{3\text{ dB}} = 70^\circ$, $A_m = 20$ dB and the maximum gain is 14 dBi,⁴ while for a 6-sector arrangement $\theta_{3\text{ dB}} = 35^\circ$, $A_m = 23$ dB and the maximum gain is 17 dBi.

Two types of mobile antenna scenarios are assumed: a small handset with two orthogonally-polarized antennas (see Figure 20.4), and a laptop with two spatially separated dual-polarized antennas (see Figure 20.5). It is assumed that the polarizations are purely horizontal and vertical in all directions when the antennas are in the nominal position. In the ‘talk position’

⁴This indicates the maximum forward gain of the antenna relative to the gain of a hypothetical omnidirectional antenna.




Figure 20.3: Base station antenna configuration.

case, the lobe is in the horizontal direction and the handset is rotated 60° (with the polarizations also being rotated). In the data position, the mobile is tilted 45° such that the radiation lobe has its maximum slanted downwards. The azimuthal orientations of the mobile antennas are selected such that the angle of arrival of the first tap occurs at +45° in all scenarios. The antenna element patterns are given by the same function as the eNodeB patterns, with the following parameters:




Figure 20.4: Handset antenna configuration.

- Handheld, talk position: $\theta_{3\text{ dB}} = 120^\circ$, $A_m = 15\text{ dB}$, maximum gain = 3 dBi (vertical), 0 dBi (horizontal);
- Handheld, data position: $\theta_{3\text{ dB}} = 120^\circ$, $A_m = 5\text{ dB}$, maximum gain = 3 dBi (vertical), 0 dBi (horizontal);
- Laptop: $\theta_{3\text{ dB}} = 90^\circ$, $A_m = 10\text{ dB}$, maximum gain = 7 dBi, spatial separation 2λ .

Spatial correlations of the LTE performance evaluation model are calculated based on the antenna characteristics and the path arrival/departure angles defined in the SCME model. Polarization covariance matrices are determined based on the antenna polarizations and a cross polarization discrimination value of 8 dB. The Kronecker assumption is applied – i.e. fully separable transmitter and receiver power azimuth spectra. Finally, 8×8 correlation matrices can be determined for the SCM-A and -B scenarios, and 16×16 correlation matrices for the SCM-C and -D scenarios (see Table 20.6).




Figure 20.5: Laptop antenna configuration.

20.3.5 Extended ITU Models with Spatial Correlation

The extended ITU models introduced in Section 20.2.3 (EPA, EVA, ETU) were developed for MIMO conformance tests in a simple way by allocating the same correlation matrix to all the multipath components [20]. Three degrees of spatial correlation between different antenna signals are defined: high, medium and low. The low correlation case actually represents a case where the antennas are fully uncorrelated. The high and medium correlations are defined by specific correlation matrices based on practical antenna constellations. In the high correlation case the eNodeB and UE antenna arrays are co-polarized and have small inter-antenna distances of 1.5 and 0.5 wavelengths respectively. Medium correlation is achieved by means of orthogonally polarized antennas. These cases are illustrated in Figure 20.6.




Figure 20.6: Antenna configurations for medium and high spatial correlation for extended ITU models.

In the 2×2 MIMO case the base station and UE correlation matrices are defined respectively as follows:

$$\mathbf{R}_{BS} = \begin{pmatrix} 1 & \alpha \\ \alpha^* & 1 \end{pmatrix}, \quad \mathbf{R}_{UE} = \begin{pmatrix} 1 & \beta \\ \beta^* & 1 \end{pmatrix} \quad (20.3)$$

and accordingly the spatial correlation is defined as

$$\mathbf{R}_{\text{spat}} = \mathbf{R}_{\text{BS}} \otimes \mathbf{R}_{\text{UE}} = \begin{pmatrix} 1 & \alpha \\ \alpha^* & 1 \end{pmatrix} \otimes \begin{pmatrix} 1 & \beta \\ \beta^* & 1 \end{pmatrix} \quad (20.4)$$

where \otimes is the Kronecker product. The spatial correlation in the different cases can be represented as shown in Table 20.7.

Table 20.7: Spatial correlation with extended ITU channel models.

Low correlation		Medium correlation		High correlation	
α	β	α	β	α	β
0	0	0.3	0.9	0.9	0.9

Further, by applying the polarization matrices for the medium and high correlation cases, the 4×4 spatial correlation matrices can be obtained for the 2×2 MIMO case [21]. Similar spatial channel models for the 4×2 and 4×4 MIMO cases are also specified. The above correlation models are relatively simple. They apply reasonably well to basic MIMO concepts but may not be adequate for advanced beamforming/MIMO concepts due to the fact that all the delay taps have the same correlation properties. In real environments each delay tap will have different directions and angular spreads which define the per-path spatial correlation properties.

20.3.6 ITU Channel Models for IMT-Advanced

In November 2008, ITU-R approved radio channel models for the evaluation of the performance of the candidate technologies for IMT-Advanced, including LTE-Advanced [22, 23].

20.3.6.1 Test Environments

ITU-R has defined four representative test environments which are designed to span sufficiently the range of possible radio environments [22]:

- Indoor: a model for isolated indoor cells (e.g. in offices or hotspots), focusing on stationary and pedestrian users, with high user throughput or high user density.
- Microcellular: an urban model for high traffic loads, focusing on high densities of pedestrian and slow vehicular users. This scenario also includes outdoor-to-indoor coverage with the base station antennas situated below rooftop level.
- Base coverage urban: a macrocellular model providing continuous coverage for users with pedestrian to fast vehicular mobility. The base station antennas are above rooftop level.
- High speed: an environment of large rural macrocells with users in high speed vehicles and trains.

The IMT-Advanced channel modelling approach includes a primary module and an extension module.

The primary module is based on the WINNER Phase II channel model, and it covers the essential parameters needed for evaluation of the IMT-Advanced candidate radio interface technologies. The primary module includes path-loss models, a generic geometry-based stochastic channel model, and optional CDL models. The CDL models are used for calibration purposes only, under which all the radio channel parameters except for the phase angles of the propagation path signals are fixed. The extension module is based on the Time-Spatial Propagation (TSP) model described in [24].

The ITU-R IMT-Advanced channel model is a geometry-based stochastic model which is targeted at multi-cell, multi-user system simulations. Similarly to the 3GPP/3GPP2 SCM it does not explicitly specify the positions of the scatterers, but instead, models the directions of the propagation paths. It allows separate definition of the antenna constellation and radiation patterns. The channel parameters for individual snapshots of the radio channel are determined stochastically, based on statistical distributions derived from channel measurements. Channel realizations are generated by summing contributions of individual propagation paths with specific small scale parameters like delay, power, Angle-of-Arrival (AoA) and Angle-of-Departure (AoD). The IMT-Advanced channel model is based on a similar drop concept as used in the SCM model. Some of the main parameters of the IMT-Advanced channel model are listed in Table 20.8.

Table 20.8: Some main parameters of the IMT-Advanced channel models [22].

Test Environment	Indoor	Microcellular	Base coverage urban	High speed
Deployment scenario	Indoor Hotspot ('InH')	Urban Micro ('UMi')	Urban Macro ('UMa')	Rural Macro ('RMa')
Network layout	Indoor floor	Hexagonal grid	Hexagonal grid	Hexagonal grid
Site-to-site distance	60 m	200 m	500 m	1732 m
Carrier frequency	3.4 GHz	2.5 GHz	2 GHz	800 MHz
BS antenna height	6 m, ceiling-mounted	10 m, below rooftop	25 m, above rooftop	35 m, above rooftop
BS antenna gain	0 dBi	17 dBi	17 dBi	17 dBi
No of BS antennas	up to 8	up to 8	up to 8	up to 8
BS Tx power	21 dBm/ 20 Mhz 24 dBm/ 40 Mhz	41 dBm/ 10 Mhz 44 dBm/ 20 Mhz	46 dBm/ 10 Mhz 49 dBm/ 20 Mhz	46 dBm/ 10 Mhz 49 dBm/ 20 Mhz
No of UE antennas	up to 2	up to 2	up to 2	up to 2
UE antenna gain	0 dBi	0 dBi	0 dBi	0 dBi
UE Tx power	21 dBm	24 dBm	24 dBm	24 dBm
UE velocity	3 km/h	3 km/h	30 km/h	120 km/h
Inter-site interference modelling	Explicitly modelled	Explicitly modelled	Explicitly modelled	Explicitly modelled

20.3.6.2 Primary Module

The primary channel models are specified in general in the frequency range from 2 GHz to 6 GHz. However, the Rural Macro cell (RMa) model extends down to 450 MHz. The path-loss depends strongly on whether LOS conditions apply or not; Figure 20.7 shows the LOS probabilities in the different propagation scenarios.




Figure 20.7: LOS probabilities for the ITU deployment scenarios.

The detailed parameters of the path-loss and shadow fading for the primary module in LOS and non-LOS cases for each deployment scenario can be found in [22].

The long-term shadow fading (in dB) follows a log-normal distribution around the mean path-loss. The normalized autocorrelation function computed at distance Δx can be approximated by an exponential function as follows

$$R(\Delta x) = \exp(-\Delta x/d_{corr})$$

where the correlation length d_{corr} depends on the environment.

In addition to the test environments described above, the primary module also includes an optional suburban model.

Table 20.9 summarizes the multi-antenna radio channel parameters in the different test environments in LOS and NLOS conditions.

20.3.6.3 Extension Module

The optional extension module can be applied for specific macrocell cases where the radio channel parameters defined in the primary module do not accurately describe the actual case

Table 20.9: Radio channel parameters for the ITU IMT-Advanced test environments (LOS and NLOS) [22]

Scenarios		InH		UMi			UMa		RMa	
		LOS	NLOS	LOS	NLOS	Outdoor-to-Indoor	LOS	NLOS	LOS	NLOS
Delay spread (DS) log ₁₀ (s)	μ	-7.70	-7.41	-7.19	-6.89	-6.62	-7.03	-6.44	-7.49	-7.43
	σ	0.18	0.14	0.40	0.54	0.32	0.66	0.39	0.55	0.48
AoD spread (ASD) log ₁₀ (degrees)	μ	1.60	1.62	1.20	1.41	1.25	1.15	1.41	0.90	0.95
	σ	0.18	0.25	0.43	0.17	0.42	0.28	0.28	0.38	0.45
AoA spread (ASA) log ₁₀ (degrees)	μ	1.62	1.77	1.75	1.84	1.76	1.81	1.87	1.52	1.52
	σ	0.22	0.16	0.19	0.15	0.16	0.20	0.11	0.24	0.13
Shadow fading (SF) (dB)	σ	3	4	3	4	7	4	6	4	8
K-factor (K) (dB)	μ	7	N/A	9	N/A	N/A	9	N/A	7	N/A
	σ	4	N/A	5	N/A	N/A	3.5	N/A	4	N/A
Cross-correlations	ASD vs DS	0.6	0.4	0.5	0	0.4	0.4	0.4	0	-0.4
	ASA vs DS	0.8	0	0.8	0.4	0.4	0.8	0.6	0	0
	ASA vs SF	-0.5	-0.4	-0.4	-0.4	0	-0.5	0	0	0
	ASD vs SF	-0.4	0	-0.5	0	0.2	-0.5	-0.6	0	0.6
	DS vs SF	-0.8	-0.5	-0.4	-0.7	-0.5	-0.4	-0.4	-0.5	-0.5
	ASD vs ASA	0.4	0	0.4	0	0	0	0.4	0	0
	ASD vs K	0	N/A	-0.2	N/A	N/A	0	N/A	0	N/A
	ASA vs K	0	N/A	-0.3	N/A	N/A	-0.2	N/A	0	N/A
	DS vs K	-0.5	N/A	-0.7	N/A	N/A	-0.4	N/A	0	N/A
SF vs K	0.5	N/A	0.5	N/A	N/A	0	N/A	0	N/A	
Delay distribution		Exp	Exp	Exp	Exp	Exp	Exp	Exp	Exp	Exp
AoD and AoA distribution		Laplacian		Wrapped Gaussian			Wrapped Gaussian		Wrapped Gaussian	
Delay scaling parameter r_τ		3.6	3	3.2	3	2.2	2.5	2.3	3.8	1.7
cross-polarization ratio (XPR) (dB)	μ	11	10	9	8.0	9	8	7	12	7
Number of clusters		15	19	12	19	12	12	20	11	10
Number of rays per cluster		20	20	20	20	20	20	20	20	20
Cluster ASD		5	5	3	10	5	5	2	2	2
Cluster ASA		8	11	17	22	8	11	15	3	3
Per cluster shadowing standard deviation ζ (dB)		6	3	3	3	4	3	3	3	3
Correlation distance (m)	DS	8	5	7	10	10	30	40	50	36
	ASD	7	3	8	10	11	18	50	25	30
	ASA	5	3	8	9	17	15	50	35	40
	SF	10	6	10	13	7	37	50	37	120
	K	4	N/A	15	N/A	N/A	12	N/A	40	N/A

¹ AoA: angle-of-arrival, AoD: angle-of-departure, std: standard deviation, DS: rms delay spread, ASA: rms azimuth spread of arrival angles, ASD: rms azimuth spread of departure angles, SF: shadow fading, and K: Ricean K-factor.

(e.g. with respect to the BS antenna height, street width, city structure, etc). The extension module is intended to cover those cases beyond the evaluations of the IMT Advanced candidate radio interface technologies [22].

20.3.7 Comparison of MIMO Channel Models

The SCM, SCME, and WINNER channel models are compared in [19]. Table 20.10 compares the main features of the SCM, 3GPP LTE-evaluation, SCME, WINNER, and IMT-Advanced channel models, and Table 20.11 shows the numerical values of some of the key parameters.

Table 20.10: Feature comparison of SCM, LTE-evaluation, SCME, WINNER, and IMT-A models.

	SCM	LTE-eval	SCME	WINNER	IMT-A
Bandwidth > 20 MHz	No	Yes	Yes	Yes	Yes
Indoor scenarios	No	No	No	Yes	Yes
Outdoor-to-indoor scenarios	No	No	No	Yes	Yes
AoA / AoD elevation	No	No	No	Yes	No
Intra-cluster delay spread	No	Yes	Yes	Yes	Yes
Cross-correlation between large-scale parameters	No	No	No	Yes	Yes
Time evolution of model parameters	No	No	Yes*	Yes	No

**Continuous time evolution.

Table 20.11: Numerical comparison of SCM, LTE, SCME, WINNER, and IMT-A models.

Parameter	Unit	SCM	LTE-eval	SCME	WINNER	IMT-A
Max. bandwidth	MHz	5	20*	100*	100**	100**
Frequency range	GHz	2	N/A	2 – 5	2 – 6	2 – 6***
No. of scenarios		3	4	3	12	5
No. of clusters		6	6	6	8 – 20	10 – 20
No. of midpaths per cluster		1	3	3 – 4	1 – 3	1 – 3
No. of subpaths per cluster		20	N/A	20	20	20
No. of taps		6	18	18 – 24	12 – 24	14 – 24
Base station angle spread	deg	5 – 19	5 – 18	5 – 18	3 – 58	6 – 42
UE angle spread	deg	68	62 – 68	62 – 68	16 – 55	30 – 74
Delay spread	ns	160 – 650	231 – 841	231 – 841	16 – 630	20 – 365
Shadow fading standard deviation	dB	4 – 10	N/A	4 – 10	3 – 8	3 – 8

* Artificial extension from 5 MHz bandwidth.

** Based on 100 MHz measurements.

*** Rural model applies from 450 MHz to 6 GHz.

The LTE evaluation channel models are simplified models which have fixed delay and angle spreads. In that sense they apply well to link-level calibration models and

also to conformance testing. However, it could be argued that they are too simplified for the most thorough of system-level performance evaluations. The SCM, SCME and WINNER models provide random delay and angle spreads for different users and therefore apply well to system level (multicell, multi-user) simulations. Thus the delay spread can take quite different values during different drops for a single user or during parallel drops to different users – this increases the variability and dynamics of the radio channels in multi-user simulations. From the MIMO modelling point of view, the LTE and SCM/SCME/WINNER/IMT-A models differ significantly in the way the spatial correlation is defined. In SCM/SCME/WINNER/IMT-A the spatial correlation is defined by the nominal direction and the angular geometry of the subpaths of each delay tap, while the LTE evaluation model generates deterministic correlation values for different clusters. Another significant difference is that the SCM/SCME/WINNER/IMT-A models specify the path-loss and shadow fading for each of the user environments, which are important in system-level performance evaluation.

20.4 Radio Channel Implementation for Conformance Testing

Radio channel phenomena can be mimicked by software (simulation) or by hardware (emulation). Software simulation enables easy programmability and low cost, and is typically utilized for link- and system-level performance evaluations. On the other hand, hardware emulation may be used to test real hardware implementations, to verify and validate products, and to speed up testing. It typically offers more realistic testing.

20.4.1 Performance and Conformance Testing

Performance testing is a general term for all testing where some performance metrics such as power, voltage, sensitivity, Bit Error Rate (BER), Block Error Rate (BLER), throughput or handover success rate are measured. Test instruments such as signal generators and radio channel emulators need to be calibrated in order to ensure the accuracy of the measurements. The standardized channel models are useful to compare different products and technologies.

The purpose of conformance testing is not to secure optimal product operation in the field, but just to validate that the various products in the network conform to specified requirements and do not cause unexpected problems. A conformance test typically results in either a ‘pass’ or a ‘fail’ and is often performed by an external organization, such as a certified conformance test laboratory. Traditionally, the models used to conduct conformance testing have been the simple tapped-delay line models in order to allow their direct implementation in hardware-based radio channel emulators. However, more recent hardware-based channel emulators implement much more realistic channel models. For the LTE conformance test, the extended ITU channel models EPA, EVA and ETU are used.

20.4.2 Future Testing Challenges

The modelling of the radio channel is one of the key challenges for accurate performance measurement and testing. A direct cable connection (conductive testing) to transfer the

RF energy between the eNodeB and the UE avoids the uncertainty of the radio channel and allows measurements within well-defined and repeatable parameters. However, as UE antennas become more highly integrated, such testing becomes more difficult since the antenna connectors are no longer available. Conductive testing using RF connectors also has the drawback that it separates the antenna testing from the overall RF characteristics of the UEs which can lead to unrealistic results; moreover, the test connectors and cables themselves may affect the RF performance. Therefore, interest in Over-The-Air (OTA) testing, which eliminates the need for cable connections, is increasing.

Studies have taken place in 3GPP to examine measurements of radiated performance for MIMO and multi-antenna reception for HSPA and LTE terminals [25]. The particular focus is on methodologies based on anechoic RF chambers and on reverberation chambers. Anechoic chambers are free from reverberations, and multipath propagation is created in a controlled manner with multiple antennas and a radio channel emulator. By contrast, reverberation chambers create random echoes to mimic multipath propagation; this can favour specific MIMO terminals due to the unrealistically rich scattering environment with short delay spread.

20.5 Concluding Remarks

It is important to specify realistic propagation conditions for the evaluation of LTE performance. Accurate modelling of the spatial characteristics of the MIMO radio channel is needed in order to determine the best-performing multi-antenna transmission schemes, and to evaluate the LTE system performance in different radio environments.

3GPP specifications define advanced geometry-based radio channel models (SCM, SCME) which fulfil most of the sophisticated demands for system performance testing of LTE. SCME, WINNER and IMT-Advanced channel models provide additional features for further evaluations and simulations of LTE. The area of MIMO propagation modelling will continue to be important as the LTE conformance tests are developed, and as new versions of LTE with even larger numbers of antennas are standardized in the future.

References⁵

- [1] Commission of the European Communities, 'IST-WINNER Project', www.ist-winner.org.
- [2] G. J. Foschini, 'Layered Space-time Architecture for Wireless Communication in a Fading Environment when Using Multi-element Antennas'. *Bell Labs Tech. Journal*, Vol. 1, pp. 41–59, 1996.
- [3] E. Telatar, 'Capacity of multi-antenna Gaussian channels'. *European Trans. on Telecommunications*, Vol. 10, pp. 585–595, November 1999.
- [4] Commission of the European Communities, 'Digital Land Mobile Radio Communications-COST 207 Final Report', 1989.
- [5] Commission of the European Communities, 'Digital Mobile Radio towards Future Generation Systems-COST 231 Final Report', 1993.
- [6] Commission of the European Communities, 'Wireless Flexible Personalised Communications-COST 259 Final Report', 2001.

⁵All web sites confirmed 1st March 2011.

- [7] W. R. Braun and U Dersch, 'A Physical Mobile Radio Channel Model'. *IEEE Trans. on Vehicular Technology*, Vol. 40, pp. 472–482, 1991.
- [8] Commission of the European Communities, 'RACE-II advanced TDMA mobile access project: An Approach for UMTS'. Springer, Berlin/Heidelberg, 1994.
- [9] ITU-R M.1225 International Telecommunication Union, 'Guidelines for evaluation of radio transmission technologies for IMT-2000', 1997.
- [10] 3GPP Technical Specification 25.101, 'UE Radio Transmission and Reception (FDD)', www.3gpp.org.
- [11] W. C. Jakes, *Microwave Mobile Communications*. New York: John Wiley & Sons Ltd/Inc, 1974.
- [12] R. H. Clarke, 'A Statistical Theory of Mobile-radio Reception'. *Bell Syst. Tech. Journal.*, pp. 957–1000, July 1968.
- [13] Ericsson, Nokia, Motorola, and Rohde & Schwarz, 'R4-070572: Proposal for LTE Channel Models', www.3gpp.org, 3GPP TSG RAN WG4, meeting 43, Kobe, Japan, May 2007.
- [14] 3GPP Technical Specification 36.141, 'Base Station (BS) conformance testing', www.3gpp.org.
- [15] 3GPP Technical Specification 25.996, 'Spatial Channel Model for Multiple Input Multiple Output (MIMO) Simulations', www.3gpp.org.
- [16] D. S. Baum, J. Salo, G. Del Galdo, M. Milojevic, P. Kyösti and J. Hansen, 'An Interim Channel Model for Beyond-3G Systems' in *Proc. IEEE Vehicular Technology Conference* (Stockholm, Sweden), May 2005.
- [17] Commission of the European Communities, IST-WINNER II project deliverable D1.1.2 'WINNER II Channel Models', www.ist-winner.org.
- [18] Ericsson, Elektrobit, Nokia, Motorola, and Siemens, 'R4-060334: LTE Channel Models and simulations', www.3gpp.org, 3GPP TSG RAN WG4, meeting 38, Denver, USA, February 2006.
- [19] M. Narandzic, C. Schneider, R. Thomä, T. Jämsä, P. Kyösti and X. Zhao, 'Comparison of SCM, SCME, and WINNER Channel Models', in *Proc. IEEE Vehicular Technology Conference* (Dublin, Ireland), April 2007.
- [20] Agilent Technologies, 'R4-071318: LTE MIMO Correlation Matrices', August 2007.
- [21] 3GPP Technical Specification 36.101, 'User Equipment (UE) radio transmission and reception', www.3gpp.org.
- [22] ITU-R Report M.2135, 'Guidelines for evaluation of radio interface technologies for IMT-Advanced', www.itu.int/itu-r.
- [23] 3GPP Technical Report 36.814, 'Further Advancements for E-UTRA Physical Layer Aspects', www.3gpp.org.
- [24] ITU-R Recommendation P.1816 'The prediction of the time and the spatial profile for broadband land mobile services using UHF and SHF bands', www.itu.int/
- [25] Vodafone, 'RP-0900352: Measurement of radiated performance for MIMO and multi-antenna reception for HSPA and LTE terminals', www.3gpp.org, 3GPP TSG RAN, meeting #43, Biarritz, France, March 2009.

21

Radio Frequency Aspects

Moray Rumney, Takaharu Nakamura, Stefania Sesia, Tony Sayers and Adrian Payne

21.1 Introduction

Radio Frequency (RF) signal processing constitutes the final interface between the baseband techniques described in earlier chapters and the transmission medium – the air. The RF processing presents its own unique challenges for practical transceiver design, arising in particular from the fact that the air interface is a shared resource between multiple RF carriers, each with their own assigned portion of spectrum. The RF transmitter must be designed in such a way as not only to generate a clean signal within the assigned spectrum portion, but also to keep Inter-Carrier Interference (ICI) within acceptable levels. The receiver likewise must reliably demodulate the wanted signal, in order to avoid requiring excessive energy to be transmitted, whilst also rejecting interference from neighbouring carriers. Performance requirements for these RF aspects aim to ensure that equipment authorized to operate on an LTE carrier meets certain minimum standards.

In general, the performance requirements for LTE transceivers are intended not to be significantly more complex than for UMTS¹ in terms of implementation and testing. Consequently, many of the RF requirements for LTE are derived from those already defined for UMTS.

There are, however, a number of key differences between LTE and UMTS which affect the RF complexity and performance.

One such difference is the use of a variable channel bandwidth in LTE, up to a maximum of 20 MHz. Even the lowest category of LTE User Equipment (UE) is required to support all the bandwidths specified for the bands in which the UE is designed to operate, which in general includes 1.4, 3.0, 5.0, 10.0, 15.0 and 20.0 MHz. This

¹Universal Mobile Telecommunications System.

LTE – The UMTS Long Term Evolution: From Theory to Practice, Second Edition.

Stefania Sesia, Issam Toufik and Matthew Baker.

© 2011 John Wiley & Sons, Ltd. Published 2011 by John Wiley & Sons, Ltd.

requires a set of RF requirements to be defined for each bandwidth, in contrast to the Frequency Division Duplex (FDD) mode of UMTS which only supports a single channel bandwidth of 5 MHz. For LTE, the variable bandwidth represents a new challenge for the RF design and testing. For example, a transceiver for a constant-bandwidth radio system can potentially employ fixed filters at a number of points in the signal processing. Such filters are designed to pass a signal with known characteristics and reject particular frequencies. If, however the bandwidth of the transceiver is variable over a wide range, it is clear that fixed filters cannot be used. Frequencies which must be passed in 20 MHz operation may need to be rejected in the narrower bandwidth modes. This implies that LTE transceivers must be more adaptable than those of previous systems, while also being cleaner in transmission and having better selectivity in reception.

A second difference between UMTS and LTE is that in LTE it is assumed that the UE has at least two receive antennas. This means that multiple RF signal paths are needed at all times.

Thirdly, LTE is even more adaptable than UMTS in terms of the range of data rates it supports in order to suit different Signal-to-Interference-plus-Noise Ratio (SINR) () conditions (e.g. from 4×2 Multiple-Input Multiple-Output (MIMO) 64QAM in high SINR conditions to Single-Input Single-Output (SISO) QPSK in low SINR conditions (see Section 10.2)). Together with the ability to vary the instantaneous bandwidth to a given user, this implies a large number of modes of operation and flexibility in signal-handling capabilities. The reception of the maximum data rate requires high SINR at the highest bandwidth, which is particularly challenging for the analogue-to-digital converter in the receiver.

Fourthly, the LTE signal structure itself alters which specific RF aspects are the most critical. As already discussed extensively in Chapters 5 and 14, the use of Orthogonal Frequency Division Multiplexing (OFDM) in the downlink and Single Carrier Frequency Division Multiple Access (SC-FDMA) in the uplink provides an inherent robustness against multipath propagation, which means that amplitude and phase distortions from receiver and transmitter filters are not as critical as for UMTS's single-carrier Wideband Code Division Multiple Access (WCDMA) modulation. On the other hand, OFDM requires better frequency synchronization and is more sensitive to phase noise as shown in Chapter 5 and further discussed in Section 21.5.

It is also worth noting that, for WCDMA, the RF requirement specifications are separate for FDD and TDD (Time Division Duplex) [1–4], while in LTE the commonality between FDD and TDD is such that they can be handled in the same specifications [5,6]. Nevertheless, a more detailed discussion of aspects specific to TDD operation is given in Chapter 23.

Some further useful background on the transmitter and receiver RF characteristics can be found in [7] and [8] for the LTE UE and eNodeB respectively.

This chapter seeks to describe the impact of the LTE physical layer radio requirements on the implementation complexities of an LTE transceiver, compared to the well-known UMTS system. The relevant spectrum bands for LTE are introduced in Section 21.2. Radio requirements for the transmitter and receiver are discussed in Sections 21.3 and 21.4 respectively, and some of the implementation challenges are highlighted. Both the UE and eNodeB are addressed, although with the greater emphasis on the UE side. This is followed by a discussion of the RF impairments for both transmitter and receiver in Section 21.5,

including mathematical models of the most common impairments to help elucidate the impact of RF imperfections on the overall radio behaviour.

21.2 Frequency Bands and Arrangements

LTE and UMTS are both defined for a wide range of different frequency bands, in each of which one or more independent carriers may be operated. Tables 21.1 and 21.2 give details of the frequency bands for FDD and TDD operation respectively. For FDD, the duplex separation is not actually defined, but typically the uplink and downlink carriers in a pair are in a similar position in their respective bands so that the duplex separation is usually approximately $F_{DL,low} - F_{UL,low}$ as shown in Table 21.1.

Table 21.1: UMTS and LTE frequency bands for FDD.

Band Number	Uplink (MHz)	Downlink (MHz)	Band Gap (MHz)	Duplex Separation (MHz)	UMTS Usage	LTE Usage
	$F_{UL,low} - F_{UL,high}$	$F_{DL,low} - F_{DL,high}$				
1	1920–1980	2110–2170	130	190	Y	Y
2	1850–1910	1930–1990	20	80	Y	Y
3	1710–1785	1805–1880	20	95	Y	Y
4	1710–1755	2110–2155	355	400	Y	Y
5	824–849	869–894	20	45	Y	Y
6*	830–840	875–885	35	45	Y	Y
7	2500–2570	2620–2690	50	120	Y	Y
8	880–915	925–960	10	45	Y	Y
9	1749.9–1784.9	1844.9–1879.9	60	95	Y	Y
10	1710–1770	2110–2170	340	400	Y	Y
11	1427.9–1447.9	1475.9–1495.9	28	48	Y	Y
12	698–716	728–746	12	30	Y	Y
13	777–787	746–756	21	31	Y	Y
14	788–798	758–768	20	30	Y	Y
17	704–716	734–746	18	30	N	Y
18**	815–830	860–875	30	45	N	Y
19**	830–845	875–890	30	45	Y	Y
20**	832–862	791–821	11	41	Y	Y
21**	1447.9–1462.9	1495.9–1510.9	33	48	Y	Y
23***	2000–2020	2180–2200	160	180	N	Y
24***	1626.5–1660.5	1525–1559	-135.5	-101.5	N	Y
25***	1850–1915	1930–1995	15	80	Y	Y
26****	814–849	859–894	10	45	Y	Y

* This band was defined in the context of Release 8; it is replaced by Band 19 for later releases (Release 9 and 10). Only legacy terminals would use band 6.

** These bands were specified in the timeframe of Release 9, although all bands are release-independent and can be implemented by UEs conforming to any release.

*** These bands were specified in the timeframe of Release 10, although all bands are release-independent and can be implemented by UEs conforming to any release.

**** This band was under consideration at the time of going to press.

Table 21.2: UMTS and LTE frequency bands for TDD.

Band	$F_{\text{low}}-F_{\text{high}}$ (MHz)	UMTS	LTE
33	1900–1920	Y	Y
34	2010–2025	Y	Y
35	1850–1910	Y	Y
36	1930–1990	Y	Y
37	1910–1930	Y	Y
38	2570–2620	Y	Y
39	1880–1920	N	Y
40	2300–2400	Y	Y
41*	2496–2690	N	Y
42*	3400–3600	N	Y
43*	3600–3800	N	Y

* These bands were specified in the timeframe of Release 10, although all bands are release-independent and can be implemented by UEs conforming to any release.

In Table 21.1 band 15 and 16 are not shown; these bands were introduced by the European Telecommunication Standard Institute (ETSI) and are applicable to Europe only.

A typical UE would support a certain subset of these bands depending on the desired market, since supporting all would be challenging for the transceiver, in particular for the front-end components such as Power Amplifiers (PAs), filters and duplexers. The set of frequency bands chosen defines the capability of the UE to switch bands, roam between national operators and roam internationally.

Each frequency band is regulated to allow operation in only a certain set of channel bandwidths, as shown in Table 21.3. Some frequency bands do not allow operation in the narrow bandwidth modes below 5 MHz, while others do not allow operation in the wider bandwidths, generally 15 MHz or above. ‘—’ denotes that the channel bandwidth is not supported in the specific band; ‘X’ indicates the channel bandwidth is too wide to be supported in the band.

In LTE-Advanced, new requirements are defined for combinations of carriers in the same and different bands for carrier aggregation. These combinations are discussed in Section 28.4.3.

Due to LTE’s flexibility in the frequency domain, two new concepts are defined: ‘transmission bandwidth configuration’, which defines the maximum number of Resource Blocks (RBs) for any channel bandwidth, and ‘transmission bandwidth’ for a specific transmission, which depends on the scheduled resource allocation and can be less than or equal to the transmission bandwidth configuration as shown in Figure 21.1.

Moreover it is important to distinguish between the transmission bandwidth configuration, generally given in RBs, and the channel bandwidth, which is expressed in MHz as shown in Table 21.4.

Table 21.3: Number of supported non-overlapping channels in each frequency band and bandwidth.

LTE band	Downlink bandwidth	Channel bandwidth (MHz)					
		1.4	3	5	10	15	20
1	60	—	—	12	6	4	3
2	60	42	20	12	6	4*	3*
3	75	53	25	15	7	5*	3*
4	45	32	15	9	4	3	2
5	25	17	8	5	2*	—	—
6	10	—	—	2	1*	X	X
7	70	—	—	14	7	4	3*
8	35	25	11	7	3*	—	—
9	35	—	—	7	3	2*	1*
10	60	—	—	12	6	4	3
11	25	—	—	4	2*	1*	1*
12	18	12	6	3*	1*	—	X
13	10	—	—	2*	1*	X	X
14	10	—	—	2*	1*	X	X
...							
17	12	—	—	2*	1*	X	X
18	15	—	—	3	1*	1*	X
19	15	—	—	3	1*	1*	X
20	30	—	—	6	3*	2*	1
21	15	—	—	3	1*	1*	X
23	20	14	6	4	2	—	—
24	34	—	—	6	3	—	—
25	65	46	21	13	6	4*	3*
26	35	25	11	7	3*	2*	—
...							
33	20	—	—	4	2	1	1
34	15	—	—	3	1	1	X
35	60	42	20	12	6	4	3
36	60	42	20	12	6	4	3
37	20	—	—	4	2	1	1
38	50	—	—	10	5	3	2
39	40	—	—	8	4	3	2
40	100	—	—	20	10	6	5
41	194	—	—	38	19	12	9
42	200	—	—	40	20	13	10
43	200	—	—	40	20	13	10

* A relaxation of the specified UE receiver sensitivity requirement is allowed.




Figure 21.1: Definition of transmission bandwidth configuration and transmission bandwidth. Reproduced by permission of © 3GPP.

Table 21.4: Transmission bandwidth configuration for each channel bandwidth.

Channel bandwidth	MHz	1.4	3	5	10	15	20
Transmission bandwidth configuration	N_{RB}	6	15	25	50	75	100
	MHz	1.06	2.7	4.5	9	13.5	18

21.3 Transmitter RF Requirements

This section explains the implications of the LTE RF performance specification for the design of its transmitters.

Each transmitter must satisfy two categories of requirements: those relating to the power level and quality of the intended transmissions, and those prescribing the level of unwanted emissions which can be tolerated. The latter is usually the more challenging for the radio designer.

21.3.1 Requirements for the Intended Transmissions

21.3.1.1 Signal Quality: Error Vector Magnitude (EVM)

The quality of the transmitted radio signal has to fulfil certain requirements. The main parameter used to measure this quality is the Error Vector Magnitude (EVM), which is a

measure of the distortion introduced by the RF imperfections of practical implementations.² It is defined as the magnitude of the difference between a theoretical reference signal (i.e. the signal defined by the physical layer specification equations) and the actual transmitted signal (normalized by the intended signal magnitude). The EVM sets the maximum possible SNR of a radio link in the absence of any noise, interference, propagation loss and other distortions introduced by the radio channel; it therefore serves to determine the maximum useful modulation order and code rate.

The EVM measures the quality of the transmitted signal across all the allocated RBs. The measurement duration is one slot for the UE (uplink) and one subframe for eNodeB (downlink) and takes into account all the symbols belonging to the modulation scheme under test. Figures 21.2 and 21.3 show the EVM measurement points in the downlink and uplink transmission chains respectively [5, 6].

The EVM measurement is taken after an equalizer in the test equipment, which carries out per-subcarrier channel correction [5, 6]. The equalizer is used in order to obtain a measurement which realistically shows what a receiver might experience. It is intended to reflect the fact that the equalizer in the receiver is capable of correcting some of the impairments of the transmitted signal to some extent. A zero-forcing equalizer is used for EVM measurement;³ however, real receiver implementations may use different equalization techniques, and therefore the measured EVM may not exactly correspond to the actual signal quality experienced by all receivers.

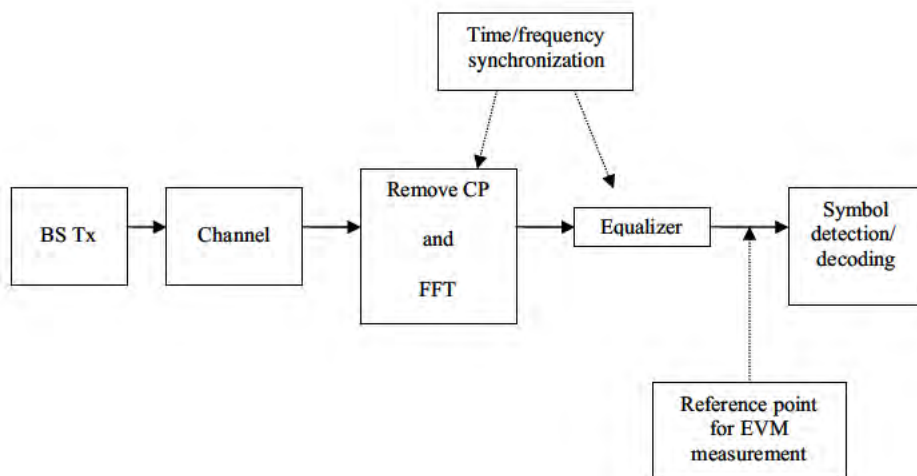


Figure 21.2: Measurement points for the EVM for the downlink signal.

²Note that distortion may also be introduced in the process of balancing in-channel and out-of-channel performance, especially in the eNodeB.

³Details of how to compute the coefficients of the equalizer can be found in [5, 6].

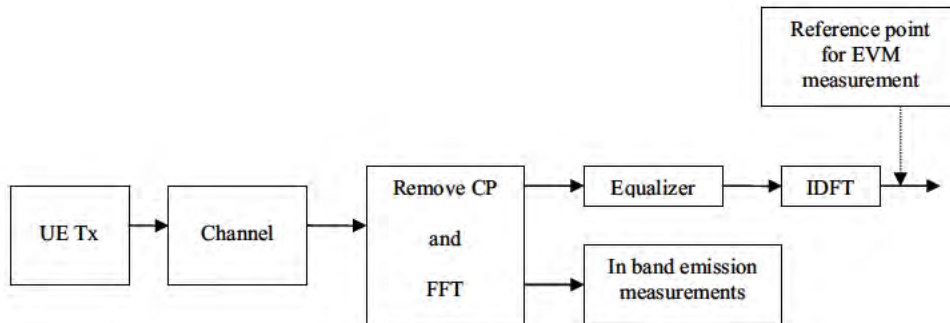


Figure 21.3: Measurement points for the EVM for the uplink signal.

It should be noted that before measuring the EVM, time and frequency synchronization must be carried out. The test equipment then computes the EVM for two extreme values of sample timing difference between the FFT processing window and the nominal timing of the ideal signal; these two extreme values correspond to the beginning and end of the window, the length of which is expressed as a percentage of the Cyclic Prefix (CP) length. Finally the measurement is averaged over 20 slots for UE (uplink) and 10 subframes for eNodeB (downlink) to reduce the impact of noise.

In LTE the EVM is required to be less than 17.5% for QPSK, 12.5% for 16QAM and 8% for 64QAM (64QAM being applicable to the downlink only) [5, 6]. These EVM values are designed to correspond to no more than a 5% loss in average and cell-edge throughputs in typical deployment scenarios. At the link level, EVM is equivalent to an SNR loss.

A description of the main transmitter impairments which generally give rise to non-zero EVM is given in Section 21.5.1.

21.3.1.2 Transmit Output Power

The transmitted output power directly influences the inter-cell interference experienced by neighbouring cells using the same channel, as well as the magnitude of unwanted emissions outside the transmission band. This affects the ability of the LTE system to maximize spectral efficiency, and it is therefore important that the transmitters can set their output power accurately.

For the eNodeB, the maximum output power must remain within ± 2 dB of the rated power, (P_{RAT}) declared by the manufacturer. There are three classes of eNodeB (see Section 24.4): the Home eNodeB with $P_{RAT} \leq 20$ dBm, the Local Area eNodeB with $P_{RAT} \leq 24$ dBm and the Wide Area eNodeB for which no upper limit has been defined. In addition, the dynamic range in the frequency domain (computed as the difference between the power of a given

Resource Element (RE) and the average RE power) must not exceed specified limits [5, 6] depending on the modulation order, in order to avoid saturating the UE receivers.⁴

For the UE, a requirement is defined for only one power class, known as ‘power class 3’ for which the maximum output power is 23 dBm. Like the eNodeB, the UE must satisfy this requirement within a range of ± 2 dB.

The UE maximum output power requirements may be modified by a number of factors [5]:

- **Maximum Power Reduction (MPR).** The purpose of MPR is to allow the UE, in some demanding configurations, to lower its maximum output power in order to meet the general requirements on signal quality and Out-Of-Band (OOB) emissions. Note that the MPR is an allowance and the UE does not have to use it. Table 21.5 shows the allowed values of MPR as a function of the modulation scheme, the channel bandwidth and the transmission bandwidth (number of transmitted RBs).

Table 21.5: Maximum power reduction for power class 3.

Modulation	Channel bandwidth (MHz)/ transmission bandwidth (RB)						MPR (dB)
	1.4	3	5	10	15	20	
QPSK	> 5	> 4	> 8	> 12	> 16	> 18	≤ 1
16QAM	≤ 5	≤ 4	≤ 8	≤ 12	≤ 16	≤ 18	≤ 1
16QAM	> 5	> 4	> 8	> 12	> 16	> 18	≤ 2

- **Additional MPR (A-MPR).** The eNodeB may inform the UE of the possibility of further lowering its maximum power by signalling an A-MPR. The need for A-MPR occurs with certain combinations of E-UTRA bands, channel bandwidths and transmission bandwidths for which the UE must meet additional (more stringent) requirements for spectrum emission mask and spurious emissions (see Section 21.3.2). As with MPR, the A-MPR is an allowance, not a requirement, and it applies in addition to MPR. Regardless of whether the UE makes use of the allowed MPR and A-MPR, the additional requirements for spectrum emission mask and spurious emissions that are signalled by the network always apply. The reason for the complex set of conditions for the relaxation is the expected intermodulation products which may fall into adjacent bands which have different levels of sensitivity (e.g. public safety bands).
- **ΔT_C .** ΔT_C is a 1.5 dB reduction in the lower limit of the maximum output power range when the signal is within 4 MHz of the channel edge.
- **Power Management MPR (P-MPR).** Introduced in Release 10, the P-MPR allows a UE to reduce its maximum output power when other constraints are present. In particular, multi-RAT⁵ terminals may have to limit the LTE transmission power if

⁴For example, a Physical Downlink Control CHannel (PDCCH) RE power can range between -6 dB and $+4$ dB around its average.

⁵Radio Access Technology.

transmissions on another RAT are taking place simultaneously. Such power restrictions may arise, for example, from regulations on Specific Absorption Rate (SAR) of radio energy into a user's body or from out-of-band emission requirements that may be affected by the inter-modulation products of the simultaneous radio transmissions. The P-MPR is not aggregated with MPR or A-MPR, since any reduction in a UE's maximum output power for the latter factors helps to satisfy the requirements that would have necessitated P-MPR.

Taking these factors into account, the UE has to configure its nominal maximum power P_{CMAX} (i.e. the highest power at which the UE will attempt to transmit) within the following upper and lower limits:

$$\begin{aligned} P_{\text{CMAX, L}} &= \min(P_{\text{EMAX}} - \Delta T_C, P_{\text{PowerClass}} - \max(\text{MPR} + \text{A-MPR}, \text{P-MPR}) - \Delta T_C) \\ P_{\text{CMAX, H}} &= \min(P_{\text{EMAX}}, P_{\text{PowerClass}}) \end{aligned} \quad (21.1)$$

where $P_{\text{PowerClass}}$ is the original power class of the UE and P_{EMAX} is a maximum power that may be signalled by the network.

Note that when carrier aggregation is configured (for UEs of Release 10 and beyond), P_{CMAX} becomes a component-carrier-specific nominal maximum power, $P_{\text{CMAX,c}}$, as explained in Section 28.3.5.

Finally, the actual transmitted power is allowed to vary over a wider range due to uncertainties in the transmit chain. The added tolerance is a function of P_{CMAX} which varies from 2 dB at high powers to 7 dB for powers below 8 dBm.

In summary, the concept of maximum power for the UE is a complex and dynamic function of many variables designed so that the needs of specific operating conditions can be met without over-specifying the transmitter design. Details of the maximum power specification can be found in [5].

21.3.1.3 Output Power Dynamics

LTE, like UMTS and HSPA, needs to ensure user orthogonality in the time domain. In WCDMA, the requirements for this were relatively straightforward being based on so-called 'on-off masks' which define the allowable transmit power during the 'off' and 'on' periods of transmission. For LTE, similar general requirements exist for the eNodeB in both FDD and TDD [6]. However, for the LTE UE the requirements are considerably more complex than those for WCDMA due to the characteristics of the SC-FDMA scheme used for the uplink transmissions. Figure 21.4 shows the general 'on/off time mask' which applies whenever the UE is required to switch on or off.⁶ Note that although the requirement is given in terms of a 'mask', it actually applies to the average power during the 'on' and 'off' periods [5].

In UMTS the transient period (20 μs) is centred on the slot boundary, while in LTE the transient period is in general shifted forwards into the next subframe. Hence, the first few SC-FDMA samples are vulnerable to being corrupted by insufficient transmit power or inter-UE interference whereas the samples at the end of the subframe are protected. For the Physical Random Access CHannel (PRACH) the transient periods are located outside the preamble in order to protect the whole PRACH preamble. The same applies to Sounding Reference

⁶This is also applicable for Discontinuous Transmission (DTX) measurement gaps – see Chapter 22.




Figure 21.4: General on-off time mask for the UE. Reproduced by permission of © 3GPP.

Signals (SRSs) (see Section 15.6) which need to be protected to ensure the eNodeB can exploit them entirely⁷ to carry out reliable uplink channel sounding without impairments arising from transient periods. For TDD, where two SRSs can be transmitted on adjacent symbols in the UpPTS field, the transient periods are located between the two SRSs.

The general time masks apply in the case of transitions into and out of the off power state. A general output power dynamic requirement also exists for continuous transmission at slot boundaries where frequency hopping occurs and at subframe boundaries for either frequency hopping or power changes. In all these cases, the transient periods are symmetric about the slot or subframe boundaries. Transient power masks also apply in the following cases:

- Transition from PUCCH/PUSCH to SRS with DTX after SRS;
- Transition from PUCCH/PUSCH to SRS to PUCCH/PUSCH;
- Transition from DTX to SRS to PUSCH/PUCCH;
- Transition from PUCCH/PUSCH to DTX in an SRS symbol to PUCCH/PUSCH.

21.3.2 Requirements for Unwanted Emissions

Ideally, the radio should transmit nothing at all outside its designated transmission channel. However, in practice this is very far from being the case.

Outside the intended channel bandwidth, the LTE specifications define two separate kinds of unwanted emission: *Out-Of-Band (OOB) emissions* and *spurious emissions*. These are shown schematically in Figure 21.5.

Out-of-band emissions are those which fall in a band close to the intended transmission, while spurious emissions can be at any other frequency. The precise boundary between the OOB range and the spurious range is different for different aspects of the LTE specifications.

LTE defines requirements for both types of unwanted emission, with those for spurious emissions being the more stringent. As mentioned in Section 21.3.1.2, the specified requirements must be fulfilled when indicated by network signalling or under particular conditions. This means that, unlike the relatively fixed RF requirements of earlier systems,

⁷Note that the SRS is only 66.67 μs in duration.




Figure 21.5: Transmitter spectrum. Reproduced by permission of © 3GPP.

the emission requirements may change in different scenarios. If the transmitter cannot satisfy a particular requirement, it can either switch the transmitter off, or adapt the transmitter characteristics (for example by reducing the transmitter power or using software-defined radio techniques) to alter the performance in the required manner. This is explained in more detail in Section 21.3.3.

21.3.2.1 Out-of-Band Emissions

Since OOB emissions occur close to the wanted transmission, increasing the power level of the wanted transmission will usually increase the level of the unwanted emissions. Conversely, reducing the transmitted power is usually an effective method for reducing the OOB emissions, thus providing one possible method for responding to network-signalled requirements as discussed above.

OOB emissions may be an almost inevitable by-product of the modulation process itself, and are also often caused by non-linearities in PAs. A summary of the main practical transmission impairments which can cause OOB emissions is given in Section 21.5.1.

In previous fixed-bandwidth radio systems (including UMTS), the OOB emissions requirements have been defined with respect to the centre frequency of the transmission. As the bandwidth of the LTE system is variable, it is more convenient to define OOB requirements with respect to the edge of the channel bandwidth rather than the centre of the channel.

In LTE, OOB emissions are defined by means of Spectrum Emission Masks (SEMs) and Adjacent Channel Leakage Ratio (ACLR) requirements. The SEM has a much narrower reference bandwidth than the ACLR, which is a stricter requirement.

Spectrum Emission Mask

The Spectrum Emission Mask (SEM) is a mask defined for out-of-channel emissions relative to the in-channel power. The spectrum emission mask of the UE applies to frequencies

within Δf_{OOB} (MHz) of the edge of the assigned LTE channel bandwidth, as shown diagrammatically in Figure 21.5. Figure 21.6 shows the basic SEM requirement for a UE transmitter.




Figure 21.6: Spectrum Emission Mask (SEM) for a UE transmitter, for each channel bandwidth.

LTE also provides a set of ‘Additional SEMs’ (A-SEMs) which the network may instruct the UE to use in particular deployment scenarios.⁸ Such an instruction may, for example, be signalled on handover to a new cell [5].

For the eNodeB, the SEM is defined in a similar manner; an example is given in Figure 21.7. It should again be noted that this represents only the basic requirement, and there are many conditions and additional cases to be taken into account in different circumstances [6].

Adjacent Channel Leakage Ratio

The second method used to measure OOB emissions in LTE is the ACLR. While the SEM measures the performance of the transmitter, the ACLR measures the power which actually leaks into certain specific nearby radio channels and thus estimates how much a neighbouring radio receiver will be affected by the OOB emissions from the transmitter.

The ACLR is defined as the ratio of the filtered mean power centred on the assigned channel frequency to the filtered mean power centred on an adjacent channel frequency.

The LTE specifications not only set ACLR requirements for adjacent 20 MHz LTE channels, but also for UMTS channels which may be located in two adjacent 5 MHz channels (i.e. a total of 10 MHz).

⁸These are Network Signals 03, 04, 06 and 07.




Figure 21.7: Spectrum emission mask for eNodeB transmitter.

The ACLR of a UE for an adjacent LTE channel is required to be >30 dB for the whole 20 MHz Bandwidth, >33 dB for the adjacent UMTS channel deployed in the closest 5 MHz of the OOB portion of the spectrum and >36 dB for the adjacent UMTS channel deployed in the next 5 MHz portion of the OOB spectrum [5].

21.3.2.2 Spurious Emissions

Spurious emissions occur well outside the bandwidth necessary for transmission⁹ and may arise from a large variety of non-ideal effects including harmonic emissions and intermodulation products.¹⁰ The magnitude of the spurious emissions may or may not vary with transmitter power.

The basic spurious emissions requirements for an LTE UE are given in Table 21.6.

Table 21.6: Spurious emissions limits.

Frequency range	Maximum level (dBm)	Measurement bandwidth
9 kHz ≤ 150 kHz	-36	1 kHz
150 kHz ≤ 30 MHz	-36	10 kHz
30 MHz ≤ 1000 MHz	-36	100 kHz
1 GHz ≤ 12.75 GHz	-30	1 MHz

⁹The spurious emission limits apply for frequency ranges that are more than Δf_{OOB} (MHz) from the edge of the channel bandwidth – see [5, Section 6.6.3.1].

¹⁰Intermodulation products are unwanted frequencies generated whenever two or more tones are present in a non-linear device, such as an amplifier or mixer. The possible combinations of generated frequencies can be described by a power series.

These basic requirements are more stringent than the OOB requirements, but are not usually difficult to meet. The main challenge for the transmitter designer lies in some 'additional spurious emissions' requirements which are specified for certain frequency bands. The additional spurious emissions requirements exist to protect the receiver on the UE, colocated receivers, or nearby receivers from the transmission. In the case of FDD radios, the receiver and transmitter operate at the same time, and the receiver relies on the duplex spacing to separate the weak received signal from its strong transmitted signal (see Section 21.4.2 for further details on transmitted signal leakage in the receiver band). If the transmitter emits anything significant on its own receiver frequency, then reception will be blocked or degraded. This results in the additional spurious emissions requirements being very stringent. In addition, some further additional spurious emissions requirements may be signalled by the network if needed for specific deployment scenarios.

21.3.3 Power Amplifier Considerations

The ACLR is typically directly related to the operating point of the power amplifier. In general, leakage into adjacent channels increases sharply as the PA is driven into its non-linear operating region at the highest output power levels, due in particular to the intermodulation products. Consequently it is important that the peak output power of the UE does not cause the PA to enter too far into this region. On the other hand, most PAs are designed to operate efficiently only in a small region at the top of the linear operating region. This corresponds to the *rated* output power of the PA, below which the efficiency falls sharply. Since high efficiency is crucial to ensuring a long battery life for the UE, it is therefore also desirable to keep the PA operating as close as possible to the top of the linear operating region.

If the ACLR and spectrum mask requirements cannot be met, typically the UE output power must be reduced to bring the leakage to acceptable levels. This can be achieved without undue loss of efficiency by reducing the peak output power of the PA – a process known as *de-rating*. As discussed in Section 21.3.1.2, the LTE specifications provide an MPR for this purpose. The amount of de-rating required is highly dependent on the Peak to Average Power Ratio¹¹ (PAPR) and bandwidth of the transmitted signal. For example, for any given channel bandwidth and PA, transmissions with a larger occupied bandwidth create more OOB emissions, resulting in larger adjacent channel leakage than transmissions with lower occupied bandwidth; this is shown in Figure 21.8, in which WCDMA and LTE occupied bandwidths and adjacent channel leakage are compared in a 5 MHz channel.

The increase in OOB emissions from the larger occupied bandwidth of the LTE signal is mainly due to increased adjacent channel occupancy by 3rd and 5th order InterModulation products (called 3rd and 5th order IM plateau in Figure 21.8).

For an LTE uplink transmitted signal, certain combinations of RB allocations and modulation schemes create more OOB emissions than others. For QPSK and 16QAM the number of RBs which can be allocated at the edge of the channel for a given amount of power de-rating is shown in Figure 21.9. For applications such as VoIP, which typically use low-bandwidth transmissions with QPSK modulation, no power de-rating is required from the normal rated maximum UE output power (see Section 21.3.1.2). This helps to ensure

¹¹The concept of peak to average power ratio is discussed later in this section and in Section 5.2.2.




Figure 21.8: Adjacent channel power increase (2–3 dB) for 4.51 versus 3.84 MHz Occupied BandWidth (OBW) [9]. Reproduced by permission of © 2006 Motorola.




Figure 21.9: Maximum power level versus the number of RBs which can be allocated. Reproduced by permission of © 2006 Motorola.

wide-area coverage for such applications, since the full rated output power of the PA can be used to counteract path-loss at the cell edge.

Moreover, the structure of the LTE uplink signal, with control information being usually situated at the channel edge and high-bandwidth data transmissions toward the middle of the band (see Section 16.3), helps to improve OOB emissions and reduce ACLR.

Taking the above considerations into account, the Total Power De-rating (TPD) required to meet a given ACLR requirement can be broken down into an element corresponding to the occupied bandwidth (as a proportion of the channel bandwidth), termed here the Occupied Bandwidth Power De-rating (OBPD), and an element corresponding to the waveform of the transmitted signal, termed here the Waveform Power De-rating (WPD).

OBPD can be approximately expressed as a ratio of the transmitted Occupied Bandwidth (OBW) and the Reference Occupied Bandwidth (OBW_{ref}) [10]:

$$OBPD \propto 10 \log_{10} \left(\frac{OBW}{OBW_{ref}} \right) \quad (21.2)$$

The other element of the required TPD, the WPD, arises in a large part from the dynamic range of the signal. This is usually quantified in terms of its *Peak to Average Power Ratio* (PAPR) and *Cubic Metric* (CM). These measures are typically used as indications of how much PA power headroom is required (i.e. how far from the rated power the PA generally has to operate) to avoid entering the non-linear region of operation. A high PAPR means that on average the operating point of the PA has to be lower in order to avoid the non-linear region and achieve the required ACLR. This results in a reduction in efficiency. A rule of thumb is that for each 1 dB increase in required power amplifier headroom, a corresponding 10 to 15% increase in PA current drain occurs, leading to a corresponding reduction in battery life. De-rating the PA is therefore usually a preferable strategy (as opposed to increasing the power headroom), since it achieves the required reduction in adjacent channel leakage by reducing the total maximum output power (including that of the wanted signal), while allowing the PA still to operate partly in the non-linear region. De-rating therefore enables a UE to meet the OOB emission requirements without a loss of efficiency, at the expense of a loss in coverage due to the reduced power in the wanted channel.

The CM of a waveform can be of interest because it characterizes the effects of the 3rd order (cubic) non-linearity of a PA on the waveform of interest relative to a reference waveform, in terms of the power de-rating needed to achieve the same ACLR as that achieved by the reference waveform at the PA's rated maximum power level. As discussed in Chapter 14, one criterion which was instrumental in selecting SC-FDMA instead of OFDMA for the LTE uplink was the low inherent CM of the SC-FDMA waveform compared to OFDMA, as shown in Figure 21.10, thus enabling more efficient operation for a given ACLR and maximum total output power.

However, it should be noted that the CM does not predict the spectral location of the non-linear distortion, and it is possible that two waveforms with different RB allocations have the same CM but one creates worse out-of-band emissions than the other; therefore, CM has not been defined in the LTE specifications.




Figure 21.10: Cubic Metric (CM) for different waveforms.
Reproduced by permission of © 2006 Motorola.

21.4 Receiver RF Requirements

Like the transmitter, the LTE receiver requirements are based largely on those of UMTS, with many of the RF requirements being identical or similar. The general intention is for the implementation of an LTE receiver not to be significantly more complex than UMTS, so as to reduce redesign effort. An underlying assumption is that there is no need for tightening the coexistence specifications, as the basic cellular scenarios remain similar.

The main differences between the LTE and UMTS RF receiver requirements arise from the variable channel bandwidth and the new multiple access schemes.

The emphasis in this section is on the UE receiver requirements for FDD operation.

21.4.1 Receiver General Requirements

The receiver requirements are based on a number of key assumptions for testing purposes:

- The receiver has integrated antennas with a 0 dBi gain.
- The receiver has two antenna ports. Requirements for four ports may be added in the future.
- Test signals of equal power level are applied to each port, with Maximum Ratio Combining (MRC) being used to combine the signals. It is assumed that the signals

come from independent Additive White Gaussian Noise (AWGN) channels, so that signal addition gives a 3 dB diversity gain.

Although LTE supports a variety of Modulation and Coding Schemes (MCSs) (see Chapter 10), the RF receiver specifications are defined for just two MCSs (referred to as 'reference channels'), near the extremes of the available range, in order to reduce the number of type approval tests which have to be performed. The 'low SNR' reference channel uses QPSK with a code rate of 1/3, while the 'high SNR' reference channel uses 64QAM with a code rate of 3/4. For each of these reference channels, the SINR requirements at which 95% throughput should be achieved are specified.

21.4.2 Transmit Signal Leakage

When an LTE receiver is tested in full-duplex FDD mode, the transmitter must also be operating so that signal leakage from the transmitter is taken into account. This does not apply to half-duplex FDD or to TDD operation, which are discussed in detail in Chapter 23.

An FDD receiver's exposure to the transmitted signal from the same equipment is largely due to insufficient isolation within the duplexer. The power of the leakage of the transmitted signal is proportional to the transmitted power.

Not only can the power of the fundamental components of the transmitted signal interfere with the receiver, but also the OOB phase noise of the transmitted signal can fall into the receive band. In LTE the spurious emissions from the UE transmitter in its own receive band are required to be -47 dBm or lower, measured in a 1 MHz bandwidth. This corresponds to -107 dBm/Hz. The maximum transmit power for a mobile device is $+23$ dBm (power class 3), so the spurious emissions requirement is -130 dBc/Hz. UMTS has a tougher requirement of -60 dBm spurious emissions in 3.84 MHz bandwidth, which is -125.8 dBm/Hz or -149.8 dBc/Hz for Class 3 $+24$ dBm output power. However, for UMTS the toughest requirement is in one specific receive band ('Band 8') which requires -153 dBc/Hz (when considering the maximum output power). The transmitter may therefore be designed to achieve this more broadly, including its own receive band.

For both UMTS and LTE, receiver sensitivity is measured with the transmitter operating at the full power for its class. In order to allow the transmitter to degrade the receiver sensitivity by no more than 0.5 dB, the transmitter noise needs to be 9 dB below the Noise Floor (NF) at the antenna ($NF - 9$ dB = -183 dBm/Hz). If we assume that the spurious emissions of the transmitter in its receive band are just at the limit of the specifications, it can be shown that the duplex isolation needs to be at least 78 dB for LTE and 59 dB for UMTS.¹²

However, making duplexers smaller and cheaper is often achieved at the expense of isolation, and typical duplexers might provide just 45 to 55 dB isolation. Potentially, therefore, the transmit signal could reduce receiver sensitivity substantially. This is not acceptable, so within the receive band the LTE transmitter has to achieve spurious emissions about 20 dB lower than the transmitter specifications require.

The transmit power at the input to the Low Noise Amplifier (LNA) is equal to the transmit power at the antenna plus the duplexer loss from transmit port to antenna (about 2 dB) minus the duplexer transmit-to-receiver isolation (assume 50 dB). Therefore, at the maximum LTE

¹²The duplex isolation is obtained as $NF - 9$ dB $- 2$ dB $- SE$, where NF is the noise floor (-174 dBm/Hz), 9 dB comes from the degradation of the receiver sensitivity, 2 dB comes from the loss from transmit port to antenna and SE is the Spurious Emission requirement.

UE output power of 23 dBm, the mean transmitter power leaking back to the LNA, $P_{\text{Tx_leak}}$ is

$$P_{\text{Tx_leak}} \approx 23 + 2 - 50 = -25 \text{ dBm.} \quad (21.3)$$

The equivalent figure for UMTS would be 1 dB higher.

However, most of the receiver specifications are defined with the transmitter power at 4 dB below maximum, so then the transmit leakage could be -29 dBm. This value is for the average transmit power, but the transmitted signal contains amplitude modulation and therefore the peak signal will be higher. This scenario is illustrated in Figure 21.11 which shows a simplified block diagram of the transceiver with direct-conversion architectures for both receiver and transmitter.




Figure 21.11: Transmit leakage problem in LTE FDD UE transceiver.

The transmit leakage problem presents a particularly challenging requirement if applied to RF bands with a small duplex separation. In FDD, when performing a selectivity or blocking test with a single interfering signal in one of the adjacent channels, the receiver is also exposed to its own transmitted signal as a second interferer, which, in the worst scenario, can cause cross-modulation to occur, in which the interferers mix together to generate in-band intermodulation distortion products.

21.4.3 Maximum Input Level

The maximum input level is the maximum mean received signal strength, measured at each antenna port, at which there is sufficient SINR for a specified modulation scheme to meet a minimum throughput requirement. Some wireless standards also specify a higher level at which it is guaranteed that no damage will be done to the receiver, but this is not covered in the 3GPP specifications. Although in UMTS the maximum input level is defined for the high and low SINR reference channels, in LTE it is only specified for the high SINR reference channel, on the assumption that the high SINR reference channel is about 4 dB more demanding due to the higher PAPR of the signal. The downlink requirement is for a maximum average input level of -25 dBm at any channel bandwidth with full RB allocation. For FDD operation, the requirement must be met when the transmitter is set at 4 dB below the maximum output power. The corresponding UMTS requirement is the same, applied to the combination of all the channels rather than full RB allocation.

The peak maximum input level is higher than the maximum average input level by an amount corresponding to the PAPR of the signal. PAPR is a quantity which increases the longer the duration of the measurement. If PAPR is measured on a bit-by-bit basis and plotted as a complementary cumulative probability distribution, then for the LTE downlink it reaches a value of around 11 dB for QPSK and 11.5 dB for 64QAM, over a window of 10^6 bits. In OFDMA, the PAPR is dominated by the multicarrier nature of the signal, and therefore does not vary much with modulation. The UMTS downlink signal consists of a combination of CDMA codes which also leads to a relatively high PAPR of around 8 dB.

The downlink peak maximum input level is therefore about -13.5 dBm for LTE and -17 dBm for UMTS.

As explained above, the maximum input level requirement must be satisfied with the transmitter operating at 4 dB below the maximum output power. For a UMTS UE, this can result in a mean transmit signal leakage into the receiver of -28 dBm,¹³ while the typical uplink PAPR of 3.4 dB results in a peak transmit signal leakage of -24.6 dBm. For LTE, the mean leakage from the uplink transmitter can be -29 dBm, resulting in a peak transmitter leakage of up to about -21.5 dBm for QPSK (or -20 dBm for 64QAM¹⁴). The peak value of the total received power is the combination of the received input signal and the leakage from the transmitter and thus reaches a peak of about -12 dBm for LTE and -15.2 dBm for UMTS.

Under the hypothesis considered above, the transmitter leakage raises the total received signal power by less than 2 dB.

Some operating margin is needed between this total peak received signal power and the non-linear operating region of the amplifiers in the receiver; typically this margin is chosen relative to the 1 dB compression point of the amplifiers, at which the output power is 1 dB below the expected linear gain.

¹³The mean transmit leakage is obtained from Equation (21.3) as the maximum output power (24 dBm for UMTS or 23 dB for LTE) minus 4 dB because the transmitter power is set at 4 dB below maximum, plus the duplexer loss from transmit port to antenna (about 2 dB) minus the isolation between transmitter and receiver (assume 50 dB).

¹⁴Note that only the highest category of LTE UE supports 64QAM transmission in the uplink.

21.4.4 Small Signal Requirements

21.4.4.1 SINR Requirements for Adaptive Modulation and Coding

The LTE specifications will define requirements for the demodulation error rate of the different modulation and coding schemes. An extra Implementation Margin (IM) is included to account for the difference in SINR requirement between theory and practicable implementation. This can include degradation of the signal due to any (digital) processing of the signal before the demodulator (such as filtering and re-sampling) and the use of a non-ideal demodulator, as well as the diversity gain being less than 3 dB. A 2.5 dB IM has been defined for QPSK with a 1/3-rate code in low SINR conditions; the IM will be higher for other modes. In general, 2.5 dB is a reasonable implementation margin for all QPSK modes, while 3 dB and 4 dB could be expected for 16QAM and 64QAM respectively.

Typical assumptions for the SINR values for different modulation and coding schemes are given in Table 21.7.

Table 21.7: Downlink SINR requirements.

System	Modulation	Code rate	SINR (dB)	IM (dB)	SINR+IM (dB)
LTE UE	QPSK	1/8	-5.1	2.5	-2.6
		1/5	-2.9		-0.4
		1/4	-1.7		0.8
		1/3	-1		1.5
		1/2	2		4.5
		2/3	4.3		6.8
		3/4	5.5		8.0
		4/5	6.2		8.7
	16QAM	1/2	7.9	3	10.9
		2/3	11.3		14.3
		3/4	12.2		15.2
		4/5	12.8		15.8
	64QAM	2/3	15.3	4	19.3
		3/4	17.5		21.5
4/5		18.6	22.6		
UMTS UE	QPSK	1/3	1.2	2	3.2

21.4.4.2 Thermal Noise and Receiver Noise Figure

In the LTE specifications the thermal noise density, kT , is defined to be -174 dBm/Hz where k is Boltzmann's constant (1.380662×10^{-23}) and T is the temperature of the receiver (assumed to be 15°C). No account is taken of the small variations in temperature over normal operating conditions (typically $+15^\circ$ to $+35^\circ\text{C}$) or extreme operating conditions (-10° to $+55^\circ\text{C}$).

kTB represents the thermal noise level in a specified noise bandwidth B , where $B = N_{RB} \times 180$ kHz in LTE, N_{RB} is the number of RBs and 180 kHz is the bandwidth of one RB.

The receiver Noise Figure (NF) is a measure of the SINR degradation caused by components in the RF signal chain. This includes the antenna filter losses, the noise introduced by the analogue part of the receiver, the degradation of the signal due to imperfections of the analogue part of the receiver (such as I/Q imbalance), the noise introduced by the Analogue to Digital Converter (ADC) and any other noise sources. The NF is the ratio of actual output noise to that which would remain if the receiver itself did not introduce noise.

LTE defines an NF requirement of 9 dB for the UE, the same as UMTS. This is somewhat higher than the NF of a state-of-the-art receiver, which would be in the region of 5–6 dB, with typically about 2.5 dB antenna filter insertion loss and an NF for the receiver integrated circuit of 3 dB or less. Thus, a practical 3–4 dB margin is allowed. The eNodeB requirement is for an NF of 5 dB.

21.4.4.3 Reference Sensitivity

The reference sensitivity level is the minimum mean received signal strength applied to both antenna ports at which there is sufficient SINR for the specified modulation scheme to meet a minimum throughput requirement of 95% of the maximum possible. It is measured with the transmitter operating at full power.

REFERENCE SENSitivity (REFSENS) is a range of values that can be calculated using the formula:

$$\text{REFSENS} = kTB + \text{NF} + \text{SINR} + \text{IM} - 3 \text{ (dBm)}$$

where kTB is the thermal noise level introduced above, in units of dBm, in the specified bandwidth (B), NF is the prescribed maximum noise figure for the receiver, 'SINR' is the SINR requirement for the chosen modulation and coding scheme, IM is the implementation margin and -3 dB represents the diversity gain.

For UMTS, the NF, SINR and IM are not specified separately, but the total is 12 dB. Typical REFSENS values can be calculated using the assumptions outlined above for a few example cases, as in Table 21.8. The REFSENS requirements can be found in [5] Table 7.3.1-1.

REFSENS is plotted for an LTE UE over the complete range of bandwidths and likely modulation and coding schemes in Figure 21.12.

Effect of receiver sensitivity on coverage

Using the path-loss equations described in Section 20.2, and knowing the downlink transmit power, the effect of receiver sensitivity on the practical operating coverage can be estimated. An example is shown in Figure 21.13, for the downlink with an eNodeB power of +46 dBm for a 5 MHz channel, ~900 MHz RF frequency (Band 5) and REFSENS of -100 dB and -80 dB for QPSK rate 1/3 and 64QAM rate 3/4 respectively. The corresponding throughput versus range for different MCSs is illustrated in Figure 21.14.

Table 21.8: Example of REFSENS computation for band 1

System	Modulation	Channel BW (MHz)	kTB (dBm)	NF (dB)	SINR (dB)	IM (dB)	REFSENS (dBm)
LTE UE	QPSK 1/3	5	-107.5	9	-1	2.5	-100
	QPSK 1/3	20	-101.4	9	-1	2.5	-94
	64QAM* 3/4	5	-107.5	9	17.5	4	-80
	64QAM* 3/4	20	-101.4	9	17.5	4	-74
LTE BS	QPSK 1/3	5	-107.5	5	1.5	2.5	-101.5
UMTS UE	QPSK 1/3	3.84	-108.2	9	1.2–21.1 (21.1 dB spreading gain)	2.5	-117

* Note that the REFSENS is specified only for QPSK modulation.




Figure 21.12: LTE UE reference sensitivity.




Figure 21.13: Downlink coverage taking receiver sensitivity into account.




Figure 21.14: Throughput versus coverage.

21.4.5 Selectivity and Blocking Specifications

Selectivity and blocking tests measure a receiver's ability to receive the wanted signal (to achieve $\geq 95\%$ of the maximum throughput of the reference measurement channels) at its assigned channel frequency in the presence of interfering signals in adjacent channels and beyond. As usual, this requirement must be met when the transmitter is set to 4 dB below the supported maximum output power. A low SINR is assumed. A summary of the selectivity and blocking tests for the LTE UE in a 5 MHz channel is illustrated in Figure 21.15. The requirements generally scale with bandwidth.




Figure 21.15: Selectivity and blocking requirements for a 5 MHz UE.

As with UMTS, the selectivity and blocking requirements include the case of a close Continuous-Wave¹⁵ (CW) blocking signal, some cases with modulated interferers in the first three adjacent channels, and some OOB blocker requirements with a spurious response allowance [5].

21.4.5.1 Adjacent Channel Selectivity

Adjacent Channel Selectivity (ACS) is a measure of a receiver's ability to receive a wanted signal at its assigned channel frequency in the presence of an adjacent channel interfering signal at a given frequency offset from the centre frequency of the assigned channel, without the interfering signal causing a degradation of the receiver performance beyond a specified

¹⁵A 'continuous wave' signal is an unmodulated tone.

limit. ACS is predominantly defined by the ratio of the receive filter attenuation on the assigned channel frequency to the receive filter attenuation on the adjacent channel.

For LTE, ACS is defined following the same principles as UMTS and requiring similar performance capability up to a 10 MHz bandwidth, but is more relaxed for 15 and 20 MHz bandwidths. The LTE ACS is defined for each bandwidth using a modulated LTE signal as the interferer, and is only defined for low SINR conditions.

In order to check the ability of the receiver to handle the full required dynamic range, the ACS requirement is specified for two cases – a small adjacent channel interferer power and a large adjacent channel interferer power. These cases are explained in more detail below.

ACS with a small adjacent interferer

In this ACS case, the wanted signal is, like in UMTS, 14 dB above REFSENS (given in Table 21.8) and therefore takes a different absolute level for each bandwidth. The Carrier-to-Interference Ratio (C/I) is set at -31.5 dB for bandwidths up to 10 MHz and at -28.5 dB and -25.5 dB for 15 and 20 MHz respectively; the ACS is quoted as being 33 dB including a 2.5 dB implementation margin for bandwidths up to 10 MHz and relaxed to 30 dB and 27 dB for 15 and 20 MHz respectively.

Up to a bandwidth of 5 MHz, the bandwidth of the interferer is the same as the bandwidth of the wanted signal. Above 5 MHz, the bandwidth of the interferer stays at 5 MHz, which means that the RF test equipment does not need to be able to generate a wide bandwidth interferer. For bandwidths above 5 MHz the interferer is positioned at the near edge of the channel. The consequence of this is that the interferer power is concentrated at the edge of the adjacent channel at which the filters used in the receiver have least selectivity. The digital filters will have a sharp cut-off because they are designed for OFDM, but the analogue filters in the RF front-end of the receiver could have much less attenuation at the near edge of the channel, which will push up the dynamic range at the ADC. To compensate, the ACS is relaxed by 3 dB and 6 dB for the 15 MHz and 20 MHz modes respectively as mentioned above.

Level diagrams for ACS requirements for two bandwidths (5 and 20 MHz) are shown in Figure 21.16.

The gap between the edge of the wanted signal and the edge of the interferer is a function of the used bandwidth of both signals. Given that both use a full allocation of RBs, this gap can be calculated as 1.25 MHz for 20 MHz bandwidth reducing to a minimum of 300 kHz for 3 MHz bandwidth. Below this bandwidth the channel usage reduces, so the gap increases. Note that in percentage terms, a 1.25 MHz gap adjacent to a 20 MHz channel is actually smaller than a 300 kHz gap adjacent to a 3 MHz channel (6.25% compared to 10%). The channel filters should be designed to scale with bandwidth, so it is this percentage ratio which determines the filter roll-off requirements. Comparing the 20 and 5 MHz modes we see that the percentage has roughly doubled but the specification is relaxed by 6 dB, so the filter roll-off requirements should be similar.

ACS with a large adjacent interferer

The large adjacent interferer requirement uses an adjacent channel interferer power of -25 dBm. Similarly to the case of the small adjacent interferer, the SIRC/I is fixed at -31.5 dB for bandwidths of 10 MHz and below (i.e. the wanted signal power -56.5 dBm)




Figure 21.16: Adjacent channel selectivity illustrated for Case 1 [5].

and -28.5 dB and -25.5 dB for 15 and 20 MHz respectively (i.e. the wanted signal power is -53.5 dBm and -51.5 dBm). The margin between the wanted signal power and REFSENS varies for each bandwidth in the range 41.0 to 50.2 dB, which is well above the noise floor and therefore not of much significance (see [5, Section 7.5, ACS Case 2]).

In practice, the toughest test for the receiver is likely to be somewhere between the small interferer and large interferer ACS tests. Assuming that the dynamic range of the receiver is limited, there is a point at which the interferer power first becomes high enough to require that the front-end gain needs to be reduced. This gain reduction will degrade the noise figure of the receiver at a point at which the wanted signal power is not very high. The gain control algorithm used by the receiver must be well planned to avoid such problems.

Adjacent Channel Interference Ratio (ACIR)

The ACS and ACLR (see Section 21.3.2.1) together give the Adjacent Channel Interference Ratio (ACIR). The ACIR is the ratio of the total power transmitted from a source to the total interference power affecting a victim receiver, resulting from both transmitter and receiver imperfections.

It follows that

$$ACIR \cong \frac{1}{\frac{1}{ACLR} + \frac{1}{ACS}}$$

ACLR and ACS have been extensively used for coexistence studies.

21.4.5.2 Narrowband Blocking (in Adjacent Channel)

The blocking characteristic is a measure of the receiver's ability to receive a wanted signal (to achieve $\geq 95\%$ of the maximum throughput of the reference measurement channels) at its assigned channel frequency in the presence of an unmodulated unwanted interferer on frequencies other than those of the spurious response or the adjacent channels, without this unwanted input signal causing a degradation of the performance of the receiver beyond a specified limit. The blocking performance applies at all frequencies except those at which a spurious response occurs.

The narrowband blocking specification is a severe test of the receiver's ability to reject 3rd order intermodulation products resulting from cross-modulation of the transmitter leakage which appears around the narrowband blocker. The frequency of the unwanted cross-modulation product depends only on the narrowband blocker frequency and not on the frequency of the transmitter, or any other modulated blocker.

The LTE 'narrowband' blocking performance requirement uses a CW interferer very close to the wanted signal at an offset less than the nominal channel spacing. For such small offsets nearly half of the transmitter leakage will appear in-band.

The CW blocker is positioned at approximately 200 kHz from the near edge of the adjacent channel. For example, for a 5 MHz bandwidth the offset is ~ 2.7 MHz, which is 450 kHz from the edge of the wanted signal. The offset of the blocker from the edge of the wanted signal reduces with bandwidth, reaching just 350 kHz for the 3 MHz channel bandwidth. Below 3 MHz bandwidth the channel occupancy reduces, which compensates for the reduced gap.

For this test, for bandwidths of 10 MHz and below, the power of the wanted signal is set at the REFSENS level plus a bandwidth-specific offset (22, 18, 16, 13, 14 and 16 dB for 1.4, 3, 5, 10, 15 and 20 MHz respectively – see [5, Section 7.6.3]); the blocker power is set to -55 dBm independently from the bandwidth. Compared to the ACS test, the gap between the wanted and interfering signals is between 30 and 50 kHz less, which makes the narrowband blocking test a little more demanding than the ACS test.

In UMTS the wanted signal power is 13 dB higher than the REFSENS and the interferer power is set to -56 dBm. The most important difference between the UMTS narrowband blocking specification and the LTE case is that the blocker for UMTS is actually a GMSK¹⁶ modulated signal, not CW. A GMSK signal is a little narrower than QPSK (although clearly not as narrow as a CW signal), and the modulation has a constant envelope so there is no PAPR variation which could increase non-linear distortion. On balance, it can be concluded that the UMTS and LTE narrowband blocking specifications are similarly demanding.

21.4.5.3 Non-Adjacent Channel Selectivity (In-Band Blocking)

Non-adjacent Channel Selectivity (NACS) is a measure of the receiver's ability to receive a wanted signal at its assigned channel frequency in the presence of unwanted interfering

¹⁶Gaussian Minimum-Shift Keying, as used in GSM.

signals falling into the receive band beyond the adjacent channel or at less than 15 MHz offset from the edge of the receive band. The interfering signals are modulated and occupy the same bandwidths as specified for ACS. The LTE specifications refer to this test as ‘in-band blocking’, although, unlike the other blocking tests, NACS does not use CW signals.

There are two requirements to be met, the first with an interferer of -56 dBm¹⁷ in the second adjacent channel or further (referred to as Case 1), the second with an interferer of -44 dBm¹⁸ in the third adjacent channel or any larger frequency offset up to 15 MHz out of band (referred to as Case 2). Furthermore, a specific requirement which applies to assigned UE channel bandwidth of 5 MHz and for Band 17 (referred to as Case 3) is also provided [5].

Unlike the ACS tests, NACS does not need to be repeated at higher signal levels, so it does not test dynamic range to the same extent. However, the wanted signal level is much lower, at just 6 dB above REFSENS for bandwidths up to 10 MHz, and 7 dB and 9 dB above REFSENS for 15 and 20 MHz bandwidths respectively. Consequently the C/I ratios are much lower, falling for example to -50 dB for the third adjacent channel.¹⁹ The total filtering requirement will therefore be of the order of -60 dB at three times the bandwidth.

A summary of the NACS requirements for LTE and UMTS UEs is shown in Tables 21.9 and 21.10. For UMTS, the wanted signal power includes 21 dB spreading gain.

Table 21.9: Non-adjacent ($N \pm 2$) channel selectivity.

System	LTE						UMTS
Bandwidth of wanted signal (MHz)	1.4	3	5	10	15	20	3.84
Own signal power above REFSENS (dB)	6	6	6	6	7	9	3
Power of interferer (dBm)	-56						
Frequency offset of interferer (MHz)	± 2.8	± 6	± 10	± 12.5	± 15	± 17.5	± 10
Bandwidth of interferer (MHz)	1.4	3	5	5	5	5	3.84

21.4.5.4 Out-of-Band Blocking

The LTE OOB blocking tests measure the receiver’s ability to receive a wanted signal at its assigned channel frequency in the presence of unwanted interfering signals falling outside the receive band, at 15 MHz or more offset from the edge of the band.

The wanted signal power is at the same level as for the in-band blocking test (at 6 dB above REFSENS for bandwidths of up to 10 MHz, and relaxed to 7 or 9 dB above REFSENS for bandwidths of 15 MHz or 20 MHz respectively).

¹⁷Same as for UMTS.

¹⁸Same as for UMTS.

¹⁹This example is for a REFSENS of -100 dBm for Band 1.

Table 21.10: Non-adjacent ($N \pm 3$) channel selectivity.

System	LTE						UMTS
Bandwidth of wanted signal (MHz)	1.4	3	5	10	15	20	3.84
Own signal power above REFSENS (dB)	6	6	6	6	7	9	3
Power of interferer (dBm)	-44						
Frequency offset of interferer (MHz)	± 4.2	± 9	± 15	± 17.5	± 20	± 22.5	± 15
Bandwidth of interferer (MHz)	1.4	3	5	5	5	5	3.84

The blocker is a CW signal with a power of -44 dBm at 15 to 60 MHz offset, -30 dBm from 60 to 85 MHz offset and -15 dBm from 85 to 12750 MHz offset.²⁰ These values are all the same as UMTS.

The actual offset of the blocker is the specified offset from the band edge plus half the wanted signal bandwidth plus the RF guard band, which is 2.4 MHz. Hence, for example, for a 5 MHz bandwidth, the -44 dBm blocker is at a minimum offset of $15 + 2.5 + 2.4 = 19.9$ MHz, which is close to four times the bandwidth.

As in UMTS, there is an additional requirement for RF bands 2 (DCS1800), 5 (GSM850), 12 and 17, comprising a -15 dBm blocker coming from the transmit band, which is at an offset of just 20 MHz for bands 2 and 5, 12 MHz for band 12 and 18 MHz for band 17, see Table 21.1. This latter requirement is clearly far tougher than all the other blocking specifications.

A summary of the selectivity and blocking requirements is shown in Figure 21.17. This includes the specified 2 or 3 dB margin added to the C/I requirements.

When normalizing all frequency offsets to the wanted signal bandwidth, it can be seen that the 20 MHz LTE bandwidth requires the most severe filter frequency response relative to its bandwidth (and also to the digital sampling frequency). Additionally, the LTE 5 MHz selectivity requirement is relatively tougher than UMTS.

21.4.5.5 Spurious Response Specifications

Frequencies for which the throughput does not meet the requirements of the OOB blocking test are called spurious response frequencies. Spurious responses occur at specific frequencies at which an interfering signal mixes with the fundamental or harmonic of the receiver local oscillator and produces an unwanted baseband frequency component. Spurious responses are measured by recording when the OOB blocking test is not passed.

At the spurious response frequencies, the LTE receiver must still achieve the required throughput with a -44 dBm CW blocker and a wanted signal level set to the power level specified in the OOB blocking test.

²⁰This blocker test is defined for UEs supporting all bands a part from 2, 5, 12 and 17.




Figure 21.17: Selectivity and blocking requirements.

21.4.6 Spurious Emissions

The spurious emissions power is the power of emissions generated or amplified in a receiver which appear at the antenna connector. For LTE, limits are set covering the range 30 MHz to 12.75 GHz. The LTE specification is the same as UMTS, namely not more than -57 dBm between 30 MHz and 1 GHz (measured in 100 kHz bandwidth), and -47 dBm from 1 GHz to 12.75 GHz (measured in 1 MHz bandwidth); these requirements correspond to -107 dBm/Hz across the full range.

For UMTS, some additional spurious emissions requirements are specified for the cellular bands. Generally the requirement is -60 dBm measured in 100 kHz bandwidth (-110 dBm/Hz), but there are some more demanding requirements, of which the toughest is for the 935–960 MHz band to protect the GSM downlink receive band, i.e. -79 dBm measured in 100 kHz (-129 dBm/Hz).

Spurious emissions are caused by power from the local oscillator or the amplified received signal leaking back to the antenna. The required isolation from the amplified received signal can be estimated as follows. The maximum received power is -25 dBm, and the minimum received bandwidth 1.08 MHz; thus at the maximum input signal density the power in the 100 kHz measurement bandwidth is -35.3 dBm. The isolation required to prevent the amplified received signal from breaking the spurious emission limit will depend on the amplifier gain, which will vary with the signal and is likely to be at its minimum at the maximum input signal level. Assuming a 20 dB amplifier gain, then the isolation would need to be about 64 dB to reach the -79 dBm spurious emission limit. This requirement could be tougher than for UMTS because spreading enables the signal levels to be lower. However,

the toughest isolation requirement is likely to be for the local oscillator signal which drives the mixer, since it typically operates at a higher power, for example around -4 dBm. The worst case isolation requirement for the local oscillator is thus about 75 dB, higher than the received signal isolation requirement.

21.4.7 Intermodulation Requirements

21.4.7.1 Intermodulation Distortion

As mentioned in Section 21.3.2.2, when two or more tones are present in a non-linear device, such as an amplifier in a receiver, intermodulation products are created. A power series describes all of the possible combinations of generated frequencies. Intermodulation products generated from mixing with a tone outside the wanted band may themselves fall into the wanted band. The problem is more severe for in-band interfering tones, since in that case the RF filter provides no attenuation.

As a device is driven further into its non-linear region, the amplitudes of the intermodulation products increase, while the power of the original tones at the output decreases. If the device was not limited in output power, then the powers of the intermodulation products would increase to the level of the original tones.

Intermodulation response rejection is a measure of the capability of the receiver to receive a wanted signal on its assigned channel frequency in the presence of two or more interfering signals, where the frequencies of the interfering signals relative to the wanted signal are such that the InterModulation Distortion (IMD) products fall into the wanted signal band.

The strongest interference is caused by third-order intermodulation products. The standard procedure for measuring the third-order intermodulation performance of a receiver is with two CW interfering tones offset from the wanted carrier frequency, with the offset of one tone being twice that of the other. Positioning one tone with twice the offset of the other results in the IMD product which occurs at the difference frequency falling into the wanted band. Such a test scenario is defined for UMTS, where it is referred to as a *narrowband intermodulation requirement*. UMTS also includes a *wideband intermodulation* test, with interferers at larger offsets, where one of the interferers is a modulated UMTS signal. This is designed to assess the effect of IMD products arising from UMTS transmissions in other channels.

For LTE, a narrowband intermodulation requirement could be defined following the same principles as in UMTS, using two CW signals as interferers. However, as a narrow bandwidth IMD product falling on top of the wanted signal would only degrade the performance of few subcarriers, the impact on throughput would be minimal. Hence, this is not defined in the current specification.

For the LTE wideband intermodulation performance requirement, the modulated interferer is defined to have the same bandwidth as the ACS test, denoted here ' BW_{int2} '.²¹ The CW interferer is at offset $BW_{channel} + 1.5 \times BW_{int2}$, where $BW_{channel}$ is the channel bandwidth; thus for the 5 MHz bandwidth case the CW interferer is at a 10 MHz offset and the modulated interferer at 20 MHz offset, twice the offset of the CW interferer, as for UMTS. The interferer powers, are both -46 dBm, the same as for UMTS. The wanted signal is at a variable level between 6 and 12 dB above REFSENS depending on the channel bandwidth. This variation

²¹ BW_{int2} is equal to the channel bandwidth for the 1.4, 3 and 5 MHz cases, and it is limited to 5 MHz for higher channel bandwidths.

of the wanted signal level is designed to keep the wanted signal at an absolute value close to -94 dBm and thus ensure a consistent IMD requirement for each bandwidth. For the 5 MHz bandwidth case, the wanted signal is 6 dB above REFSENS, 3 dB higher than UMTS (because of the 3 dB diversity gain).

For bandwidths of 10 MHz and above, the modulated interferer is always 5 MHz. The consequence of this is that the power of the interferer is spread over a narrower bandwidth than the wanted signal, so the interfering power on each subcarrier is higher than the wanted power of each subcarrier by the ratio of the bandwidths. This increases the IMD requirement: for example, for the 20 MHz bandwidth, a value of $10 \times \log_{10}(\text{BW}_{\text{channel}}/\text{BW}_{\text{int2}}) = 6$ dB should be taken into account when calculating IMD requirements.

The usual way to quantify the IMD performance of a receiver is to determine the *third-order intercept point*, or IP3. IP3 is a theoretical point where the amplitudes of the intermodulation tones are equal to the amplitude of the fundamental tones. The third-order IMD products, IM3, increase as the cube of the power of the input tones (i.e. at 60 dB/decade, three times the rate of the fundamental). Therefore, at some power level the distortion products will overtake the fundamental signal; the point at which the curves of the fundamental signal and the IM3 intersect on a log-log scale is the IP3. The corresponding input power level is known as Input IP3 (IIP3), and the output power when this occurs is the Output IP3 (OIP3) point, which is higher than IIP3 by an amount equal to the gain of the receiver. The IIP3 can be computed as follows:

$$\text{IIP3} = (3P_{\text{in}} - P_{\text{IMD3,in}})/2 = \text{OIP3} - G \quad (21.4)$$

where P_{in} is the power of a single tone at the input of the system, $P_{\text{IMD3,in}}$ the input-referred third-order intermodulation distortion product and G the passband gain (all in dB).

It is not possible to measure IP3 directly, because by the time the non-linear amplifier reached this point it would be heavily overloaded. Instead the amplifier is measured at a lower input tone power and IP3 is found by extrapolation.

The wide-band intermodulation test can be used to derive an equivalent IIP3 requirement. The IM3 products fall in-band and add to the existing receiver noise. The maximum power of the IM3 products that can be tolerated, referred back to the input (i.e. so that the gain of the receiver is factored out), $P_{\text{IMD3,in}}$, is such that, when combined with the existing noise, it reaches a level that is equal to the maximum tolerable noise floor. This noise floor is below the wanted signal power by the SINR requirement plus implementation margin. Thus, if the wanted signal is 3 dB above REFSENS the IM3 products and existing noise can be the same, but as the wanted signal level is further raised above the REFSENS level the existing noise floor will be less significant and the IM3 products a little higher.

As an example, we can calculate the IIP3 requirement for a 15 MHz bandwidth UE supporting band 1. The SINR requirement plus implementation margin is assumed to be 1.5 dB. The wanted signal is 7 dB above REFSENS (i.e. at -88.2 dBm), so it can be calculated that the IM3 products must be 0.97 dB²² below the maximum tolerable noise floor. The factor of three difference in bandwidth must be accounted for by a factor of 4.77 dB.²³ Thus, $P_{\text{IMD3,in}} = -88.2 - 1.5 - 0.97 - 4.77 = -95.4$ dBm and $\text{IIP3} =$

²²The IM3 noise plus the thermal noise combines to use up to 7 dB margin, $0.97 = -10 \cdot \log_{10}(1 - 10^{-7/10})$, therefore $10 \cdot \log_{10}(10^{-7/10} + 10^{-0.97/10}) = 0$.

²³This comes from the fact that the modulated interferer always has a 5 MHz bandwidth when the channel bandwidth is 10 MHz or more.

$[(3 \times (-46)) + 95.4]/2 = -21.3$ dBm.²⁴ IIP3 can be calculated for the other bandwidths and reaches a maximum of -20.6 dBm for 5 or 10 MHz bandwidth.

The IIP3 requirement for UMTS can be calculated in a similar manner.

21.4.7.2 Out-of-Band Intermodulation Distortion

As explained above, the LTE transmit signal may leak into the receiver. If a blocker is also present, IM3 products may result. If the blocker is half-way between the transmit and receive bands (i.e. between 15 and 200 MHz from the wanted band, depending on the duplex spacing) then the IM3 products will fall into the wanted band. For Band 5, the duplex spacing is 45 MHz, so a blocker at a 22.5 MHz offset, which can have a power of -44 dBm, would be the worst case. For the OOB blocking tests, the wanted signal is 3 dB above REFSENS. Therefore $P_{\text{IMD3,in}}$ can be calculated as described above (equal to -101.5 dBm for the 5 MHz bandwidth). If we assume the transmitter leakage is -29 dBm and the blocker is at -44 dBm, then the average input power is -31.9 dBm. Thus the OOB IP3 requirement is $\text{IIP3} = [(3 \times (-31.9)) + 101.5]/2 = +2.9$ dBm and the maximum IIP3 is $+6$ dBm for the 1.4 MHz bandwidth.

For Band 1 the duplex spacing is 190 MHz, so a blocker at 95 MHz, which can be at -15 dBm, would cause IMD. If the RF filter gives less than 29 dB attenuation of the blocker, then this band would present an even tougher OOB IIP3 requirement.

A similar calculation for UMTS for Band 5, assuming $+20$ dBm transmit signal leaking into the receiver at -28 dBm and a -44 dBm blocker, would yield a $P_{\text{IMD3,in}}$ requirement of -98 dBm and an OOB IIP3 of $+2.7$ dBm.

Since the receiver's IP3 is dominated by that of the mixer, the problem of cross-modulation can be reduced for the most demanding bands by using a band-pass filter between the LNA and the mixer, so that most of the cross-modulation occurs in the LNA.

21.4.8 Dynamic Range

The input dynamic range is a key factor affecting the cost of the receiver. The larger the dynamic range, the larger must be the linear operating region of the receiver components. There are many ways to analyse the dynamic range and resulting signal handling requirements for a receiver, with widely differing results.

The simplest definition of the dynamic range of a receiver is the ratio of the maximum and minimum signal levels required to maintain a specified throughput or error rate. For UMTS the maximum and minimum wanted signal powers are -44 and -117 dBm respectively, giving a dynamic range of 73 dB. For an LTE UE, the maximum input is -25 dBm (assumed to be applicable to any modulation), while the minimum signal level for a 5 MHz bandwidth is -100 dBm (minimum REFSENS value across the different bands), giving a maximum dynamic range of 75 dB.

Dynamic range can also be specified as the ratio between the maximum signal and the noise floor. Even more usefully, the maximum signal level could include a margin to allow for the variation of the peak signal power above its average (PAPR). Such a measure of dynamic range gives an indication of the total signal handling requirements of the receiver in the absence of gain control. In Section 21.4.3, the total peak received signal power was

²⁴ -46 dBm is the interferer power as specified in [5].

derived, which includes the PAPR and the contribution from the transmitter leakage; the values were -12.0 dBm for LTE and -15.2 dBm for UMTS. For the LTE 5 MHz bandwidth the noise floor at the antenna is -107.5 dBm for UMTS it is -108.2 dBm. Thus this dynamic range estimate is $-12.0 - (-107.5) = 94.9$ dB for LTE and $-15.2 - (-108.2) = 92.9$ dB for UMTS (note that for UMTS the REFSSENS, -117 dBm, is below the noise floor).

In practice, gain control for large signals is normally used in a receiver to reduce the dynamic range. The minimum dynamic range that the receiver should handle linearly, measured at the antenna, is the maximum SINR requirement, plus implementation loss, plus NF, plus a margin for PAPR (assumed to be 11.5 dB). For LTE this gives $17.5 + 4 + 9 + 11.5 = 42$ dB, while for UMTS using QPSK it is only $12.2 + 8 = 20.2$ dB.

However, this does not take account of any interferers which could be present, nor leakage from the transmitter. At REFSSENS, for an FDD LTE UE, there could be transmitter leakage at -25 dBm (peak about -16 dBm); the difference between the peak transmitter leakage and the noise floor is 91.5 dB for LTE and 83.6 dB for UMTS, although the frequency offset is large so the analogue filters will reduce the interference.

In conclusion, it is clear that the higher SINR and higher PAPR requirements of LTE present a higher signal-handling dynamic range requirement than for UMTS, but the linearity requirements due to interferers are similar.

21.5 RF Impairments

The RF parts of the transmitter and receiver are comprised of non-ideal components which can have a strong impact on the ultimate demodulation performance. Different access technologies have different sensitivities to these RF non-idealities. For example, OFDM-based systems are particularly sensitive to any distortion which may remove the orthogonality between the subcarriers, resulting in Inter-Carrier Interference (ICI), as discussed in Section 5.2.3. The RF impairments can have a non-negligible impact on Bit Error Rate (BER), as shown for example in [11, 12] and references therein, where some compensation algorithms are also discussed.

The goal of this section is to show how the most common RF impairments introduce errors in an OFDM signal. First a simplified model of the impairments is discussed for the purpose of establishing SINR limitations for LTE performance assessment. This is followed in Section 21.5.2 by a mathematical framework for analysing the sensitivity of the LTE to particular impairments.

21.5.1 Transmitter RF Impairments

A generic model of the typical RF impairments of a transmitter is given in Figure 21.18. It comprises the following components:²⁵

- **Digital to Analogue Converter (DAC):** a source of quantization noise – here we assume a uniform linear quantizer;
- **Up-sampling function:** up-samples the OFDMA/SC-FDMA sample rate process by a factor (not shown in diagram);

²⁵It is not implied that a real implementation would necessarily use exactly the components specified; the aim here is to give a representative view of the typical impairments.

- **Baseband equivalent filter:** equivalent to the concatenation of linear filtering components in the transmit path, including digital and analogue elements;
- **Quadrature error component:** corresponding to the loss of I/Q orthogonality in the frequency conversion process;
- **d.c. offset:** arising from, for example, direct conversion Local Oscillator (LO) leakage effects;
- **Phase noise:** corresponding to the LO phase noise process;
- **Frequency error:** due to LO long-term frequency error;
- **InterModulation Distortion (IMD):** attributable to, for example, PA non-linearity.

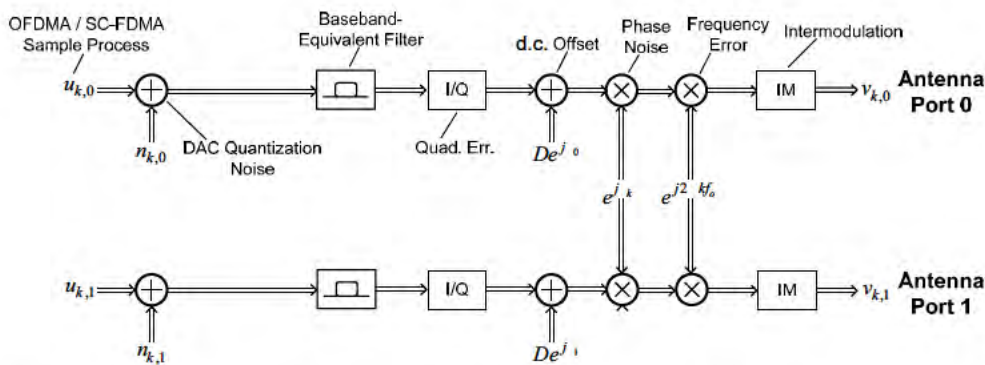


Figure 21.18: Model of multi-antenna transmitter impairments. Reproduced by permission of © 2006 Motorola.

A representative model of the RF impairments in the receiver, shown in Figure 21.19 for two receiver antennas, comprises the following conceptual components:

- **Adjacent Channel Interference (ACI):** important when establishing deployment and coexistence guidelines and for estimating total system spectral efficiency. The effect of ACI may be neglected, however, in an initial LTE system performance analysis by assuming that Co-Channel Interference (CCI) is the dominant interference-limiting effect.
- **Antenna Gain Imbalance (AGI):** in practical multiantenna UE designs, the antenna gains can be affected by multiple factors including user grip pattern, orientation and proximity to the user’s body.
- **InterModulation (IM):** receiver non-linearity in general, is a critical consideration when operating at the upper limits of the receiver’s dynamic range and in the presence of strong adjacent channel interference.
- **Thermal Noise:** applicable noise figures for LTE UE and eNodeB devices are specified in [13]. One further practical consideration is the potential for non-uniform noise figures applicable to each LTE antenna port.

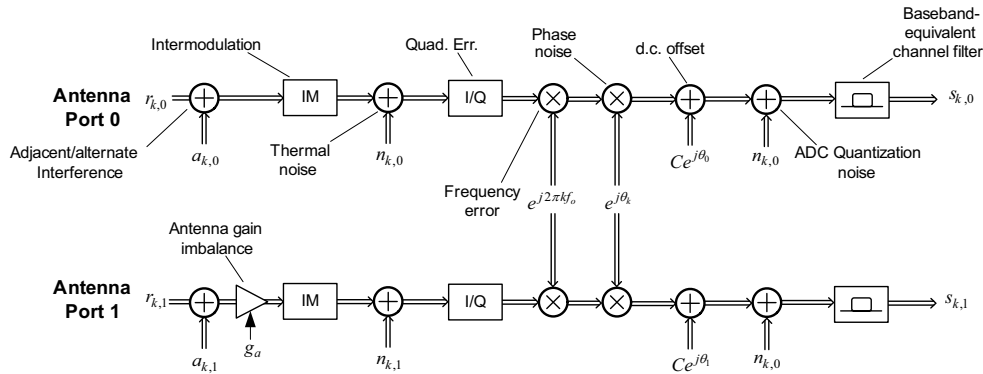


Figure 21.19: Model of multi-antenna receiver impairments. Reproduced by permission of © 2006 Motorola.

- **Quadrature error component:** as with the transmitter, this element models the loss of quadrature in the frequency conversion process. As an initial assumption, quadrature error may be neglected in eNodeB receivers, but is an essential element in direct conversion UE receiver modelling.
- **Frequency error:** the eNodeB receiver frequency error attributed to eNodeB LO error may be neglected since the UE uses the downlink waveform as a frequency reference. Clearly, in some circumstances there can be a significant frequency shift between the downlink signal received by the UE and the resulting uplink signal observed by the eNodeB.
- **Phase noise:** this corresponds to the eNodeB and UE LO phase noise process.
- **d.c. offset:** as for the transmitter model, this can arise due to LO leakage effects.
- **Analogue to Digital Converter (ADC):** similarly to the transmitter, this can be modelled as a quantization noise source.
- **Baseband equivalent filter:** equivalent to the concatenation of linear filtering components in the receive path, including digital and analogue elements.
- **Down-sampling function (not shown):** decimates the OFDMA/SC-FDMA sample rate process by a factor.

Often, a simple additive white Gaussian model can be considered to simplify the analysis of the impact of the RF impairments on link performance (see, for example, [11]). In fact the distortion generated by any non-linear device present for example in the downlink transmitter²⁶ can be modelled by using Bussgang's theorem [14] as follows:

$$\tilde{\mathbf{x}} = \alpha \mathbf{x} + \mathbf{d} \quad (21.5)$$

where $\tilde{\mathbf{x}}$ is the OFDM signal vector distorted by non-linearities, \mathbf{x} is the column vector of the ideal OFDM symbol, \mathbf{d} is the vector of the equivalent interference term due to distortions

²⁶The analysis here considers the downlink OFDM transmitter, but it can equally well be applied to the uplink transmitter.

which are uncorrelated with the signal \mathbf{x} and α is a complex gain factor accounting for the attenuation and phase rotation.

The parameter α can be derived as

$$\alpha = \frac{\mathbb{E}[\tilde{\mathbf{x}}^H \mathbf{x}]}{\mathbb{E}[\mathbf{x}^H \mathbf{x}]} \quad (21.6)$$

where $\mathbb{E}[\cdot]$ is the expectation operator.

Assuming, for the sake of simplicity, an AWGN²⁷ propagation channel, the distorted signal $\tilde{\mathbf{x}}$ generated by the RF transmitter is further corrupted by complex white circular Gaussian noise $\mathbf{n} \sim \mathcal{N}(0, \sigma_n^2 \mathbf{I})$ such that the signal received by the OFDM receiver at the output of the FFT can be expressed as

$$\mathbf{R} = \alpha \mathbf{F} \mathbf{x} + \mathbf{F} \mathbf{d} + \mathbf{F} \mathbf{n}$$

Assuming coherent detection in the OFDM receiver,²⁸ the phase rotation induced by the scaling factor α has no effect as it will be ideally compensated in the receiver's channel estimation process. However, in both coherent and non-coherent OFDM systems, the impact of the magnitude of α translates into a power penalty for the useful signal. The reader is referred to [14] for further details.

Assuming FFT processing of large blocks and applying the central limit theorem, the frequency-domain distortion term $\mathbf{D} = \mathbf{F} \mathbf{d}$ can be closely approximated by complex white circular Gaussian noise $\mathbf{D} \sim \mathcal{N}(0, \sigma_d^2 \mathbf{I})$, and the SINR of the distorted received OFDM signal can be simply expressed as

$$\text{SINR}_d = \frac{|\alpha|^2 \text{tr}\{\mathbf{C}_{\text{XX}}\}}{\text{tr}\{\mathbf{C}_{\text{DD}}\} + \sigma_n^2} = \frac{|\alpha|^2}{\sigma_d^2 + \sigma_n^2} \quad (21.7)$$

where $\mathbf{C}_{\text{XX}} = \mathbf{F} \mathbf{C}_{\text{xx}} \mathbf{F}^H$ and $\mathbf{C}_{\text{DD}} = \mathbf{F} \mathbf{C}_{\text{dd}} \mathbf{F}^H$, assuming that the signal power \mathbf{X} is normalized.

This model can be used in determining the impact of RF impairments on the LTE performance, with the additive distortion term being modelled as a transmitter EVM source.

21.5.2 Model of the Main RF Impairments

In order to gain insight into the impact of typical impairments, not only in terms of the impact on a single-user's communication link but also in terms of their overall effect on system-level spectral efficiency, it is instructive to develop a more precise analytical model. In particular, three sources of errors are included here:

- Phase noise;
- Modulator image interference (amplitude and phase imbalance);
- Non-linear distortion and intermodulation products which fall into adjacent RBs.

²⁷Additive White Gaussian Noise. Note that the reasoning can be straightforwardly extended to the case of a frequency-selective convolutive channel and an OFDM system with CP.

²⁸Coherent detection means that the phase of the received signal is known to the receiver a priori, for example from Reference Signals (RSs) as discussed in Chapter 8.

The model here is developed considering an OFDM signal (downlink) but it can be easily generalized for an SC-FDMA signal (uplink). The OFDM signal can be written as the sum of closely-spaced tones which do not interfere with each other due to their orthogonality. In particular the discrete-time signal for symbol ℓ is written as

$$x_\ell[n] = \sum_{k=-N/2}^{N/2-1} S_{k,\ell} \exp\left[j2\pi k\Delta f \frac{n}{N}\right] \quad (21.8)$$

where N is the number of subcarriers, $S_{k,\ell}$ is the constellation symbol sent on the k^{th} subcarrier for the ℓ^{th} OFDM symbol and n represents the component within the ℓ^{th} OFDM symbol, $n \in [0, N - 1]$. In the following, we can drop the dependency from the variable ℓ without loss of generality.

21.5.2.1 Amplitude and Phase Imbalance

Let us consider the continuous time version of the signal $x(t)$ modulated to frequency f_c , i.e.

$$r(t) = \text{Re}\{x(t)e^{j2\pi f_c t}\} = x(t)e^{j2\pi f_c t} + x^*(t)e^{-j2\pi f_c t} \quad (21.9)$$

With a non-ideal I/Q (de-)modulator, the recovered signal can be written as follows:

$$\begin{aligned} y_{\text{IQ}}(t) &= \text{LP}\{r(t)(\cos(2\pi f_c t) - j\beta \sin(2\pi f_c t + \phi))\} \\ &\stackrel{(1)}{=} \text{LP}\{(x(t)e^{j2\pi f_c t} + x^*(t)e^{-j2\pi f_c t})(\gamma_1 e^{-j2\pi f_c t} + \gamma_2 e^{j2\pi f_c t})\} \\ &\stackrel{(2)}{=} x(t)\gamma_1 + x^*(t)\gamma_2 \end{aligned} \quad (21.10)$$

where LP means Low-Pass filtering, (1) is because of Equation (21.9) with $\gamma_1 = (1 + \beta e^{-j\phi})/2$ and $\gamma_2 = (1 - \beta e^{-j\phi})/2$ and (2) is obtained by low-pass filtering the signal. By substituting Equation (21.8) into (21.10), after a variable change it follows that

$$y_{\text{IQ}}(t) = \sum_{k=-N/2}^{N/2-1} \left(\frac{1 + \beta e^{-j\phi}}{2} S_{k,\ell} + \frac{1 - \beta e^{-j\phi}}{2} S_{-k,\ell}^* \right) \exp\left[j2\pi k\Delta f \frac{t}{N}\right] \quad (21.11)$$

By letting $\beta = 1 + q$ and $Q(\beta, \phi) = (q - j\phi - qj\phi)/2$, using the approximation that $\cos(\phi) = 1$ and $\sin(\phi) = \phi$ for small angle ϕ , it follows that

$$\begin{aligned} y_{\text{IQ}}(t) &= \sum_{k=-N/2}^{N/2-1} S_{k,\ell} \exp\left[j2\pi k\Delta f \frac{t}{N}\right] \\ &\quad + Q(\beta, \phi) \sum_{k=-N/2}^{N/2-1} (S_{k,\ell} - S_{-k,\ell}^*) \exp\left[j2\pi k\Delta f \frac{t}{N}\right] \end{aligned} \quad (21.12)$$

Equation (21.12) shows clearly that the I/Q imbalance creates two different error terms: the first is the self-interference created by the same signal at the same subcarrier frequency, while the second is the signal at the frequency mirror-image subcarrier. The ideal case corresponds

to $q = 0$ and $\phi = 0$. Note also that the amplitude imbalance and the phase imbalance create the same effect, i.e. only the factor $Q(\beta, \phi)$ changes, in case of only amplitude imbalance being present ($\phi = 0$) $Q(\beta, \phi) = q/2$, and when only phase imbalance is present the multiplicative coefficient becomes $Q(\beta, \phi) = j\phi/2$.

For high-order modulation (increasing number of states in the constellation) the error term acts as noise and it spreads the constellation points as shown in [15].

In general the amplitude and phase imbalance can vary depending on the frequency, on a subcarrier basis. This happens when there is a timing offset between the in-phase and quadrature signal paths.

21.5.2.2 Phase Noise

When the LO frequency at the transmitter is not matched to the LO at the receiver, the frequency difference implies a shift of the received signal spectrum at the baseband. In OFDM, this creates a misalignment between the bins of the FFT and the peaks of the sinc pulses of the received signal. This results in a loss of orthogonality between the subcarriers and a leakage between them. Each subcarrier interferes with every other (although the effect is strongest on adjacent subcarriers) and, as there are many subcarriers, this is a random process equivalent to Gaussian noise. Thus, a LO frequency offset lowers the SINR of the receiver. An LTE receiver will need to track and compensate for this LO offset and quickly reduce it to substantially less than the 15 kHz subcarrier spacing. This is a tougher requirement than for UMTS. The phase noise impairment has been widely studied both for a single antenna [15–17] and for the MIMO case [18], especially in the context of WiMAX.

The ideal baseband transmitted signal, neglecting the additive white Gaussian noise, is given in Equation (21.8). The received baseband signal, in the presence of only phase noise can be written as

$$y_{\theta}(t) = \sum_{k=-N/2}^{N/2-1} S_{k,\ell} \exp\left[j2\pi k\Delta f \frac{t}{N}\right] e^{j\theta(t)} \quad (21.13)$$

where θ is the phase noise. In [11] it is shown that the single-sideband phase noise power follows a Lorentzian spectrum [11, 19]:

$$L(f) = \frac{2}{\pi\Delta f_{3\text{dB}}} \frac{1}{1 + [2f/(\Delta f_{3\text{dB}})]^2} \quad (21.14)$$

where $\Delta f_{3\text{dB}}$ is the two-sided 3 dB bandwidth of phase noise. The power spectrum given in Equation (21.14) can be considered as an approximation to practical oscillator spectra.

After the FFT, the symbol transmitted on the m^{th} subcarrier in the ℓ^{th} OFDM symbol can be written as

$$z_{m,\ell} = S_{m,\ell} \frac{1}{N} \sum_{n=-N/2}^{N/2-1} e^{j\theta(n)} + \frac{1}{N} \sum_{k=-N/2, k \neq m}^{N/2-1} S_{k,\ell} \sum_{n=-N/2}^{N/2-1} \exp\left[j2\pi\Delta f \frac{n}{N}(k-m)\right] e^{j\theta(n)} \quad (21.15)$$

By defining $C(k) = (1/N) \sum_{n=-N/2}^{N/2-1} e^{j2\pi\Delta f nk/N + j\theta(n)}$, as in [16], Equation (21.15) can be rewritten as

$$z_{m,\ell} = S_{m,\ell} C(0) + \sum_{k=-N/2, k \neq m}^{N/2-1} S_{k,\ell} C(k-m) \quad (21.16)$$

Equation (21.16) shows that the effect of phase noise is twofold: the useful symbol transmitted on the m^{th} subcarrier is scaled by a coefficient $C(0)$ which depends on the phase noise realization on the l^{th} OFDM symbol, but it is independent of the subcarrier index – i.e. all the subcarriers are rotated by the same quantity $C(0)$. This is referred to as the Common Phase Rotation (CPR). This term can be estimated from the RSs and removed. In LTE, four symbols per subframe contain RS (see Section 8.2), so that in theory low frequency phase noise up to about 2 kHz can be compensated. However common loop bandwidths of PLLs²⁹ are in general in the order of 10 to 100 kHz, so that a major part of the phase noise energy is outside the region which can be compensated. The second term causes Inter-Carrier Interference (ICI), the amount of which is given by the coefficient $C(k - m)$ for subcarrier k interfering on subcarrier m . The ICI due to phase noise creates a fuzzy constellation [15]. In [18] the above equations are generalized for the MIMO case, showing that phase noise has similar effects in that case.

Time-domain offsets

Any timing synchronization error results in a misalignment between the samples contained in the OFDM symbol and the samples actually processed by the FFT. If the timing error is large, some samples could be from the wrong symbol, which would cause a serious error. It is more likely that the timing error will be just a few samples and the presence of the cyclic prefix should give enough margin to prevent samples from the wrong symbol being used. Assuming this is the case, the shift in time is equivalent to a linearly-increasing phase rotation of the constellation points.

The use of higher modulation schemes and wider channel bandwidths mean that timing synchronization needs to be performed more accurately for LTE than for UMTS.

Another source of timing error is a sampling frequency error. This can occur when the baseband sampling rate of the receiver is offset from that of the transmitter, or if there is jitter of the receiver's clock. When the sampling rate is too high this contracts the spectrum of the signal; similarly, slow sampling widens the spectrum. Either of these effects misaligns the FFT bin locations, resulting in a loss of orthogonality and hence reduced SINR. A sampling rate converter, perhaps driven by an error estimate, could correct for a systematic sampling rate error or sampling rate drift. Sampling jitter should be kept low by the choice of a clean crystal reference oscillator.

Group delay distortion

Filter group delay and amplitude ripple variation create deviation from the wanted impulse response of the receiver and cause inter-symbol interference. Unfortunately analogue filters cannot have both flat group delay and flat amplitude response, so a compromise is needed. To achieve this the analogue channel filter is usually slightly larger than the channel to allow reduced in-band distortion at the expense of out-of-band attenuation and in-band noise suppression.

One of the inherent advantages of multicarrier OFDM is that it is quite resilient to group delay variation and amplitude ripple, more so than the single-carrier QPSK modulation used by UMTS.

²⁹Phase-Locked Loop.

Additionally, the OFDM symbol rate is relatively low and thus there are greater margins for delays introduced by large filters. Overall there is more freedom with LTE to design the receiver analogue and digital filters to achieve high selectivity.

Reciprocal mixing

An unwanted phase offset of the received signal may include a large fixed (or slowly varying), phase offset, which can be corrected by the equalizer. However, the LO also introduces small random variations and jitter of the frequency and phase, which manifests itself as shaped phase noise, tending to reduce with offset from the carrier frequency. This cannot be corrected by the equalizer. Crystal oscillators have low levels of phase noise, but synthesizers are not so clean, especially if they are PLL-based.

Each of the elements of a frequency synthesizer produces noise which contributes to the overall noise appearing at the output. The noise within the PLL loop bandwidth arises from the phase detector and the reference, whereas outside the loop bandwidth the VCO³⁰ is the main noise source. The phase noise tends to reduce from the edge of the PLL loop bandwidth until eventually it reaches a flat noise floor.

A serious problem for receivers caused by phase noise is reciprocal mixing. Reciprocal mixing occurs when the phase noise on the receiver LO mixes with a strong interfering signal to produce a signal which falls inside the pass-band of the receiver. Intermediate-Frequency (IF) filters, if used, do not give any rejection of such signals; instead the problem must be solved by keeping the LO phase noise at a sufficiently low level. The various interference requirements (see Sections 21.4.6 to 21.4.7.2) together create an overall requirement for the receiver LO phase noise. The interferer which gets mixed onto the received signal needs to be weaker by a margin of the SINR requirement plus the implementation margin, plus a further margin of about 10 dB if it is to have no impact on the received signal. Therefore, the LO noise, L , needs to satisfy the following requirement:

$$L \leq \text{SINR} + \text{IM} - C/I + 10 + 10 \cdot \log_{10}(B) \text{ (dBc/Hz)}$$

21.5.2.3 Non-Linear Distortion and Intermodulation Products

IMD results from non-linearities of the PAs. The high PAPR (see Section 5.2.2) associated with multicarrier signals is one of the principal challenges in the implementation of OFDM systems. It requires linear operation of the PA over a large dynamic range, and this imposes a considerable implementation cost and reduced efficiency. The linearity and efficiency of a power amplifier are mutually exclusive specifications.

As already discussed, practical power amplifiers have a non-linear response. Numerous models exist, a selection of which can be found in [20]. Here we describe a simple polynomial memoryless model, by which the output signal can be written as

$$y_{\text{PA}}(t) = a_1 y(t) + a_2 y^2(t) + a_3 y^3(t) \quad (21.17)$$

³⁰Voltage-Controlled Oscillator.

where a_1 , a_2 , a_3 are independent coefficients which can be found by measurement. From Equation (21.8), (21.17) can be written as³¹

$$\begin{aligned}
 y_{\text{PA}}(t) = & a_1 \sum_{k=-N/2}^{N/2-1} S_{k,\ell} \exp\left[j2\pi k\Delta f \frac{t}{N}\right] \\
 & + a_2 \sum_{k=-N/2}^{N/2-1} \sum_{p=-N/2}^{N/2-1} S_{k,\ell} S_{p,\ell} \exp\left[j2\pi(k+p)\Delta f \frac{t}{N}\right] \\
 & + a_3 \sum_{k=-N/2}^{N/2-1} \sum_{p=-N/2}^{N/2-1} \sum_{v=-N/2}^{N/2-1} S_{k,\ell} S_{p,\ell} S_{v,\ell} \exp\left[j2\pi(k+p+v)\Delta f \frac{t}{N}\right] \quad (21.18)
 \end{aligned}$$

The non-linear response of the PA creates ICI. The intermodulation products contribute to a noise-like cloud surrounding each constellation point.³² For higher-order modulation in particular (such as 64QAM), these constellation clouds contribute to an increase in error rate for each subcarrier. Thus, in an OFDM modem design, linearity must be carefully controlled.

21.6 Summary

In this chapter, we have reviewed the RF requirements affecting the practical implementation of LTE equipment, especially the UE. We have seen how the effects of typical RF impairments of the transmitter and receiver can be analysed, including the use of a mathematical model for quantitative analysis.

The LTE RF requirements can be compared to those of UMTS. In many aspects, the RF requirements of LTE are similarly demanding as for UMTS, including for example narrowband blocking of the receiver and self-interference in case of full-duplex FDD operation. However, the LTE specifications are more challenging for other aspects, such as

- The channel bandwidth is variable over a wide range;
- At least two receive antennas are expected for receive diversity, and only the lowest UE category does not have to support reception of MIMO spatial multiplexing;
- There is a large variety of modulation and coding schemes, with some of them requiring a high SNR;
- OFDM is more sensitive to phase noise.

On the other hand, LTE is more robust than UMTS against amplitude and phase distortions from receiver and transmitter filters.

³¹This can be derived by considering ideal (de-)modulator components; thus, the signal in Equation 21.9 can be recovered by a simple low-pass filtering.

³²Note that the interference is correlated with the symbol transmitted on subcarrier k , and hence strictly it cannot be considered as white noise.

References³³

- [1] 3GPP Technical Specification 25.101, ‘User Equipment (UE) Radio Transmission and Reception (FDD)’, www.3gpp.org.
- [2] 3GPP Technical Specification 25.104, ‘Base Station (BS) Radio Transmission and Reception (FDD)’, www.3gpp.org.
- [3] 3GPP Technical Specification 25.102, ‘User Equipment (UE) Radio Transmission and Reception (TDD)’, www.3gpp.org.
- [4] 3GPP Technical Specification 25.105, ‘Base Station (BS) Radio Transmission and Reception (TDD)’, www.3gpp.org.
- [5] 3GPP Technical Specification 36.101, ‘Evolved Universal Terrestrial Radio Access (E-UTRA); User Equipment (UE) Radio Transmission and Reception’, www.3gpp.org.
- [6] 3GPP Technical Specification 36.104, ‘Evolved Universal Terrestrial Radio Access (E-UTRA); Base Station (BS) Radio Transmission and Reception’, www.3gpp.org.
- [7] 3GPP Technical Report 36.803, ‘Evolved Universal Terrestrial Radio Access (E-UTRA); User Equipment (UE) Radio Transmission and Reception’, www.3gpp.org.
- [8] 3GPP Technical Report 36.804, ‘Evolved Universal Terrestrial Radio Access (E-UTRA); Base Station (BS) Radio Transmission and Reception’, www.3gpp.org.
- [9] Motorola, ‘R1-060144: UE Power Management for E-UTRA’, www.3gpp.org, 3GPP TSG RAN WG1, LTE Ad-hoc, Helsinki, Finland, January 2006.
- [10] Motorola, ‘R1-060023: Cubic Metric in 3GPP-LTE’, www.3gpp.org, 3GPP TSG RAN WG1, LTE Ad-hoc, Helsinki, Finland, January 2006.
- [11] B. E. Priyanto, T. B. Sorensen, O. K. Jensen, T. Larsen, T. Kolding and P. Mogensen, ‘Assessing and Modeling the Effect of RF Impairments on UTRA LTE Uplink Performance’, in *Proc. IEEE Vehicular Technology Conference*, Baltimore, MD, USA, September 2007.
- [12] M. Moonen and D. Tander, ‘Compensation of RF Impairments in MIMO OFDM Systems’ in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing*, Las Vegas, NV, USA, April 2008.
- [13] 3GPP Technical Report 25.814, ‘Physical Layer Aspects for Evolved UTRA, (Release 7)’, www.3gpp.org.
- [14] J. Bussgang, ‘Crosscorrelation Function of Amplitude Distorted Gaussian Signals’, Technical Report 216, Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA, March 1952.
- [15] Agilent Technologies, ‘R4-061276: Effects of Physical Impairments on OFDM Signals’, www.3gpp.org, 3GPP TSG RAN WG4, meeting 33, Riga, Latvia, November 2006.
- [16] F. Munier, T. Eriksson and A. Svensson, ‘Receiver Algorithms for OFDM Systems in Phase Noise and AWGN’ in *Proc. IEEE International Symposium on Personal, Indoor and Mobile Radio Communications*, Barcelona, Spain, September 2004.
- [17] S. Wu and Y. Bar-Ness, ‘Performance Analysis on the Effect of Phase Noise in OFDM Systems’, in *Proc. IEEE International Symposium on Spread Spectrum Techniques and Applications*, Prague, Czech Republic, 2002.
- [18] T. C. W. Schenk, T. Xiao-Jiao, P. F. M. Smulders and E. R. Fledderus, ‘On the Influence of Phase Noise Induced ICI in MIMO OFDM Systems’. *IEEE Communications Letters*, August 2005.
- [19] E. Costa and S. Pupolin, ‘M-QAM-OFDM System Performance in the Presence of a Nonlinear Amplifier and Phase Noise’. *IEEE Trans. on Communications*, pp. 462–472, March 2002.

³³All web sites confirmed 1st March 2011.

- [20] S. C. Cripps, *RF Power Amplifiers for Wireless Communications*. Artech House, Norwood, MA, USA, 1999.

Radio Resource Management

Muhammad Kazmi

22.1 Introduction

Radio Resource Management (RRM) encompasses a wide range of techniques and procedures, including power control, scheduling, cell search, cell reselection, handover, radio link or connection monitoring, and connection establishment and re-establishment. Advanced features like interference management, location services, Self-Optimizing Networks¹ (SON) and some network planning methods make use of RRM-related techniques based on radio related measurements made by the User Equipment (UE) or eNodeB. In this chapter, we address the RRM techniques and reporting mechanisms that support UE mobility in the LTE network (E-UTRAN), including cell search, radio measurements, cell reselection, handover and radio link monitoring. We focus here on the performance requirement aspects, while the procedures themselves are described in Chapters 3, 4, 7, 12, 17 and 18.

The RRM-related actions undertaken by the UE can be broadly divided into those relevant in the RRC_IDLE state and those relevant in the RRC_CONNECTED state, as illustrated in Figures 22.1 and 22.2 respectively.

RRM in E-UTRAN is designed to handle the challenges posed by the fundamental characteristics of the LTE system, including:

- The packet-oriented transmission of LTE, realized by fast time- and frequency-domain scheduling, may lead to large and swift interference fluctuations. This may affect the accuracy of the signal quality estimates required for mobility decisions in certain scenarios. Time-domain filtering is used to help smooth out the interference variations.
- The provision of a wide range of Discontinuous Reception (DRX) cycle lengths, including periods up to 2.5 seconds, yields substantial benefits in terms of UE power saving. However, such DRX cycles also mean that it is not feasible to enforce the same mobility performance in all cases.

¹Details on user equipment positioning and self optimization of the network are described in Chapters 19 and 25 respectively.




Figure 22.1: UE actions related to RRM Procedures in RRC_IDLE state.

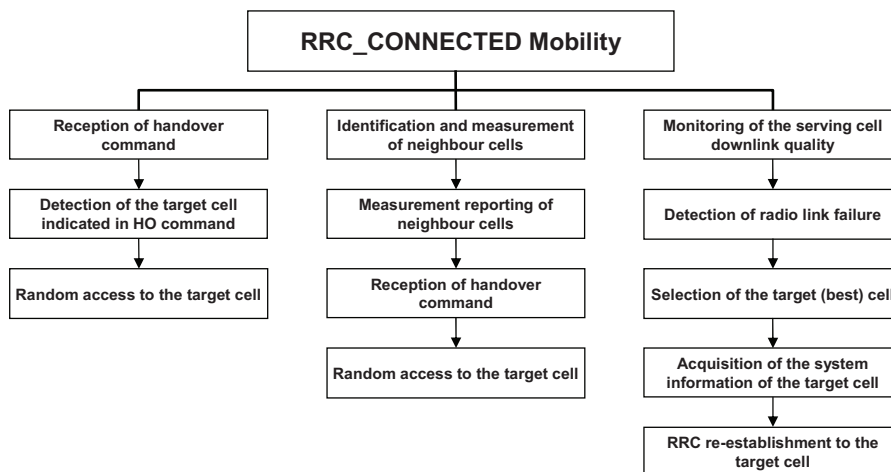


Figure 22.2: UE actions related to RRM Procedures in RRC_CONNECTED state.

- LTE is designed to support seamless mobility not only within itself (both intra- and inter-frequency) but also with the legacy 3GPP Radio Access Technologies (RATs) (e.g. GERAN² and UTRAN³) and with certain non-3GPP RATs (e.g. CDMA2000

²GSM EDGE Radio Access Network.

³Universal Terrestrial Radio Access Network.

1xRTT and HRPD⁴). Mobility between these different RATs requires the UE to detect and measure cells on the target technologies as requested by the network, each with its own different channel structure.

- LTE is designed to support a wide range of cell sizes (ranging from a few tens of metres to tens of kilometres) and deployment/propagation scenarios, yet the RRM techniques are required to be as generic as possible and to exhibit consistent mobility performance.
- Despite the differences between the Frequency Division Duplex (FDD) and Time Division Duplex (TDD) frame structures in LTE (see Section 6.2), and the fact that in FDD deployments the eNodeBs may be either synchronized or unsynchronized, mobility requirements are expected to be the same (or at least similar) in order to minimize UE implementation complexity and simplify network planning.
- The low overall latency requirements in LTE place constraints on the time taken to make and report measurements and to perform handover. Additionally, to meet the challenges of certain specific mobility procedures, enhanced requirements with substantially lower latency are also specified.
- The time-multiplexing of Multimedia Broadcast/Multicast Service (MBMS) and unicast data on a single LTE carrier is attractive in that it obviates the need for a separate carrier frequency to offer mobile broadcast services. However, it reduces the opportunities for the UEs to perform downlink cell measurements to support mobility; nevertheless, the RRM performance should remain consistent.

22.2 Cell Search Performance

Cell search is one of the most fundamental aspects of mobility. As explained in Chapter 7, it enables the UE to acquire the carrier frequency, timing and cell identity of cells. In LTE the cell search performance requirements relate to the case of neighbour cell search under the assumption that the UE has already acquired carrier frequency synchronization; the requirements are only applicable in RRC_CONNECTED.

22.2.1 Cell Search within E-UTRAN

The objective of cell search within E-UTRAN is to identify one of the 504 unique Physical Cell Identities (PCIs) (see Chapter 7 and [1, Section 6.11]). The cell search algorithm is not specified and is left for UE implementation; however, typically, the UE performs cell search in a hierarchical manner (see Section 7.2).

An important characteristic of the cell identification requirements is that the same requirements are applicable in a wide range of propagation conditions and for both FDD (with or without synchronization of the eNodeBs) and TDD (where synchronization of the eNodeBs can be assumed).

The requirements are specified in terms of the maximum permissible cell identification delay, which includes the time taken for Reference Signal Received Power (RSRP) or Reference Signal Received Quality (RSRQ) physical layer measurements (see Sections 22.3.1.1 and 22.3.1.2 respectively).

⁴High Rate Packet Data.

22.2.1.1 E-UTRAN Intra-frequency Cell Search

In case of intra-frequency cell search, the UE identifies E-UTRA cells on the same carrier frequency as that of the serving cell. The time required to detect a cell depends upon a number of factors, most notably the received quality of the synchronization signals, the received level of the Reference Signals (RSs) and the time available for performing the search. The latter factor stems from the fact that the available time for intra-frequency measurements may be reduced by measurement gaps for inter-frequency or inter-RAT measurements as explained in Section 22.2.1.2. The cell search delay also depends upon the configured DRX cycle period.

If no DRX is configured, and for DRX cycles up to 40 ms,⁵ the UE is required to detect an E-UTRA FDD or TDD intra-frequency target cell within 800 ms if no inter-frequency measurement gaps are configured, provided that the target cell's received synchronization signal quality \hat{E}_s/I_{ot} (defined as the energy per Resource Element (RE) of the synchronization signals divided by the total received energy of noise and interference on the same RE) is at least -6 dB. This is the 'minimum', or worst case, requirement.

The cell search delay can be shorter if the received signal quality is higher than the minimum cell detection threshold. The performance in some typical deployment conditions is illustrated in Figures 22.3 and 22.4. Here, scenarios covering both synchronized and unsynchronized eNodeBs are analysed, as summarized in Table 22.1. ETU5 (Extended Typical Urban with UE speed 5 km/h), ETU300 (UE speed 300 km/h) and EPA5 (Extended Pedestrian A with UE speed 5 km/h) propagation models are used,⁶ and two receive antennas are assumed at the UE.⁷ Further details of the modelled scenarios can be found in [2].

Table 22.1: Cell identification test parameters.

	Unit	Cell1	Cell2	Cell3 (target cell)
Relative delay of 1 st path for synchronized case	ms	0	0	Half CP length
Relative delay of 1 st path for unsynchronized case	ms	0	1.5	3
SNR	dB	5.18	0.29	-0.75 (worst case)
PSS for case of different PSS		PSS1	PSS2	PSS3
PSS for case of same PSS		PSS1	PSS2	PSS1

The cell search performance is measured in terms of the 90-percentile cell identification delay, i.e. the maximum time required to detect the target cell 90% of the time.

Various scenarios are analysed to examine the impact on the detection performance of different combinations of PSS⁸ and SSS⁹ sequences as indicated in Table 22.2. More detailed performance results can be found in [3].

⁵This is designed to ensure robust mobility performance for delay-sensitive services like Voice over IP (VoIP), which typically requires short DRX cycles.

⁶Further details of these propagation models are given in Chapter 20.

⁷No margin is included for non-ideal UE receiver implementation or reporting delay for the RSRP measurement to the network.

⁸Primary Synchronization Signal.

⁹Secondary Synchronization Signal.

Table 22.2: Cell identification test scenarios.

Test case (synch, asynch eNodeBs)	Cell3 (Target)	Cell1 (Interference)	Cell2 (Interference)
1,5	PSS3 SSS3a, SSS3b	PSS1 SSS1a, SSS1b	PSS2 SSS2a, SSS2b
2,6	PSS1 SSS3a, SSS3b	PSS1 SSS1a, SSS1b	PSS2 SSS2a, SSS2b
3,7	PSS1 SSS1a, SSS3b	PSS1 SSS1a, SSS1b	PSS2 SSS2a, SSS2b
4,8	PSS3 SSS1a, SSS1b	PSS1 SSS1a, SSS1b	PSS2 SSS2a, SSS2b




Figure 22.3: Cell search performance with synchronized eNodeBs.
 Reproduced by permission of © NXP Semiconductors.

When measurement gaps are configured for inter-frequency or inter-RAT measurements the UE has less opportunity to detect a cell. Hence, in this case the cell identification delay may be larger than the baseline 800 ms delay, with the actual value depending upon the periodicity of the gaps. Furthermore, for DRX cycles larger than 40 ms, the cell identification delay increases in proportion to the length of the DRX cycle, allowing UE to save battery power.

22.2.1.2 E-UTRAN Inter-frequency Cell Search

In the case of inter-frequency cell search, the UE identifies E-UTRA cells operating on carrier frequencies other than that of the serving cell (and possibly also in different frequency bands and/or with different duplex modes). Inter-frequency measurements, including cell identification, are performed during periodic measurement gaps unless the UE has more than one receiver. Two possible gap patterns can be configured by the network, each with a gap length of 6 ms: in gap pattern #0, the gap occurs every 40 ms, while in gap pattern #1 the




Figure 22.4: Cell search performance with unsynchronized eNodeBs. Reproduced by permission of © NXP Semiconductors.

gap occurs every 80 ms, as shown in Figure 22.5. There is an obvious trade-off between these two gap patterns: the former yields a shorter cell identification delay but a greater interruption in data transmission and reception. Only one gap pattern can be configured at a time for measuring all frequency layers (both inter-frequency and inter-RAT).




Figure 22.5: Measurement gap patterns for inter-frequency and inter-RAT cell search and measurements.

If no DRX is used, or if the DRX cycle length is less than or equal to 160 ms, the UE is required to identify an E-UTRA FDD or TDD inter-frequency cell within 3.84 s, provided the received synchronization signal quality is at least -4 dB (assuming gap pattern #0). As for intra-frequency cell search, for DRX cycles larger than 160 ms the cell identification delay increases proportionally.

22.2.2 E-UTRAN to E-UTRAN Cell Global Identifier Reporting Requirements

In addition to the basic PCI cell search requirements specified in LTE Release 8, E-UTRAN Cell Global Identifier¹⁰ (ECGI) reporting requirements were introduced in Release 9 for both FDD and TDD. The requirements are the same for FDD and TDD for their respective frequency bands, and therefore the description in the subsequent sections is generic for both.

The UE is required to identify and report the ECGI of a target E-UTRA cell upon request from the serving eNodeB. The ECGI measurement report can be used by the eNodeB for various features. For instance, it can assist the eNodeB to establish the neighbour cell relations automatically (see Chapter 25.2) when a new neighbour cell is added or an existing one is removed [5]. The ECGI measurement report can also be used by the serving eNodeB for performing inbound mobility to a Home eNodeB¹¹ (HeNB) in RRC_CONNECTED state (see Chapter 24 for more details on HeNBs). Due to the small size of HeNB cells, it is not unlikely that a PCI may be used more than once among the HeNBs located within the coverage area of a macro-eNodeB. Hence an important objective of the ECGI report is to enable the serving eNodeB to identify unambiguously the target HeNB.

The acquisition of ECGI requires the UE first to read the Master Information Block (MIB), which is transmitted on Physical Broadcast Channel (PBCH) with a periodicity of 40 ms (see Section 9.2.1). This is followed by the reading of the System Information Block type 1 (SIB1), which contains the ECGI and is transmitted with a periodicity of 80 ms on the DownLink Shared Channel (DL-SCH) (see Section 3.2.2). The UE is not required to receive or transmit in the serving cell while acquiring the ECGI of a target cell; a UE is permitted to create gaps autonomously in downlink reception and uplink transmission to read the MIB and SIB1 of a target cell.

HeNBs can be deployed on either a shared carrier or a dedicated carrier. Similarly the SON Automatic Neighbouring Relation (ANR) is applicable to both intra- and inter-frequency cells, and therefore ECGI requirements are specified for both intra-frequency and inter-frequency ECGI reporting.

22.2.2.1 Intra-frequency E-UTRAN Cell Global Identifier Reporting

The intra-frequency ECGI reporting requirement is specified in terms of the maximum time allowed for the UE to identify the ECGI of a new intra-frequency target E-UTRA cell whose synchronization signal quality \hat{E}_s/I_{ot} (as defined in Section 22.2.1.1) is at least -6 dB; the identification delay is not allowed to exceed 150 ms with 90% confidence, excluding the procedure delay of the ECGI request message. The same ECGI identification delay requirements are applicable with and without DRX.

The initial autonomous gap created by the UE for correcting the frequency and for decoding the first MIB block of the target cell comprises 9 subframes. Each subsequent autonomous gap for decoding the remaining MIB and SIB1 blocks to read the whole of the target cell's ECGI typically consists of 4 subframes. The maximum aggregated duration of autonomous gaps created by the UE is not allowed to exceed 90 subframes when identifying

¹⁰The E-UTRAN Cell Global Identifier is a globally unique identifier broadcast by the cell. It is composed of 3-bytes of Public Land Mobile Network (PLMN) Identity and 28 bits to identify the cell within that PLMN – see [4, Section 6.3.4], and Chapter 25.

¹¹Inbound mobility to Home eNodeB refers to the case of a UE moving into a Home eNodeB cell.

an intra-frequency cell's ECGL; this is designed to ensure that certain minimum serving cell reception and transmission performance is maintained.

22.2.2.2 Inter-frequency E-UTRAN Cell Global Identifier Reporting

The inter-frequency ECGL requirements (i.e. ECGL identification delay and serving cell reception/transmission performance) are the same as those specified for the intra-frequency requirement, except that the synchronization signal quality under which the requirement should be satisfied is relaxed to -4 dB. An autonomous gap created by the UE for decoding the initial MIB block of the target inter-frequency cell comprises 9 subframes, and each of the subsequent gaps typically consists of 4 subframes. The maximum aggregated duration of autonomous gaps created by the UE is likewise not allowed to exceed 90 subframes when identifying an inter-frequency cell's ECGL. This ensures that certain minimum serving cell reception and transmission performance is maintained while reading the inter-frequency target cell ECGL.

22.2.3 E-UTRAN to UTRAN Cell Search

The E-UTRAN to UTRAN cell search procedure allows the UE to identify a target FDD or TDD UTRA cell. The E-UTRAN-UTRAN cell search requirements apply until an explicit neighbour cell list is received by the UE from the serving E-UTRA cell. This list indicates the primary scrambling codes of up to 32 neighbour cells per UTRA FDD or TDD carrier. The signals provided by UTRA FDD and TDD differ, and therefore separate cell identification requirements are specified for each technology. In each case, the cell search algorithm is not specified, but the search is typically carried out in hierarchical manner. In UTRA FDD, where

Primary and Secondary Synchronization Channels (P-SCH and S-SCH) are provided, the UE typically identifies the target UTRA cell via the following steps:

Slot timing acquisition: The UE acquires the slot timing by performing a matched filter correlation over the Primary Synchronization Code (PSC), which is transmitted in the first 256 chips in every slot (0.666 ms). The PSC is common to all UTRA cells.

Frame timing acquisition and code group identification: The UE acquires the frame timing and the identity of the group of 8 codes to which the cell's primary scrambling code belongs by performing correlations over the Secondary Synchronization Codes (SSCs) transmitted in the initial 256 chips of every slot (i.e. at the same time as the PSC). A sequence of 15 SSCs in one 10 ms radio frame uniquely identifies the code group. In good propagation conditions the information contained within three slots is sufficient to identify both the frame timing and the code group.

Primary scrambling code identification: A given UTRA cell uses one code from the group indicated by the SSC as the scrambling code for all downlink channels, including the Primary Common Pilot Channel (P-CPICH). The UE therefore performs a simple correlation against the known CPICH sequence (which is the same in all UMTS cells) scrambled in turn by each of the eight possible scrambling sequences in the code group, in order to determine which primary scrambling code is being used in the cell, and hence the cell identity.

In order to report an identified cell, the UE needs to measure the P-CPICH E_c/N_0 (where E_c is the CPICH energy per chip and N_0 the noise power spectral density) or Received Signal Code Power (RSCP) over the 'physical layer measurement period'.

The same measurement gap patterns are used for E-UTRA inter-frequency and inter-RAT measurements. In both DRX and non-DRX scenarios, a target UTRA FDD cell is considered detectable if the CPICH E_c/I_0 is at least -20dB and the synchronization channel SCH_E_c/I_0 is at least -17dB, where I_0 is the power spectral density of the total input signal at the UE antenna connector and SCH_E_c is the transmit energy per chip of the SCH, including the downlink signal transmitted by the serving cell [6]. However, in terms of reporting delay, separate sets of UTRA cell identification requirements are specified for the cases with and without DRX.

If no DRX is used, or for DRX cycles up to 40 ms, the UE is required to identify a detectable UTRA FDD cell within $2.4 * N_{\text{freq}}$ seconds and $4.8 * N_{\text{freq}}$ seconds using gap patterns #0 and #1 respectively (where N_{freq} is the number of carrier frequencies to be monitored). For DRX cycles larger than 40 ms, the UTRA cell search delay increases proportionally and is expressed in terms of a number of DRX cycles. For example in case of DRX cycle length of 1.28 seconds, the maximum permissible cell search delay is $25.6 * N_{\text{freq}}$ seconds or $20 * \text{DRX cycle length} * N_{\text{freq}}$; where the factor of 20 incorporates the required UTRA FDD cell search occasions and CPICH measurement samples.

22.2.4 E-UTRAN to GSM Cell Search

In GSM the Broadcast Control CHannel (BCCH) frequency reuse factor is greater than one. Therefore in principle the GSM BCCH Received Signal Strength Indicator (RSSI) measurement, which is measured on the BCCH carrier (see Section 22.3.3) can identify a GSM cell, since a different RSSI measurement can be reported for each carrier. However, if small cluster sizes are used with BCCH frequency reuse, the GSM carrier RSSI alone cannot always reliably identify the correct GSM cell: for example, the network may not be able to distinguish RSSI measurement reports for GSM cells whose BCCHs use the same carrier. This may lead to calls being dropped, especially in deployment scenarios with tight reuse. Therefore, the UE is also required to decode the Base Station Identity Code (BSIC)¹². In E-UTRAN, GSM measurements are always reported together with the verified BSIC [4], in contrast to the legacy GSM system, where measurements can be reported with or without a verified BSIC. The BSIC verification comprises initial BSIC identification and BSIC re-confirmation. These are summarized below; A more extensive description of the GSM common channels is provided in [7].

Initial BSIC identification comprises the BSIC decoding for a GSM cell for the first time when the UE does not have any knowledge about the relative timing between the E-UTRA and GSM cells. The UE is required to decode the BSIC of the 8 strongest BCCH carriers

¹²The BSIC allows a UE to distinguish two different GSM cells which share the same beacon frequency. It is a 6-bit field composed of two 3-bit fields: the Base station Colour Code (BCC) and the Network Colour Code (NCC). The BCC is used to identify the Training Sequence Code (TSC) to be used when reading the BCCH. The NCC is used to differentiate between operators utilizing the same frequencies, e.g. on an international border when both operators have been allocated the same frequency or frequencies.

of the GSM cells that are provided in the GSM inter-RAT cell information list¹³ [4]. The UE decodes the BSIC of the BCCH carriers in order of decreasing RSSI. Through the BSIC decoding the UE acquires the initial timing information of these GSM cells.

BSIC re-confirmation consists of the BSIC decoding of a GSM cell after the initial BSIC identification has been performed. The UE updates its stored values of the timing of up to 8 identified GSM cells every time the BSIC is decoded, in order to compensate for cell timing drift relative to the serving cell to which the UE is locked.

BSIC verification requirements are applicable down to the reference sensitivity level defined in [8]. Initial BSIC identification and BSIC re-confirmation requirements are expressed in terms of the time required to decode the BSIC for the first time, $T_{\text{identity,GSM}}$ and the time required to re-confirm the initially identified BSIC $T_{\text{re-confirm,GSM}}$ under the above constraints, which depend on the total number of carrier frequency layers of all RATs since the same gap pattern is shared by all RATs.

22.2.5 Enhanced Inter-RAT Measurement Requirements

As mentioned earlier, the measurement requirements are generally minimum requirements for worst-case scenarios. In order to support Circuit-Switched Fall Back (CSFB)¹⁴ scenarios in E-UTRAN, which require shorter overall delay, enhanced measurement requirements for UTRAN FDD and GSM were introduced in LTE Release 9.

For both UTRA FDD and GSM, the enhanced measurement requirements are specified under the assumption that CSFB from a serving E-UTRA cell to a target UTRA FDD or GSM cell is used when the serving cell and the target cell have overlapping coverage. Under these conditions the received signal quality of the target cell is typically higher than the minimum levels specified for normal cell search.

In the case of UTRA FDD, the enhanced UTRA FDD cell identification requirements are therefore specified assuming that both received CPICH E_c/I_0 and synchronization channel SCH E_c/I_0 levels from the UTRA FDD cell are at least -15 dB. At these levels when no DRX is used or when $\text{DRX} \leq 40$ ms, the UTRA FDD cell search delay is about 1 s (assuming gap pattern #0). Hence, compared to the minimum requirements, the enhanced UTRA FDD measurement reporting delay of an unknown UTRA FDD cell is reduced by a factor of 2.

In the case of GSM, the enhanced BSIC verification requirements are specified assuming that the target cell is received at about 10 dB above the reference sensitivity level or reference interference levels (as specified in [8]). Under this condition and without DRX or with $\text{DRX} \leq 40$ ms, the UE can complete initial BSIC decoding in about 1.3 s, assuming gap pattern #0 is used only for monitoring GSM measurements.

¹³When camped on an LTE cell, the UE is provided with a Neighbour Cell List (NCL) containing at least 32 GSM carrier numbers (i.e. Absolute Radio Frequency Channel Numbers (ARFCNs)) indicating the frequencies of neighbouring cells, and optionally an associated BSIC for each GSM carrier in the NCL.

¹⁴CSFB enables voice and other circuit-switched services to be provided to UEs served by E-UTRAN by reusing the existing circuit switched infrastructure such as GERAN or UTRAN [9] – see Section 2.4.2.1.

22.3 Mobility Measurements

In order to support mobility within E-UTRAN and between E-UTRAN and other RATs (UTRAN FDD and TDD, GERAN and CDMA2000), a number of radio-related UE measurements are specified. All the E-UTRAN and inter-RAT measurements described in the following sections are reported by the UE to the serving eNodeB only in RRC_CONNECTED state. In RRC_IDLE, the measurements are not reported but may be used autonomously by the UE for cell reselection. Therefore although a particular measurement may be used in both RRC_CONNECTED and RRC_IDLE states, the corresponding requirements only apply to RRC_CONNECTED.

22.3.1 E-UTRAN Measurements

For mobility within E-UTRAN and from other RATs to E-UTRAN, two UE measurements are defined: RSRP and RSRQ [10].

22.3.1.1 Reference Signal Received Power (RSRP)

RSRP is measured by the UE over the cell-specific Reference Signals (RSs) within the measurement bandwidth over a measurement period. RSRP is a type of signal strength measurement and is indicative of the cell coverage. It is defined as the linear average over the power contributions (in Watts) of the REs that carry cell-specific RSs within the considered measurement frequency bandwidth. Normally the RSs transmitted on the first antenna port are used for RSRP determination, but the RSs on the second antenna port can also be used if the UE can reliably determine that they are being transmitted. If receive diversity is in use by the UE, the combined RSRP must be at least as large as the RSRP of any of the individual diversity branches. It is applicable in both RRC_IDLE and RRC_CONNECTED states and is used for cell reselection and handover within E-UTRAN (intra-frequency and inter-frequency) and to E-UTRAN from any of UTRAN FDD or TDD, GSM, CDMA2000 1xRTT or HRPD.

RSRP is therefore considered to be the most important measurement quantity for E-UTRAN.

An accuracy requirement is defined for the RSRP measurement (both intra- and inter-frequency); this is applicable for received signal quality \hat{E}_s/I_{ot} as low as -6 dB. A measurement bandwidth equivalent to the central 6 Resource Blocks (RBs) is assumed. In the time domain the physical layer measurement periods when no DRX is used (and for short DRX cycles) are 200 ms and 480 ms for intra-frequency and inter-frequency RSRP respectively. The inter-frequency measurement naturally takes longer as it can only be performed during the measurement gaps; furthermore, the inter-frequency RSRP physical layer measurement period increases linearly with the number of carrier frequencies that has to be measured during the gaps, N_{freq} . The physical layer measurement period also increases in proportion to the DRX cycle for DRX cycles larger than 40 ms for intra-frequency RSRP and larger than 80 ms for inter-frequency RSRP. The measurement sampling rate is not specified but is left to the UE implementation.

The UE is required to be able to measure RSRP from at least 8 identified intra-frequency cells over the physical layer measurement period.¹⁵ Similarly, the UE is also required to measure RSRP from at least 4 identified inter-frequency cells per inter-frequency carrier for up to 3 carriers (i.e. a total of 12 inter-frequency cells).

22.3.1.2 Reference Signal Received Quality (RSRQ)

RSRQ is the ratio of RSRP to RSSI for an E-UTRA carrier. The RSSI part of RSRQ is the total received power including interference from all sources, including serving and non-serving cells, adjacent channel interference and thermal noise. Unlike RSRP, it is measured on all REs in the OFDM symbols containing RSs for antenna port 0 (not just the REs containing RSs themselves), within the measurement bandwidth.¹⁶

This interference component of RSRQ enables the UE to quantify the received signal quality considering both signal strength and interference, which may vary with the UE's location in the cell [11].

In the first release of LTE (Release 8), RSRQ was applicable only in RRC_CONNECTED state. It is therefore used for handover within E-UTRAN, and from other RATs to E-UTRAN.

However, in order to prevent outages caused by high interference situations, in Release 9 RSRQ was also introduced for RRC_IDLE state; this gives the network the option to configure the UE to use RSRQ as a metric for performing cell reselection, at least in the cases of cell reselection within E-UTRAN, from UTRAN FDD to E-UTRAN and from GSM to E-UTRAN.

Both intra- and inter-frequency RSRQ measurement requirements are specified in [12]. The RSRQ and RSRP together have been shown to be particularly beneficial for performing inter-frequency quality-based handover [13]. RSRQ is inherently a relative quantity which to some extent eliminates absolute measurement errors and leads to better accuracy than is possible for RSRP. Like RSRP, the RSRQ accuracy requirements are applicable for \hat{E}_s/I_{ot} down to -6 dB (based on a measurement bandwidth equivalent to the central 6 RBs). The time domain physical layer measurement periods for RSRP are equivalent to those of RSRP for the corresponding DRX cases. The measurement sampling rate is also UE implementation dependent.

The UE is also required to measure RSRQ from the same number of intra- and inter-frequency cells as for RSRP.

22.3.2 UTRAN Measurements

The UTRAN measurements can be categorized into those specific to UTRAN FDD and those for E-UTRAN TDD. For mobility to UTRAN FDD, CPICH¹⁷RSRP, CPICH E_c/N_0 ¹⁸ and RSSI measurements are defined [10]. The CPICH measurements are the most important for

¹⁵When measurement gaps are used, and depending upon the gap periodicity, the number of measured cells may be lower than 8 due to the fact that the measurement gaps reduce the time available for the UE to make intra-frequency measurements.

¹⁶Note that in Release 10 the RSRQ can be configured to be measured on all OFDM symbols if mechanisms for time-domain inter-cell interference coordination are configured – see Section 31.2.4.2.

¹⁷Common Pilot Channel.

¹⁸Energy per chip divided by the noise power spectral density.

mobility between E-UTRAN and UTRAN FDD systems. For mobility to UTRAN TDD P-CCPCH¹⁹ RSCP and RSSI measurements are used.

All these UTRAN measurements can be used in both RRC_IDLE and RRC_CONNECTED states. However, the network has the freedom to configure specific measurements in different mobility scenarios. Therefore, all the measurement requirements (such as accuracies) are applicable only in RRC_CONNECTED state.

22.3.2.1 UTRAN FDD CPICH RSCP

UTRAN FDD CPICH RSCP is measured on the Primary-CPICH of the target UTRAN cell and is used for both cell reselection and handover from E-UTRAN. For inter-RAT mobility, the CPICH RSCP absolute accuracy requirement should be fulfilled over the physical layer measurement period, whose length depends upon the gap pattern and DRX cycle (e.g. 480 ms for gap pattern #0 with DRX cycle not greater than 40 ms). The UE is required to measure the CPICHs from 6 identified UTRA FDD cells.

22.3.2.2 UTRAN FDD CPICH E_c/N_0

The UTRAN FDD CPICH E_c/N_0 is the ratio of the CPICH RSCP to the UTRA carrier RSSI. CPICH E_c/N_0 can be used for both cell reselection and handover.

The measurement periods for CPICH E_c/N_0 with and without DRX are the same as for CPICH RSCP.

22.3.2.3 UTRAN FDD Carrier RSSI

The UTRAN FDD carrier RSSI is the total received wideband power, including thermal noise and noise generated in the receiver, within the bandwidth defined by the receiver pulse shaping filter. In practice this measurement quantity is rarely used in RRC_IDLE, and its usage is mainly for performing RRC_CONNECTED handovers between E-UTRAN and UTRAN FDD. The measurement periods are the same as those described in Section 22.3.2.1.

22.3.2.4 UTRAN TDD P-CCPCH RSCP

The UTRAN TDD P-CCPCH RSCP is the received power of the P-CCPCH of a UTRA TDD neighbour cell. It is considered to be the most fundamental measurement quantity for all mobility scenarios between E-UTRAN and UTRAN TDD. It is therefore used for performing both cell reselection and handover to UTRA TDD. The P-CCPCH RSCP physical layer measurement period depends on the configured gap periodicity and the DRX cycle in the same way as for UTRAN FDD.

22.3.2.5 UTRAN TDD carrier RSSI

The UTRAN TDD carrier RSSI is the total received wideband power, including thermal noise and noise generated in the receiver, within the bandwidth defined by the receiver pulse shaping filter within a specified timeslot. In principle it can be used for both cell reselection

¹⁹Primary Common Control Physical CHannel.

and handover from E-UTRAN to UTRAN TDD. The measurement periods of UTRAN TDD carrier RSSI with and without DRX are the same as for the P-CCPCH RSCP measurement.

22.3.3 GSM Measurements: GSM Carrier RSSI

For mobility from E-UTRAN to GSM in RRC_IDLE mode (i.e. cell reselection) and RRC_CONNECTED mode (i.e. handover) the GSM carrier RSSI is used.

The GSM carrier RSSI is measured on a GSM BCCH carrier [10]. In accordance with the GSM core specification [14] the UE is required to take at least 3 GSM carrier RSSI samples per GSM carrier as evenly spaced as possible during the physical layer measurement period. The GSM carrier RSSI physical layer measurement period without DRX and for DRX cycles up to 80 ms is 480 ms regardless of the gap pattern, and increases linearly with the number of carrier frequencies, N_{freq} , monitored during the measurement gaps. For DRX cycles larger than 80 ms, the measurement period increases with the length of the DRX cycle.

22.3.4 CDMA2000 Measurements

For mobility from E-UTRAN to CDMA2000 systems the CDMA2000 1xRTT pilot strength and CDMA2000 High Rate Packet Data (HRPD) Pilot Strength measurement quantities are used in RRC_IDLE and RRC_CONNECTED states.

22.3.4.1 CDMA2000 1xRTT Pilot Strength

CDMA2000 1xRTT Pilot Strength is the strength of the received pilot signals as defined in [15, Section 2.6.6.2.2]. The required measurement accuracy, specified in [16, Section 3.2.4], is to be fulfilled over the measurement periods of about $1.1 * N_{\text{freq}}$ seconds and $2.1 * N_{\text{freq}}$ seconds for gap patterns #0 and #1 respectively. It is used for cell reselection and handover from E-UTRAN to CDMA2000 1xRTT. A CDMA2000 1xRTT-capable UE is required to monitor CDMA2000 1xRTT Pilot Strength of CDMA2000 1xRTT cells for up to five CDMA2000 1xRTT carriers.

22.3.4.2 CDMA2000 HRPD Pilot Strength

CDMA2000 HRPD pilot strength is the strength of the received pilot signals as defined in [17, Section 8.7.6.1.2.3]. It is used for cell reselection and handover from E-UTRAN to CDMA2000 HRPD. An HRPD-capable UE is required to monitor CDMA2000 HRPD Pilot Strength of HRPD cells for up to five HRPD carriers.

22.4 UE Measurement Reporting Mechanisms and Requirements

As explained in Section 22.3, in LTE the UE reports measurements to the serving eNodeB only in RRC_CONNECTED state using the Dedicated Control CHannel (DCCH). The measurement reporting mechanism can be periodic, event-triggered or event-triggered and periodic.

These reporting mechanisms are applicable both with and without DRX. The interval between periodic reports is configurable by the eNodeB and ranges from 120 ms to 3600 s [4]. The events which can trigger reporting by the UE are described in Section 3.2.5.1.

The total number of measurement reporting criteria which can be evaluated by the UE in parallel is limited to 21. This includes 9 E-UTRA intra-frequency reporting criteria, 7 E-UTRA inter-frequency reporting criteria and 5 inter-RAT reporting criteria. The actual number and types of reporting criteria are configured by the network and can be based on any of the three possible reporting mechanisms mentioned above.

The event triggered reporting delay requirements, which also apply to the first report in event-triggered periodic reporting, are explained in the following sections for different E-UTRA and inter-RAT cases.

22.4.1 E-UTRAN Event Triggered Reporting Requirements

The E-UTRAN event-triggered reporting consists of intra- and inter-frequency event reporting. Five such events (A1-A5) and their corresponding triggering thresholds are configured in the UE by the serving eNodeB as explained in section 3.2.5.2. When an event is triggered, the UE reports the event to the eNodeB, which may take an appropriate mobility decision.

For both intra-frequency and inter-frequency event-triggered measurements, the reporting delay²⁰ is less than the cell search delay (see Sections 22.2.1.1 and 22.2.1.2) for the corresponding cases. Due to signal variations and user mobility, the measured quality of a cell may vary between detectable and undetectable threshold levels.

To cater for this typical scenario, requirements are specified for both intra- and inter-frequency event reporting as follows: Regardless of the number of times the cell quality varies between the detectable and undetectable levels, whenever an event (i.e. any of the intra- or inter-frequency events A1-A5) is triggered the UE is required to send the event-triggered measurement report within a duration less than the relevant physical layer measurement period without layer 3 filtering (e.g. within 200 ms RSRP/RSRQ for the intra-frequency non-DRX case), provided certain conditions are met. When layer 3 filtering is used, a reporting delay longer than the physical layer measurement period is expected. For an intra-frequency event reporting, these conditions require that the cell which has been detectable over the intra-frequency cell search delay, does not become undetectable for more than 5 s. It is also required that the timing to the intra-frequency cell does not change by more than $\pm 50T_s$.²¹ In case of inter-frequency event reporting the corresponding condition requires that the timing to the inter-frequency cell does not change by more than $\pm 50T_s$ during periods when the measurement gaps are not available (i.e. between successive measurement gaps). This is because in the absence of measurement gaps the UE cannot appropriately track the timing of an inter-frequency cell.

22.4.2 Inter-RAT Event-Triggered Reporting

Event-triggered reporting requirements are specified for E-UTRAN to UTRAN FDD, E-UTRAN to UTRAN TDD and E-UTRAN to GSM inter-RAT scenarios. Two inter-RAT

²⁰Here we refer to the delay for the transmitted measurement reports without layer 3 filtering being applied. The behaviour of the Layer 3 filters is standardized and their configuration provided by RRC signalling. See [4, Section 5.5.3.2] for more details.

²¹ T_s is the basic time unit and is equal to $1 / (15000 \times 2048)$ s.

events (B1 and B2) and their corresponding triggering thresholds are also configured in the UE by the serving eNodeB as explained in section Section 3.2.5.2. When one of these events is triggered, the UE reports it to the eNodeB, allowing it to take an appropriate inter-RAT mobility decision.

E-UTRAN to UTRAN event-triggered reporting. The UTRA FDD and TDD event-triggered measurement reporting delays without Layer 3 filtering are less than the corresponding cell search delays (see Section 22.2.3). However, if a UTRA FDD cell, which was previously detectable over a period equal to the UTRA FDD cell search delay, enters or leaves the reporting range, then the event-triggered measurement reporting delay is less than the UTRA FDD physical layer measurement period without layer 3 filtering provided that the timing to that UTRA FDD cell has not changed more than ± 32 chips when the measurement gap has not been available. Note also that the inter-RAT physical layer measurement period is typically a few times shorter than the corresponding inter-RAT cell search delay.

E-UTRAN to GSM event-triggered reporting. The event-triggered measurement reporting delay for a GSM cell with BSIC verified without layer 3 filtering is less than twice the GSM measurement period (see Section 22.3.3).

22.5 Mobility Performance

The mobility procedures comprising of cell selection and cell reselection in RRC_IDLE state and handover in RRC_CONNECTED state are described in Chapter 3. Here we focus on the performance requirements for these mobility procedures.

22.5.1 Mobility Performance in RRC_IDLE State

The performance requirements for mobility in RRC_IDLE state aim to ensure that a UE camps on a cell which guarantees good paging reception, that substantial UE battery power saving is achieved and that the interruption in paging reception during cell reselection is minimized.

All intra-frequency, inter-frequency and inter-RAT cell reselection requirements have the following characteristics:

- Requirements are specified for a selected set of typical DRX cycles: 0.32 s, 0.64 s, 1.28 s and 2.56 s;
- Cell reselection involves detection of new neighbour cells (both E-UTRA and inter-RAT), and measurement of those cells and of previously detected neighbour cells;
- Measurement and evaluation of cells are carried out at specific rates, which depend upon the DRX cycle in use;
- Cell reselection decisions are autonomously taken by the UE but are governed by pre-defined standardized rules, network control parameters and performance requirements;
- No performance requirements are specified for E-UTRA or inter-RAT cell identification, RSRP/RSRQ or inter-RAT measurements in RRC_IDLE, since in RRC_IDLE the UE is not required to report any event or measurement to the network.

The key elements of cell reselection in E-UTRAN are described in the following subsections.

22.5.1.1 Measurement and Evaluation of Serving Cell

The UE is required to measure both RSRP and RSRQ of the serving cell in order to evaluate the serving cell selection criterion (the ‘S-criterion’ – see Section 3.3) at least once every DRX cycle.

Failure of the S-criterion may imply that the UE is on the verge of losing the serving cell and it is therefore important that the UE identifies a potential new serving cell. Therefore if the S-criterion for the serving cell is not met over a certain number of consecutive DRX cycles (depending on the DRX cycle length in use), the UE initiates measurements of all the neighbour cells (i.e. over all frequency layers indicated by the serving cell) regardless of the measurement rules. For example in case of a 1.28 s DRX cycle, the UE starts measuring all the neighbour cells if the S-criterion for the serving cell is not met for 2 consecutive DRX cycles (i.e. for 2.56 s).

22.5.1.2 Intra-frequency Cell Reselection

Cell reselection to a neighbour cell on the same frequency is governed by a set of requirements for measurement and evaluation of intra-frequency cells.

Measurement of intra-frequency cells. Unlike the measurement of the serving cell, the neighbour cell measurements may not be performed every DRX cycle. The UE initiates the measurement of intra-frequency neighbour cells when the serving cell’s RSRP or RSRQ fall below their respective thresholds.

However, the cell ranking for cell reselection is only based on RSRP. The UE is required to detect and measure the RSRP of neighbour cells whose received quality is above the following thresholds, without an explicit intra-frequency neighbour cell list:

- Synchronization signal $\hat{E}_s/I_{ot} \geq -4$ dB
- RSRP and synchronization signal received power \geq respective thresholds, which depend on the frequency band e.g. -124 dBm for Band 1.

The RSRP of the identified intra-frequency cells is measured once every $T_{\text{measure, E-UTRA_intra}}$, as shown in Table 22.3.

Table 22.3: Measurement, detection and evaluation rates for intra-frequency cells.

DRX cycle length (s)	$T_{\text{detect, EUTRAN_intra}}$ (s) (number of DRX cycles)	$T_{\text{measure, EUTRAN_intra}}$ (s) (number of DRX cycles)	$T_{\text{evaluate, EUTRAN_intra}}$ (s) (number of DRX cycles)
0.32	11.52 (36)	1.28 (4)	5.12 (16)
0.64	17.92 (28)	1.28 (2)	5.12 (8)
1.28	32(25)	1.28 (1)	6.4 (5)
2.56	58.88 (23)	2.56 (1)	7.68 (3)

Evaluation of intra-frequency cells. The measured intra-frequency cells are evaluated for possible cell reselection based on cell ranking (see Section 3.3.4.3).

If an intra-frequency cell is detectable but not yet detected then the UE is required to evaluate whether it meets the reselection criteria based on ranking within $T_{\text{detect,EUTRAN_Intra}}$ (see Table 22.3) when $T_{\text{reselection}} = 0$. On the other hand if the cell is already detected then the UE is required to evaluate whether this intra-frequency cell meets the reselection criteria based on ranking within $T_{\text{evaluate,EUTRAN_Intra}}$ (see Table 22.3) when $T_{\text{reselection}} = 0$, provided that the target intra-frequency cell is ranked at least 3 dB above the serving cell. If the timer $T_{\text{reselection}}$ has a non-zero value and the target intra-frequency cell is found to be better ranked than the serving cell over the $T_{\text{reselection}}$ time, then that intra-frequency cell is selected.

22.5.1.3 Inter-frequency Cell Reselection

Similarly to the reselection of intra-frequency cells, requirements are also defined for measurement and evaluation of inter-frequency cells.

Measurement of inter-frequency cells. An inter-frequency layer may have a lower, equal or higher priority than that of the serving frequency layer. The UE is required to detect and measure the relevant measurement quantity under the same conditions as are applicable for intra-frequency cells.

If the serving cell's RSRP and RSRQ are above their respective thresholds the UE searches higher-priority inter-frequency layers at least once every $T_{\text{higher_priority_search}}$:

$$T_{\text{higher_priority_search}} = (60 * N_{\text{layers}}) \text{ seconds} \quad (22.1)$$

where N_{layers} is the total number of configured higher-priority inter-frequency and inter-RAT frequency layers. The relevant measurement quantities should be measured at least every $T_{\text{measure, E-UTRAN_Inter}}$ as defined in Table 22.4.

If the serving cell's RSRP and RSRQ become equal to or fall below their respective thresholds, the UE searches and measures all inter-frequency cells regardless of their priority. The relevant measurement quantity of the identified inter-frequency cells is measured every $K_{\text{carrier}} * T_{\text{measure, E-UTRAN_Inter}}$ seconds, as shown in Table 22.3, where K_{carrier} is the number of inter-frequency carriers configured by the serving cell.

Table 22.4: Measurement and Evaluation of Inter-frequency Cells.

DRX cycle length (s)	$T_{\text{detect,EUTRAN_Inter}}$ (s) (number of DRX cycles)	$T_{\text{measure,EUTRAN_Inter}}$ (s) (number of DRX cycles)	$T_{\text{evaluate,EUTRAN_Inter}}$ (s) (number of DRX cycles)
0.32	11.52 (36)	1.28 (4)	5.12 (16)
0.64	17.92 (28)	1.28 (2)	5.12 (8)
1.28	32 (25)	1.28 (1)	6.4 (5)
2.56	58.88 (23)	2.56 (1)	7.68 (3)

Evaluation of inter-frequency cells. The evaluation of inter-frequency cells for possible cell reselection is based on cell ranking for layers of equal priority and on absolute priorities for layers of different priority. Only RSRP is allowed for cell ranking, while the network can configure either RSRP or RSRQ for the evaluation of cells with unequal priorities [18].

If an inter-frequency cell of lower or equal priority is detectable but not yet detected then the UE is required to evaluate whether it meets the reselection criteria (see Sections 3.3.4.2 and 3.3.4.3) within $K_{\text{carrier}} * T_{\text{detect,EUTRAN_Inter}}$ seconds (see Table 22.4) provided that the cell reselection criteria can be met by a certain margin (5 dB for cell ranking and 6 dB for absolute priority based reselection) and the timer value $T_{\text{reselection}}$ is zero.

However, if an inter-frequency cell is already detected then the UE is required to evaluate whether it meets the reselection criteria within a shorter duration, $K_{\text{carrier}} * T_{\text{evaluate,EUTRAN_Inter}}$ seconds (see Table 22.4) provided the same two conditions above are satisfied.

If the timer $T_{\text{reselection}}$ has a non-zero value and the target inter-frequency cell is found to be better ranked than the serving cell over the time $T_{\text{reselection}}$, then that inter-frequency cell is reselected.

22.5.1.4 Inter-RAT Cell Reselection

Inter-RAT cell reselection covers reselection to UTRAN FDD, UTRAN TDD, GSM, CDMA2000 1xRTT and CDMA2000 HRPD.

As for inter- and intra-frequency reselection, cell reselection to an inter-RAT neighbour cell is also governed by the measurement and evaluation of the inter-RAT cells.

Measurement of inter-RAT cells. An inter-RAT frequency layer may have either lower or higher priority than that of the serving frequency layer. If the quality of the serving cell is above the threshold 'Snonintrasearch'²² the UE searches higher priority inter-RAT frequency layers at least once every $T_{\text{higher_priority_search}}$ according to Equation (22.1). The detected higher-priority inter-RAT cells of RAT_j are to be measured at least every $T_{\text{measure,RAT}_j}$ as defined in Table 22.5.

If the serving cell quality becomes equal to or falls below the Snonintrasearch threshold the UE searches and measures all inter-RAT cells of RAT_j (i.e. cells of higher, lower or equal priority frequency layers of RAT_j). The relevant measurement quantity of the identified inter-RAT cells is measured every $N_{\text{RAT},j} * T_{\text{measure,RAT}_j}$, as illustrated in Table 22.5, where $N_{\text{RAT},j}$ is the number of frequency layers of RAT_j configured by the serving cell (for GSM, $N_{\text{RAT}} = 1$). The relevant inter-RAT measurement quantities in RRC_IDLE states are described in Section 22.3.

Evaluation of inter-RAT cells. Evaluation of inter-RAT cells for possible cell reselection is based only on priority (see Section 3.3.4.2).

If an inter-RAT cell is detectable but not yet detected then the UE is required to evaluate whether it meets the reselection criteria within $N_{\text{RAT},j} * T_{\text{detect,RAT}_j}$ (see Table 22.5) provided that the cell reselection criteria can be met by a 6 dB margin (in the case of UTRA) and the timer value $T_{\text{reselection}} = 0$.

²²See Section 3.3.4.1; 'Snonintrasearch' is defined in [18, Section 5.2.4.7].

Table 22.5: Measurement and Evaluation of Inter-RAT Cells.

DRX cycle length (s)	apply to UTRA FDD and TDD	$T_{\text{detect, RAT},j}$ (s); $T_{\text{measure, RAT},j}$ (s) (number of DRX cycles); apply to all RATs	$T_{\text{evaluate, RAT},j}$ (s) (number of DRX cycles); apply to UTRA FDD, TDD and CDMA2000
0.32	30	5.12 (16)	15.36 (48)
0.64	30	5.12 (8)	15.36 (24)
1.28	30		6.4 (5) 19.2 (15)
2.56	60	7.68 (3)	23.04 (9)

For an inter-RAT cell that has already been detected, the UE is required to evaluate that it meets the reselection criteria within a shorter duration, $N_{\text{RAT},j} * T_{\text{evaluate, RAT},j}$ (see Table 22.4) provided that the cell reselection criteria of the respective RAT are satisfied.

22.5.1.5 Paging Interruption during Cell Reselection

During intra-frequency, inter-frequency or inter-RAT cell reselection procedures, the UE monitors the serving cell for paging messages until the UE is able to monitor the paging channels of the target cell. In order to complete the cell reselection successfully and camp on the new cell the UE has to acquire the relevant system information of the target cell. Therefore, during this period paging interruption might occur as the UE is not required simultaneously to receive paging and acquire the system information of the target cell.

The interruption in paging reception should not exceed $T_{\text{Interrupt}} = T_{\text{SI-RAT}} + 50(\text{ms})$, where $T_{\text{SI-RAT}}$ is the time required for acquiring all the relevant system information of the target RAT.

22.5.2 Mobility Performance in RRC_CONNECTED State

In E-UTRAN, only hard handovers are possible (see Section 3.2.3.4) resulting in a delay including a short interruption. In order to limit the length of the interruption, requirements are specified for various E-UTRAN handover scenarios as described in the following subsections. These requirements are expressed in terms of handover delay, which is the sum of the RRC procedure delay and the interruption time. This principle applies whether the target cell is known to the UE (referred to as *non-blind handover*), or unknown (*blind handover*). The interruption time is defined as the time from the end of the last subframe containing the handover command on the Physical Downlink Shared CHannel (PDSCH) from the serving cell and the moment the UE starts transmission on the relevant uplink physical channel in the target cell. Figure 22.6 illustrates the radio interface signalling of the handover procedure to a known target.

As explained in Section 3.2.3.4, the network does not require measurement reports from target cells for performing a blind handover. This is particularly useful in case of multiple frequency layers, which can only be monitored using the same gap pattern (i.e. either pattern #0 or pattern #1). In such a scenario the network may request the UE to perform measurements on only a subset of the frequency layers while relying on blind handovers for

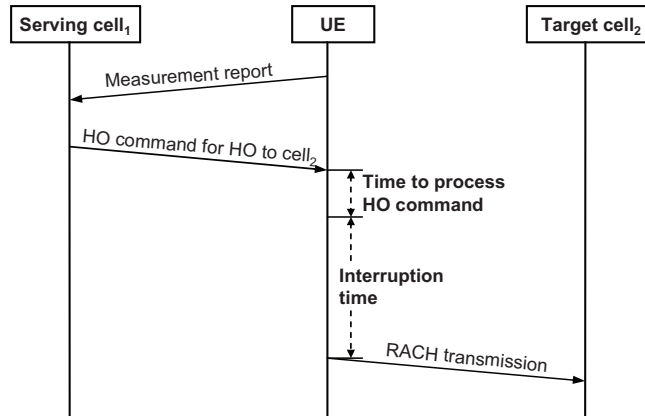


Figure 22.6: Handover to a known target cell.

the remaining ones. Handover to an unknown target cell results in a longer interruption time since the UE has to detect the target cell prior to accessing it. The process of handover to an unknown target cell is shown in Figure 22.7.




Figure 22.7: Handover to an unknown target cell.

22.5.2.1 E-UTRAN to E-UTRAN Handover

An E-UTRAN to E-UTRAN handover is completed when the UE starts transmission on the Physical Random Access CHannel (PRACH) in the target E-UTRA cell. The interruption time is expressed as $T_{interrupt} = T_{search} + T_{IU} + 20$ (ms), where:

- T_{search} is the cell search delay, which is 0 for known target cells and 80 ms for unknown target cells. The target cell is considered to be known if the cell search requirements were met in the last 5 s; otherwise the cell is considered to be unknown.
- T_{IU} is the timing uncertainty due to the timing of the PRACH occasions, and thus depends upon the PRACH configuration (see Section 17.4.2.5).

As an example, suppose the PRACH is allowed once in the middle of every frame (i.e. in every 4th subframe). If the target cell is known, the total intra-frequency or inter-frequency handover delay is 50 ms including 15 ms of RRC procedure delay (see [4, Section 11.2]). For the same PRACH configuration, in case of blind handover, the total handover delay is 130 ms including the 15 ms RRC procedure delay. In this case the cell search is much shorter than the usual cell search requirements (see Section 22.2.1) due to the fact that it is assumed that the target cell is sufficiently strong to be detected by the UE on the first correlation attempt.

22.5.2.2 Handover to Other 3GPP RATs

E-UTRAN to UTRAN handover

UTRA FDD or TDD cells can be known or unknown for the purposes of handover. The target cell is known if it has been measured by the UE during the last 5 s; otherwise it is considered unknown. An E-UTRAN-UTRAN FDD handover is completed when the UE starts transmission on the uplink Dedicated Physical Control CHannel (DPCCH) in the target cell. An E-UTRAN-UTRAN TDD handover is completed when the UE starts transmission of either the uplink DPCCH or an uplink synchronization code ('SYNC-UL') in the target cell.

Depending upon the choice of parameters, the interruption time when the target cell is known is typically in the order of 100–150 ms. When the target cell is unknown, the interruption increases by about 100 ms for UTRA FDD and 160 ms for UTRA TDD. The RRC procedure delay is 50 ms (see [12, Sections 5.3.1 and 5.3.2]) as this requirement is based on the UTRAN requirement [19]. Hence the overall handover delay can be of the order of 150–200 ms when the target UTRA (FDD or TDD) cell is known and 250–300 ms when the target cell is unknown. The exact handover delay can be derived from the relevant expressions in [12, Sections 5.3.1 and 5.3.2].

E-UTRAN to GSM handover

Handovers from the serving eNodeB to both known and unknown target GSM cells are supported in E-UTRAN. A handover to GSM is completed when the UE starts uplink transmission in the target GSM cell.

The total handover delays to known and unknown target GSM cells are 90 ms and 190 ms respectively, which includes 50 ms delay for processing the handover command (see [12, Section 5.3.3.1]).

22.5.2.3 Handover to CDMA2000 1xRTT or HRPD

Non-3GPP RAT handovers defined from E-UTRAN are for CDMA2000 1xRTT and HRPD. Handovers to both known and unknown target CDMA2000 1xRTT or HRPD cells are

possible. In both of these CDMA2000 technologies, a cell is considered known if it has been measured by the UE during the last 5 s. When performing handover to CDMA2000 1xRTT or HRPD the interruption time is influenced primarily by the following factors :

- Uncertainty due to changing the timing from the old E-UTRA serving cell to the new CDMA2000 1xRTT or HRPD cell (the delay can be up to one frame, i.e. 20 ms for CDMA2000 1xRTT and 26.7 ms for HRPD);
- The number of known and unknown target cells, which is signalled to the UE in the E-UTRAN System Information Block 8 (SIB8);
- The search window sizes for searching the known and unknown target cells; this is expressed in CDMA2000 1xRTT chips and signalled to the UE in SIB8.

The RRC procedure delay for processing the handover command is 130 ms for CDMA2000 1xRTT and 50 ms for HRPD. In order to derive the handover delay for different combinations of parameters, the reader is referred to [12, Section 5.4.2]. As a typical example the handover delay can be about 200 ms for CDMA2000 1xRTT and 130 ms for HRPD, assuming one target cell (regardless of whether it is known or unknown) and a typical window size of 60 chips.

22.6 RRC Connection Mobility Control Performance

In LTE the establishment and maintenance of the RRC connection is governed by the two main control plane functions: RRC Connection Re-establishment and Random Access.

22.6.1 RRC Connection Re-establishment

RRC connection re-establishment is initiated when a UE in RRC_CONNECTED state loses its RRC connection (e.g. due to radio link failure, handover failure or RRC connection reconfiguration failure), as specified in [4, Section 5.3.7.3]. The UE tries to re-establish the RRC connection with the strongest E-UTRA cell. Successful RRC re-establishment means that the UE is able to send the 'RRCConnectionReestablishmentRequest' message within $T_{\text{re-establish_delay}}$, including the delay in acquiring the uplink grant for sending the message and $T_{\text{UE-re-establish_delay}}$ which is defined as the delay from the moment when the UE detects the need for RRC re-establishment until it transmits a random access signal to the target cell:

$$T_{\text{UE-re-establish_delay}} = 50\text{ms} + N_{\text{freq}} * T_{\text{search}} + T_{\text{SI}} + T_{\text{PRACH}} \quad (22.2)$$

where N_{freq} is the total number of E-UTRA carrier frequencies available for RRC re-establishment, T_{search} is the target cell search delay and depends upon whether the target cell is known or unknown to the UE, T_{SI} is the time required to read the target cell System Information (SI) and T_{PRACH} is the delay due to random access.

22.6.2 Random Access

In LTE the random access procedure serves several purposes as described in Sections 17.2 and 17.3. Requirements are specified for both contention-based and contention-free random

access. The primary objectives of the requirements are to ensure the correct UE behavior when performing random access, and that the UE transmit power accuracy and transmit timing error when sending random access are within suitable limits. These requirements include:

- If the UE does not receive a random access response matching its PRACH preamble identity within the random access response window, it should retransmit a preamble, up to and not exceeding the maximum number of preamble transmissions configured by the eNodeB.
- The transmit power accuracy of the UE's PRACH preamble transmissions should fulfil the requirements defined in [20, Sections 6.3.5.1 and 6.3.5.2].
- The required transmit timing accuracy of the UE's PRACH preamble transmissions depends upon the transmission bandwidth, since at lower bandwidth the UE typically uses lower sampling rate. The transmit timing error due to all PRACH transmissions (initial and subsequent preamble transmissions) should be within $\pm 24T_s$ and $\pm 12T_s$ for system bandwidths equal to 1.4 MHz and ≥ 3 MHz respectively, as specified in [12, Section 7.1.2].

22.7 Radio Link Monitoring Performance

The purpose of the radio link monitoring function in the UE is to monitor the downlink radio link quality of the serving cell in RRC_CONNECTED state and is based on the cell-specific RSs (see Section 8.2.1). This in turn enables the UE when in RRC_CONNECTED state to determine whether it is *in-sync* or *out-of-sync* with respect to its serving cell (see [21, Section 4.2.1]).

In case of a certain number of consecutive out-of-sync indications (called 'N310'), the UE starts a network-configured radio link failure timer 'T310'. The timer is stopped if a number 'N311' of consecutive in-sync indications are reported by the UE's physical layer. Both the out-of-sync and in-sync counters (N310 and N311) are configurable by the network. Upon expiry of the timer T310, Radio Link Failure (RLF) occurs. As a consequence the UE turns off its transmitter to avoid interference and is required to re-establish the RRC connection within $T_{UE-re-establish_delay}$ as explained in Section 22.6.1. The various actions pertaining to radio link monitoring and the subsequent RRC re-establishment to the target cell are shown in Figure 22.8.

22.7.1 In-sync and Out-of-sync Thresholds

The UE's estimate of the downlink radio link quality is compared with out-of-sync and in-sync thresholds, Q_{out} and Q_{in} , for the purpose of radio link monitoring. These thresholds are expressed in terms of the BLock Error Rate (BLER) of a hypothetical Physical Downlink Control Channel (PDCCH) transmission from the serving cell (see [12, Section 7.6] for details). Specifically, Q_{out} corresponds to a 10% BLER while Q_{in} corresponds to a 2% BLER. The same threshold levels are applicable with and without DRX.




Figure 22.8: Radio link monitoring of the serving cell followed by RRC re-establishment to the target cell.

The mapping between the cell specific RS based downlink quality and the hypothetical PDCCH BLER is up to the UE implementation. However, the performance is verified by conformance tests defined for various environments [22].

22.7.2 Requirements without DRX

When no DRX is configured, out-of-sync occurs when the downlink radio link quality estimated over the last 200 ms period becomes worse than the threshold Q_{out} . Similarly without DRX the in-sync occurs when the downlink radio link quality estimated over the last 100 ms period becomes better than the threshold Q_{in} . Upon detection of out-of-sync, the UE initiates the evaluation of in-sync. The occurrences of out-of-sync and in-sync are reported internally by the UE's physical layer to its higher layers, which in turn may apply layer 3 (i.e. higher layer) filtering for the evaluation of RLF.

22.7.3 Requirements with DRX

When DRX is in use, in order to enable sufficient UE power saving the out-of-sync and in-sync evaluation periods are extended and depend upon the configured DRX cycle length. The UE starts in-sync evaluation whenever out-of-sync occurs. Therefore the same period ($T_{Evaluate_Q_{out_DRX}}$) is used for the evaluation of out-of-sync and in-sync. However, upon starting the RLF timer (T310) until its expiry, the in-sync evaluation period is shortened to 100 ms, which is the same as without DRX. If the timer T310 is stopped due to N311 consecutive in-sync indications, the UE performs in-sync evaluation according to the DRX-based period ($T_{Evaluate_Q_{out_DRX}}$).

22.7.4 Requirements during Transitions

In LTE a transition phase is caused by switching between DRX and non-DRX operation or switching between short and long DRX or vice versa. These scenarios can occur often, and

therefore the UE behaviour is specifically defined for the evaluation of the radio link quality during the transition period. There are two main aspects of the requirements:

- Length of the transition period, T_P ;
- Evaluation period, $T_{\text{Evaluate_Transition}}$, during the transition.

The transition period T_P is equal to the evaluation period of the mode after the transition. During this phase the evaluation period $T_{\text{Evaluate_Transition}}$ is defined as follows:

$$T_{\text{Evaluate_Transition}} \geq \min(T_{\text{Evaluate_mode1}}, T_{\text{Evaluate_mode2}}) \quad (22.3)$$

where $T_{\text{Evaluate_mode1}}$ and $T_{\text{Evaluate_mode2}}$ correspond to the evaluation periods of the first and the second mode respectively.

Equation (22.3) applies to both in-sync and out-of-sync evaluations. After the transition period, the UE uses an evaluation period corresponding to the second mode. The evaluation periods during the transition from short to long DRX and vice versa are illustrated in Figure 22.9.




Figure 22.9: Radio link monitoring evaluation period (a) during transition from long DRX to short DRX, and (b) during transition from short DRX to long DRX.

22.8 Concluding Remarks

In this chapter we have explained the main aspects of the LTE RRM performance requirements. These requirements play a key role in ensuring that the UE meets the desired mobility performance in a wide range of practical scenarios envisaged for LTE. More specifically, by virtue of these requirements, robust mobility in both RRC_IDLE and RRC_CONNECTED

states is ensured, both within E-UTRAN and between E-UTRAN and other RATs including UTRAN FDD/TDD, GSM and CDMA2000 technologies.

References²³

- [1] 3GPP Technical Specification 36.211, 'Evolved Universal Terrestrial Radio Access (E-UTRA); Physical Channels and Modulation', www.3gpp.org.
- [2] Texas Instruments, NXP, Motorola, Ericsson, and Nokia, 'R4-072215: Simulation Assumptions for Intra-frequency Cell Identification', 3GPP TSG RAN WG4, meeting 45, Jeju, Korea, November 2007, www.3gpp.org.
- [3] NXP, 'R4-080691: LTE Cell Identification Performance in Multi-cell Environment', www.3gpp.org 3GPP TSG RAN WG4, meeting 46bis, Shenzhen, China, February 2008.
- [4] 3GPP Technical Specification 36.331, 'Evolved Universal Terrestrial Radio Access (E-UTRA); Radio Resource Control (RRC); Protocol specification', www.3gpp.org.
- [5] 3GPP Technical Specification 36.300, 'Evolved Universal Terrestrial Radio Access (E-UTRA) and Evolved Universal Terrestrial Radio Access Network (E-UTRAN); Overall description; Stage 2', www.3gpp.org.
- [6] 3GPP Technical Specification 25.101, 'User Equipment (UE) Radio Transmission and Reception (FDD)', www.3gpp.org.
- [7] M. Mouly and M.-B. Pautet, *The GSM System for Mobile Communications*. Cell & Sys, 1992.
- [8] 3GPP Technical Specification 45.005, 'Radio transmission and reception', www.3gpp.org.
- [9] 3GPP Technical Specification 23.272, 'Circuit Switched (CS) fallback in Evolved Packet System (EPS); Stage 2', www.3gpp.org.
- [10] 3GPP Technical Specification 36.214, 'Evolved Universal Terrestrial Radio Access (E-UTRA); Physical Layer Measurements', www.3gpp.org.
- [11] Ericsson, 'R1-073041: Reference Signal Received Quality, RSRQ, Measurement', www.3gpp.org, 3GPP TSG RAN WG1, meeting 49bis, Orlando, USA, June 2007.
- [12] 3GPP Technical Specification 36.133, 'Evolved Universal Terrestrial Radio Access (E-UTRA); Requirements for Support of Radio Resource Management', www.3gpp.org.
- [13] M. Kazmi et al., 'Evaluation of Inter-Frequency Quality Handover Criteria in E-UTRAN', IEEE 69th Vehicular Technology Conference, Spring 2009.
- [14] 3GPP Technical Specification 45.008, 'Radio subsystem link control', www.3gpp.org.
- [15] 3GPP2 Technical Specification C.S0005-D, 'Upper Layer (Layer 3) Signaling Specification for cdma2000 Spread Spectrum Systems', www.3gpp2.org.
- [16] 3GPP2 Technical Specification C.S0011-C, 'Recommended Minimum Performance Standards for cdma2000 Spread Spectrum Mobile Stations', www.3gpp2.org.
- [17] 3GPP2 Technical Specification C.S0024-A, 'cdma2000 High Rate Packet Data Air Interface Specification', www.3gpp2.org.
- [18] 3GPP Technical Specification 36.304, 'Evolved Universal Terrestrial Radio Access (E-UTRA); User Equipment (UE) Procedures in Idle Mode', www.3gpp.org.
- [19] 3GPP Technical Specification 25.331, 'Radio Resource Control (RRC) Protocol Specification', www.3gpp.org.
- [20] 3GPP Technical Specification 36.101, 'Evolved Universal Terrestrial Radio Access (E-UTRA); User Equipment (UE) Radio Transmission and Reception', www.3gpp.org.

²³All web sites confirmed 1st March 2011.

- [21] 3GPP Technical Specification 36.213, 'Evolved Universal Terrestrial Radio Access (E-UTRA); Physical Layer Procedures', www.3gpp.org.
- [22] 3GPP Technical Specification 36.521-3, 'Evolved Universal Terrestrial Radio Access (E-UTRA); User Equipment (UE) conformance specification; Radio Transmission and Reception; Part 3: Radio Resource Management Conformance Testing', www.3gpp.org.

Paired and Unpaired Spectrum

Nicholas Anderson

23.1 Introduction

Expansion of consumer demand for cellular communications, as well as for the multitude of other wireless systems and applications, places a corresponding strain on the basic physical resource needed to support them: spectrum. The suitability of spectral resources to a particular application is governed by a range of inter-related factors of a technological, commercial or regulatory nature. Technological considerations influencing the choice of frequency band include propagation characteristics, antenna size and separation, the viability of Radio Frequency (RF) circuitry, and design implications resulting from the need to coexist with systems operating in neighbouring spectrum without causing (or suffering from) undue interference. These considerations determine in part the commercial viability of a system. Signal range and spectral efficiency govern coverage and capacity, and hence determine the required number of base station sites and the capital outlay, whilst terminal costs and form factor affect the acceptability of products in the marketplace.

From a regulatory perspective, in spite of significant coordination at an international level through the International Telecommunications Union (ITU), regional variations in governmental policy regarding spectrum and technology are inescapable due to the differing needs of each region and the different historical factors which have shaped their present-day spectrum allocations. Thus, the availability of globally harmonized spectrum for a particular technology cannot be guaranteed. Furthermore, in many regions, a tendency towards technology neutrality within viable spectrum assignments (see, for example, [1]) is beginning to liberalize the traditional one-to-one mappings between technologies and their addressable spectrum and is hence opening up new markets.

LTE – The UMTS Long Term Evolution: From Theory to Practice, Second Edition.

Stefania Sesia, Issam Toufik and Matthew Baker.

© 2011 John Wiley & Sons, Ltd. Published 2011 by John Wiley & Sons, Ltd.

Against this background, a key design goal of LTE was to enable deployment in a diverse range of spectrum environments in terms of bandwidth, uplink-downlink duplex spacing¹ and uplink-downlink asymmetries. By supporting paired spectrum allocations (separate uplink and downlink carriers) in addition to stand-alone unpaired allocations (uplink and downlink operating on the same carrier frequency), wastage of valuable spectrum can be avoided.

It is also the case that traffic is increasingly data-centric and often asymmetric. In unpaired spectrum, asymmetry may be provided through the use of unequal duty cycles in the time domain for uplink and downlink. In paired spectrum, asymmetry is also possible via the deployment of unequal bandwidths for uplink and downlink.

Considering the nature of the regulatory environment, an increase in demand and competition for spectrum and the trends of mobile traffic and service usage, it is advantageous for the LTE system to be designed such that it may be flexibly adapted to diverse spectral assignments including both paired and unpaired bands. The market addressable by the technology is thereby increased.

23.2 Duplex Modes

The term ‘duplex’ refers to bidirectional communication between two devices, as distinct from unidirectional communication which is referred to as ‘simplex’. In the bidirectional case, transmissions over the link in each direction may take place at the same time (‘full duplex’) or at mutually exclusive times (‘half-duplex’).

It should be noted that a communication link being half-duplex need not imply that only half-duplex user services are supported (such as ‘push-to-talk’ voice applications). Full-duplex services such as a normal telephone conversation may of course be carried over half-duplex communication systems. The key factor differentiating the duplex nature of the service from that of the underlying communication link is the timescale over which the communication direction is cycled in relation to the service duplex timescale: if the link direction is cycled at a sufficiently high rate then the duplex nature of the link can be concealed when viewed from the perspective of the user application.

In the case of a full-duplex transceiver, the frequency domain is used to separate the inbound and outbound communications; that is, different carrier frequencies are employed for each link direction. This is referred to as Frequency Division Duplex (FDD). As discussed in Section 21.4.2, the ability of a full-duplex transceiver to transmit and receive at the same time instant is enabled via the use of a duplexer – a tuned filter network able to provide a high degree of isolation between the inbound and outbound signals on the different carrier frequencies (often sharing the same antenna). However, these filter networks incur some signal attenuation. For the receiver, this attenuation occurs before the low-noise amplifier in the signal path and hence contributes directly to the receiver noise figure and degrades its sensitivity. For the transmitter, the duplexer follows the high power amplification stage, requiring either a higher-powered amplification device to overcome the loss, or tolerance of the corresponding reduction in communication range.

Conversely, in the case of a half-duplex transceiver, the time domain provides the necessary separation between the inbound and outbound communications. When the same

¹ ‘Duplex spacing’ is the term used to describe the size of the frequency separation between uplink and downlink carriers.

carrier frequency is used for each link direction, the system is said to be purely Time Division Duplex (TDD) and incorporates half-duplex transceivers at each end of the radio link. Alternatively different carrier frequencies may be used, in which case the system is often known as ‘Half-Duplex FDD’ (HD-FDD). The popular GSM system utilizes HD-FDD, with uplink and downlink communications taking place for a particular user not only on distinct carrier frequencies but also at different times. HD-FDD can therefore be seen as a hybrid combination of FDD and TDD. By scheduling users (or sets of users) at mutually exclusive times, full occupancy of the transmission resources may be achieved without requiring simultaneous transmission and reception at each mobile terminal. Thus for HD-FDD, the base station is full-duplex whilst the terminals are half-duplex.

The three duplex modes outlined above are represented diagrammatically in Figure 23.1.




Figure 23.1: FDD, TDD and HD-FDD duplex modes.

Both HD-FDD and TDD carry the advantage that terminals may be developed without the need for duplexers, thus simplifying the design. HD-FDD operation is also potentially useful for paired bands in which the duplex spacing is relatively narrow (e.g. less than a few times the system bandwidth) since it avoids the need for a technically challenging duplexer design in the terminal and can thereby help in reducing its cost. At the base station there is more scope to implement such technically challenging designs due to less onerous size and cost constraints and because base station products are usually targeted towards a specific frequency band.

The spectrum bands defined for LTE are listed in Section 21.2. The paired bands are all currently symmetrically dimensioned between uplink and downlink, although this may change in the future.

23.3 Interference Issues in Unpaired Spectrum

As discussed in Chapter 21, the spectral emissions from an eNodeB or from a UE are unfortunately not strictly band-limited within the desired carrier bandwidth. This is due to practical limitations on filter technology and on the linearity of amplification.

In addition, receivers suffer from non-ideal suppression of signals falling outside the desired reception band. There is thus the potential for interference between transmitter and receiver. This is shown in Figure 23.2, where all sources of attenuation of the signal between transmitter and receiver (due to phenomena such as propagation, antenna radiation patterns and cable losses) are denoted by a single ‘coupling loss’. In the region marked ‘A’, energy from the transmitter falls directly within the desired passband of the receiver, while in region ‘B’ the non-ideal characteristic of the receiver collects energy from inside the passband of the transmitter. In region ‘C’ transmitted energy falling outside both the transmit and receive passbands is collected by the non-ideal receiver characteristic. The aggregated effects of regions A, B and C represent the total unwanted energy from the interfering transmitter that is captured by the receiver.²




Figure 23.2: Frequency-domain interference between transmitter and receiver.

Interference arising from regions A and B is controlled by performance requirements imposed on the transmitting and receiving devices respectively, while in region C both transmitter performance and receiver performance have an impact. The interference effects are generally small for large frequency separations (as in the majority of FDD cases), and when there is a large coupling loss between the interfering transmitter and a ‘victim’ receiver, as is the case when there is a large physical separation between the two.

Naturally, the interference effects are most pronounced when the transmitter and receiver operate on adjacent carrier frequencies (excepting the co-channel case) and with low coupling loss (e.g. due to small spatial separation). These worst-case adjacent-channel scenarios form the basis upon which key requirements for transmitter and receiver are generally set.

²The limits on these emissions according to the LTE specification are explained in Chapter 21.

The overall ‘leakage’ (whether at the transmit side or the receive side) from a transmission on one carrier into a receiver operating on an adjacent carrier is described by the *Adjacent Channel Interference Ratio* (ACIR), which is derived from the transmitter’s *Adjacent Channel Leakage Ratio* (ACLR) and the receiver’s *Adjacent Channel Selectivity* (ACS) as defined in Section 21.3.2.1. It is worth noting that it is often the case that the ACIR is dominated by either the ACS or the ACLR. For example, it is technologically difficult to design a mobile transmitter with an ACLR that approaches or exceeds the ACS of a base station receiver.

As is discussed in the following section, the presence of imperfect ACIR has implications for the deployment of systems at a boundary between unpaired and paired spectrum allocations, and also for unsynchronized systems operating in closely spaced unpaired allocations.

23.3.1 Adjacent Carrier Interference Scenarios

For an FDD cellular system, adjacent channel frequency separation of an interfering transmitter and a victim receiver naturally implies that the interferer and victim are of differing equipment types (i.e. one is a mobile terminal whilst the other is a base station). Transmitter–receiver interference between one User Equipment (UE) and another, or between one eNodeB and another is avoided by virtue of the duplex spacing.

The same is also generally true in TDD systems if they are time-synchronized so that overlap between uplink and downlink transmission periods is avoided. However, when synchronization is not or cannot be provided, or when TDD systems operate on carriers adjacent to an FDD system, the possibility arises for interferer and victim to be of the same device type. Figure 23.3 depicts a relatively common scenario in which an unpaired spectral allocation is located in a region between an FDD downlink band and an FDD uplink band (as is the case for the 2.5–2.6 GHz UMTS extension band, for example).

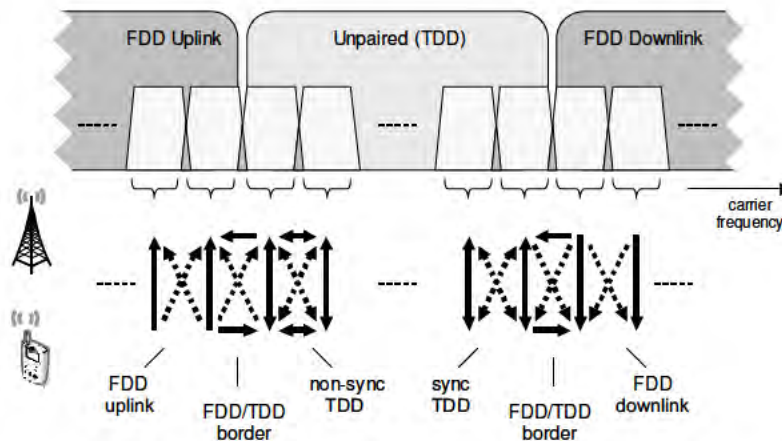


Figure 23.3: Possibilities for Adjacent Carrier Interference (ACI) between FDD and TDD systems.

The vertical arrows in Figure 23.3 represent the desired communication between the base station and the mobile terminal, unidirectional on a per-carrier basis for FDD and bi-directional for TDD. The diagonal dotted lines represent base-station-to-mobile and mobile-to-base-station adjacent channel interference that results from imperfect ACIR.

Two TDD scenarios are shown, synchronized and non-synchronized – the latter encompassing any general possibility for partial or full uplink-downlink overlap in time between two systems. For the synchronized case, the interference scenarios between the base station and mobile terminal (and vice versa) are the same as for the corresponding FDD-to-FDD adjacent channel cases. However, for the non-synchronized case additional interference paths exist between TDD mobiles and between TDD base stations, represented by the horizontal bidirectional arrows. At the FDD/TDD border regions, these same ‘horizontal’ interference paths exist but are unidirectional in nature.

There are many facets of a deployment which affect the severity of these various interference paths. For example, the locations of the interfering transmitter and victim receiver, as well as the characteristics of the propagation between them, clearly influence the overall coupling that exists. Macrocellular deployments typically use base station transmit antennas mounted on masts located above roof-top level, thereby resulting in an increased likelihood of Line-Of-Sight (LOS) propagation between base stations, with a correspondingly low path-loss exponent. The common use of macrocell base station antennas with vertical directivity and hence high gain can further worsen this situation. These aspects are less problematic for microcellular and dense-urban deployments in which LOS propagation between base stations is less likely due to their antennas being located below rooftop level.

Coupling between mobile terminals is often mitigated by the surrounding local clutter, and due to the lower antenna gain in the terminals. One typical scenario is depicted in Figure 23.4 in which non-co-located macro base stations have some potential to exhibit stronger mutual coupling than between the mobiles which they respectively serve. However, the figure also shows that it is not always possible to rely on local clutter to provide the necessary isolation between terminals, due to the fact that when terminals are closely spaced (for example, in the same office or café), LOS propagation again becomes more likely and the potential for interference is increased – as a result of both the lower path-loss exponent between the terminals and the small physical separation between them.

The base-to-base and UE-to-UE interference scenarios that are particular to unsynchronized TDD deployments and to TDD deployments adjacent to FDD deployments are reviewed in more detail in the following two subsections.

23.3.1.1 Base-Station to Base-Station Interference

Base stations of relevance to a particular base-to-base interference scenario may be either co-located (i.e. antennas mounted at the same cell-site) or non-co-located. Nevertheless, base-to-base interference is generally deterministic. This is because the locations of the base stations are fixed, and furthermore the link adaptation strategy typically employed for the LTE downlink usually results in all available transmit power being used to maximize the throughput of the link. It is therefore reasonable to analyse the interference assuming full transmit power from each base station.

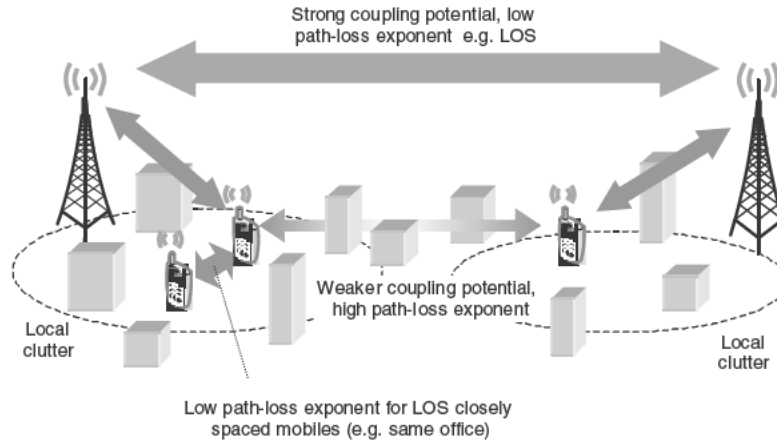


Figure 23.4: Typical RF interference scenario for a TDD system.

In general, systems operating on adjacent carriers may use differing deployment topologies and cell sizes, giving rise to varying distances between the base stations on the adjacent carriers. We therefore consider the worst-case base-to-base distance only, and assume that at some point in the network the transmit and receive antenna patterns are aligned to provide maximum gain at this worst-case distance. The co-channel Power Spectral Density (PSD) of the received interference at the victim base station antenna connector can then be written as:

$$PSD_{Rx} = PSD_{Tx} + 2G_{BS} - \rho_{BS-BS}(x_0) \tag{23.1}$$

where PSD_{Tx} is the transmitted PSD at each interfering base station antenna connector, G_{BS} is the antenna gain at each base site, $\rho_{BS-BS}(x)$ is the path-loss between base sites as a function of distance x in metres and x_0 is the worst-case (smallest) distance between base sites of different carriers.

In order for the inter-system interference to have only a minor effect, one can assume that the PSD of the interference after benefiting from any available ACIR should be of the order of $PSD_N + NF_{BS} - 6$ dB or less if it is to produce no more than a 1 dB desensitization of the base station receiver (where NF_{BS} denotes the base station noise figure and PSD_N is the PSD of thermal noise, e.g. -174 dBm/Hz at typical temperatures):

$$(PSD_{Rx} - ACIR) \leq (PSD_N + NF_{BS} - 6) \text{ dB} \tag{23.2}$$

Note that in the case of an SC-FDMA³ uplink victim receiver, consideration needs to be paid not only to the ACIR averaged over the system bandwidth, but, more challengingly, to the localized frequency resource blocks located closest to the interfering carrier, especially if the important uplink control signalling on the Physical Uplink Control CHannel (PUCCH) is to be protected (see Section 16.3.1). In general, however, the ACIR requirement varies directly with PSD_{Tx} :

$$ACIR \geq PSD_{Tx} + 2G_{BS} - \rho_{BS-BS}(x_0) - PSD_N - NF_{BS} + 6 \text{ dB} \tag{23.3}$$

³Single Carrier Frequency Division Multiple Access.

In order to arrive at an ACIR requirement, we must therefore know the transmitted PSD and the intervening path-loss. To do so, it is reasonable to assume that the base station transmit power capabilities are dimensioned in order that each of the two systems are interference-limited on the downlink (at least this can apply for small- and medium-sized cells without exceeding the eNodeB output power capabilities). We therefore assume here that the spectral density of the downlink signal for 95% of the total area is γ_{DL} dB larger than the spectral density of the thermal noise in the UE receivers. Thus

$$\text{PSD}_{Tx} = \text{PSD}_N + \text{NF}_{UE} + \gamma_{DL} + \rho'_{BS-UE} - G_{BS} \quad (23.4)$$

where NF_{UE} is the noise figure of the UE receiver, γ_{DL} is the received downlink signal-to-noise ratio in dB which is exceeded at 95% of the UE receivers, and ρ'_{BS-UE} is the 95-percentile path-loss between UEs and their serving eNodeBs in the interfering network.

For ease of representation, an empirical approximation to ρ'_{BS-UE} is applied here specific to this particular example deployment: let σ be the standard deviation of the log-normal shadow fading between eNodeB and UE and L_b represent the additional building penetration loss (assuming indoor coverage); then with $\rho_{BS-UE}(x)$ denoting the path-loss between base stations and UEs separated by distance x metres, we can write

$$\rho'_{BS-UE} \approx \rho_{BS-UE}(0.58s) + 0.7\sigma + L_b \quad (23.5)$$

Substituting Equations (23.5) and (23.4) into (23.3) we obtain an approximate expression for the necessary ACIR to maintain an acceptable adjacent channel interference level at a victim base station:

$$\text{ACIR} \geq \text{NF}_{UE} - \text{NF}_{BS} + \gamma_{DL} + \rho_{BS-UE}(0.58s) + 0.7\sigma + L_b + G_{BS} - \rho_{BS-BS}(x_0) + 6 \text{ dB} \quad (23.6)$$

This is plotted in Figure 23.5 as a function of the smallest inter-carrier eNodeB–eNodeB separation x_0 for several selected values of the inter-site spacing s between co-channel base-stations on the interfering carrier under the following assumptions:

- uniform base station deployment on the interfering carrier;
- free-space propagation between eNodeBs $\rho_{BS-BS}(x)$;
- $\rho_{BS-UE}(x) = 128.1 + 37.6 \text{ dB}$, from [2];
- $\gamma_{DL} = 6 \text{ dB}$, $\text{NF}_{UE} = 9 \text{ dB}$, $\text{NF}_{BS} = 5 \text{ dB}$, $G_{BS} = 14 \text{ dBi}$, $L_b = 20 \text{ dB}$.

The transmit PSD of the eNodeB (P_{eNB}) for 95% coverage is also listed in the legend for each Inter-Site Distance (ISD).

For a given ISD, the required ACIR naturally decreases as the worst-case separation between interfering eNodeBs is increased. Notice, however, that the base-to-base problem worsens significantly as the ISD in the interfering network increases, due to the fact that the path-loss exponent from eNodeB to UE is higher than the path-loss exponent between eNodeBs. The transmit power needed by the interfering eNodeB to reach UEs at its cell edge increases at a faster rate than can be compensated by the path-loss to a victim eNodeB receiver at the same cell-edge location. It should be remembered, however, that this is representative of a macrocellular scenario with eNodeB antennas mounted above rooftop level; for the smaller cell sizes (characteristic of microcells) the eNodeB-to-eNodeB situation




Figure 23.5: Required ACIR for 1 dB desensitization in an eNodeB-to-eNodeB interference scenario.

will be significantly improved by the higher propagation exponent between eNodeBs whose antennas are mounted below rooftop level.

For co-sited eNodeBs (i.e. very small x_0), a Minimum Coupling Loss (MCL) value of 30 dB has been used to replace $2G_{BS} - \rho_{BS-BS}(x_0)$ in Equation (23.1), resulting in

$$ACIR_{co-siting} \geq PSD_{Tx} - MCL - PSD_N - NF_{BS} + 6 \text{ dB} \tag{23.7}$$

For wide area base stations, the ACIR required for co-siting can rise towards a challenging 120 dB or so. This problem, however, is not new or specific to LTE and has been encountered previously for FDD/TDD coexistence in WCDMA. Some practical solutions to this problem have been documented in [3], in which RF bandpass cavity resonator filters were used to improve greatly the ACLR and ACS of base station transmitters and receivers respectively either side of a TDD/FDD boundary. These significantly exceed the standardized minimum requirements which were not intended to cope unaided with the case of co-sited base stations.

Similar techniques also apply to LTE yet remain significantly challenging. With careful design, however, adjacent channel deployment of FDD and TDD LTE base stations, or of two non-synchronized LTE TDD base stations, should be feasible, even for co-sited arrangements, provided that appropriate measures are adopted in both the interfering and victim base stations.

23.3.1.2 Mobile-to-Mobile Interference

The UE-to-UE interference scenario requires a more probabilistic approach than for base-to-base interference for the following reasons:

- the locations of the interferers and victims are variable and dynamic;
- the physical resources assigned by the scheduler to the UEs are variable;
- the transmit powers of the interfering UEs are a function of their channel conditions and of the power control policy implemented;
- the received levels of the wanted base station signals at the victim UEs are also variable as a function of the UEs' channel conditions.

It is difficult, therefore, to formulate a definitive analysis of UE-to-UE interference for LTE. Nonetheless, a basic analysis is presented here together with some discussion of the attributes of the LTE system which have some bearing on the magnitude of the interference effects.

We consider two similar overlaid tri-sectored LTE deployments on adjacent 5 MHz carriers. One deployment contains the interfering UEs while the other contains the victim UEs. Both carriers have regions in which uplink transmissions overlap in time with downlink transmissions. The base stations of the deployments are either co-located or maximally spaced non-co-located. Cells in the interfering network each schedule groups of four contiguous uplink Resource Blocks (RBs) to a number of randomly selected UEs, resulting in six simultaneously scheduled interferers per subframe.

The impacts of the scheduled interferers' transmissions on a UE receiver in the victim network are calculated for the case in which a randomly selected victim UE is scheduled a downlink transmission resource in the one RB next to the band edge separating the two carriers. The impact caused by the interfering adjacent-carrier UEs is analysed in terms of the mean percentage reduction in victim UE downlink throughput, R_{loss} , caused by the presence of the interfering UEs.

As assumed in [2], the ACIR increases by an additional 13 dB for localized SC-FDMA interferer transmissions located anywhere other than the four RBs next to the band edge of the interferer network. Additionally, for the purposes of this analysis an uplink power control strategy is employed whereby the transmit PSD of each UE is set such that it is not received at any co-channel non-serving eNodeB receiver any higher than 6 dB above the eNodeB receiver's thermal noise floor.

The transmit PSD of the eNodeB in the victim network is set via Equations (23.4) and (23.5) in the same manner as for the eNodeB-to-eNodeB analysis, such that the downlink is in an interference-limited region of operation (but is not excessively 'over-powered'). The same path-loss model between eNodeBs and UEs is assumed. The path-loss between UEs is assumed to be given by Equation (23.8) (with a carrier frequency f_c of 2000 MHz), based upon a simple two-slope microcellular model from [4] with break point at $x_b = 45$ m to reflect the likelihood of free-space-like propagation (with exponent 2) for low separation distances (i.e. $x \leq x_b$), and increased attenuation exponent $z = 6.7$ due to local clutter at higher

distances ($x > x_b$).

$$\rho_{UE-UE}(x) = \begin{cases} -27.56 + 10 \log_{10}(f_c^2 x^2) \text{ dB} & \text{for } x \leq x_b \\ -27.56 + 10 \log_{10}\left(\frac{f_c^2 x^z}{x_b^{z-2}}\right) \text{ dB} & \text{for } x > x_b \end{cases} \quad (23.8)$$

A microcellular model is considered applicable to UE-to-UE interference as it reflects the case where both the transmitting and receiving antennas are below rooftop level. Other system parameters assumed for this analysis are generally in line with those of [4] for macrocell simulation. With these assumptions, and for $L_b = 20$ dB, the results of Figure 23.6 are obtained, displaying the relationship between the band-edge ACIR and the throughput loss R_{loss} for various values of ISD s .




Figure 23.6: Throughput reduction due to UE-to-UE interference.

The results in Figure 23.6 are given for the case of co-located base stations. Those for the non-co-located case are very similar showing only a small further degradation for low ACIR values; the absence of a significant difference is a result of the uplink power control strategy employed as described above (whereby the mobile transmit power is correlated with the strongest non-serving cell path-loss rather than with the path-loss of the serving cell).

The LTE specifications are based upon a 30 dB ACLR which would provide an ACIR of 28 dB assuming an ACS of 33 dB. In this case, it can be seen from Figure 23.6 that the

worst-case throughput loss for the band-edge downlink resource block would be between 7% and 16% depending on cell size.

This analysis is, however, rather sensitive to certain assumptions, especially the value of the in-building penetration loss L_b . When the penetration loss is increased, the serving eNodeB may instruct the UE to increase its transmit power by the same amount in order to maintain its received Signal-to-Interference-plus-Noise Ratio (SINR) at the serving eNodeB (subject to maximum UE output power constraints) without causing additional interference to non-serving eNodeBs which are also protected by the same building penetration loss. However, the path-loss to a worst-case nearby victim UE (e.g. with LOS propagation between the UEs) is not affected by the increased building penetration loss. Thus, increased building penetration loss can have the effect of increasing worst-case UE-to-UE interference levels. This effect is clearly evident from the fact that the curves of Figure 23.6 for the case of $L_b = 20$ dB show a greater R_{loss} than the dotted curves for $L_b = 0$ dB.

This suggests that the susceptibility of the system to UE-to-UE interference is lowered considerably in an outdoor scenario. In these cases, the system's need to control the uplink inter-cell interference between UEs and neighbouring eNodeBs in a frequency-reuse-1 network constrains the quantity of interference power that is injected by those UEs into adjacent carriers. System throughput loss is then minimal ($\sim 2\%$) for commonly expected ACIR levels. The fact that a reduction in penetration loss mitigates UE-to-UE interference also points towards possible cell-planning solutions to alleviate the problem, for example using picocells or home base stations rather than macrocells to provide in-building coverage.

The statistical nature of UE-to-UE interference is also of relevance when assessing its impact. LTE allows for randomization of the allocated radio resources for both the interferer and the victim in both the time and frequency domains. Uplink frequency hopping is able to provide the necessary randomization in frequency, and in the time domain a degree of randomization can be provided by different resource scheduling strategies, as well as the possibility for differing retransmission delays due to the fact that the downlink retransmissions to the victim UE are dynamically scheduled rather than synchronous. The use of Hybrid Automatic Repeat reQuest (HARQ) also provides robustness against those instantaneous events in which high interference levels are experienced. Thus, in the case of frequency-adjacent LTE systems, UE-to-UE interference may be heavily randomized. Its effects can therefore be 'smoothed' and shared amongst all users of the system on a probabilistic basis, helping to avoid persistent effects on specific pairs of users with close RF coupling.

UE-to-UE interference may also be alleviated by receive processing at the UE. As discussed in Section 26.2.6, Interference Rejection Combining (IRC) receivers [5] can be used to maximize the received SINR, taking into account the instantaneous direction of arrival of the wanted and interfering signals. Forms of the IRC receiver that make use of averaged correlation (e.g. in time or in frequency) of the received signals across antennas (e.g. [6]) are particularly applicable to the adjacent-channel UE-to-UE interference scenario in which explicit channel estimation of interferers is likely to be impractical. A nearby closely coupled UE on an adjacent channel would typically present a single dominant interference source, enabling an IRC receiver to provide a gain when the instantaneous interference-to-thermal-noise ratio is relatively high. Thus, one could anticipate that use of the IRC receiver may provide some additional robustness against UE-to-UE interference. Figure 23.7 confirms that this is the case (here shown for a 500 m ISD). This figure shows the benefits that an IRC

receiver can deliver compared to a Maximum Ratio Combining (MRC) receiver in terms of mean user throughput in the presence of UE-to-UE interference from an adjacent carrier. The figure also shows that an IRC receiver provides benefits even in the absence of adjacent-carrier UE-to-UE interference (due to the removal of interference from co-channel eNodeBs).




Figure 23.7: Comparison of IRC and MRC receivers with and without adjacent channel UE-UE interference.

More generally, Figure 23.7 reveals that for the scenario considered, a system using ideal IRC receivers in the presence of UE-to-UE interference with 29 dB ACIR could perform as well as a system using Maximum Ratio Combining (MRC) receivers with no UE-to-UE interference at all. Thus, the detrimental effects of UE-to-UE interference may be mitigated via a combination of moderate ACIR together with IRC receivers.

23.3.2 Summary of Interference Scenarios

The preceding sections have provided some discussion of the issues facing the deployment of LTE TDD on carriers adjacent either to other non-synchronized TDD systems, or to FDD downlink or uplink carriers. In these deployment scenarios, one must consider not only eNodeB-to-UE and UE-to-eNodeB interference but also the nature and severity of eNodeB-to-eNodeB and UE-to-UE interference.

In the case of eNodeB-to-eNodeB interference, the discussions of Section 23.3.1.1 have shown that very stringent ACLR and ACS are needed, especially in the case of co-siting. Nonetheless, co-siting would appear to be technically feasible, for example by means of cavity-based low-loss RF filtering solutions at the base stations.

In the case of UE-to-UE interference, system-level analysis suggests that by using appropriate radio resource management on the uplink, together with interference randomization and possibly also interference suppression at the receiver, the overall system throughput loss and user experience in the presence of the adjacent carrier interference are likely to remain acceptable for many deployment scenarios.

Although it is generally desirable to enable alignment of uplink–downlink switching points between the LTE TDD system and LTE or legacy TDD systems on an adjacent carrier, it should be remembered that this cannot always be relied upon.

23.4 Half-Duplex System Design Aspects

For half-duplex (including TDD and HD-FDD) operation, the restriction that a UE may not transmit and receive at the same time instant has consequences for the physical layer (and to a lesser extent, higher-layer) design of LTE.

In order to enable the exploitation of economies of scale, LTE has generally followed a design principle in which differences between the duplex modes are introduced only where necessary for correct system operation, or where they offer a significant performance advantage when used for a particular duplex mode. With careful design, only a relatively small set of attributes need to be modified. Nonetheless, these have the potential to alter certain behaviours and structures of the physical layer, primarily in terms of frame structure, HARQ operation and control signalling, which are discussed in the following sections.

23.4.1 Accommodation of Transmit–Receive Switching

For TDD systems, switching between transmit and receive functions occurs on the transition from uplink to downlink (for the UE) and on the transition from downlink to uplink (for the eNodeB). For half-duplex FDD systems, switching only occurs at the UE, as the eNodeB is assumed to be full-duplex.

In order to preserve the frequency-domain orthogonality of the LTE uplink multiple access scheme, propagation delays between an eNodeB and the UEs under its control are compensated by means of timing advance as explained in Section 18.2.

At a half-duplex UE, the timing-advanced uplink transmission cannot be allowed to overlap with reception of any preceding downlink. For TDD, to prevent the overlap, a transmission gap or ‘guard period’ between transmission and reception at the eNodeB is created (T_{G1}) to accommodate the greatest possible timing advance and any required switching delay (including power amplifier ramp-up or ramp-down to avoid excessive wideband emissions). A further guard period (T_{G2}) is also required at the TDD eNodeB transition between uplink and downlink to cater for switching and power ramping delays only (this being independent of the propagation delay or timing advance). These are illustrated in Figure 23.8.

Four switching times are therefore of relevance in the case of TDD operation. These correspond to the transmit-to-receive and receive-to-transmit delays at the UE (denoted

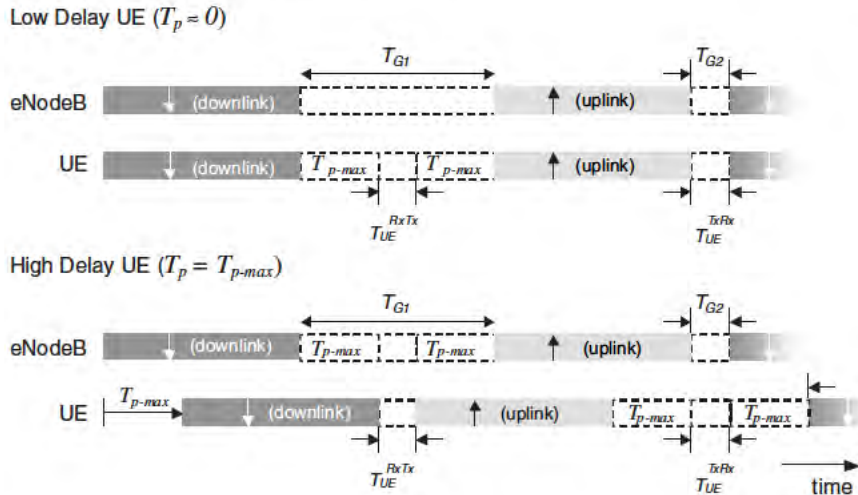


Figure 23.8: TDD signal timings in the presence of uplink timing advance.

as time intervals T_{UE}^{TxRx} and T_{UE}^{RxTx} respectively) and likewise at the eNodeB (denoted T_{eNB}^{TxRx} and T_{eNB}^{RxTx}). Figure 23.8 depicts two cases corresponding to the two extremes of propagation delay T_p within a TDD cell ($T_p = 0$ for a UE physically close to the eNodeB and $T_p = T_{p_max}$ for a UE at the border of a cell, where $T_{p_max} = d_{max}/c$ corresponds to the maximum one-way propagation delay supported by the cell, occurring at distance d_{max}). Note that the switching delays are exaggerated for diagrammatical clarity and that UE switching delays are assumed to be longer than those at the eNodeB.

It is apparent from Figure 23.8 that the time available at the UE for downlink to uplink transition is a function of the propagation delay T_p (most stringent for the case of high delay) whereas the time available at the eNodeB for the same transition is constant and equal to T_{G1} :

$$T_{G1} = 2T_{p_max} + T_{UE}^{RxTx} \tag{23.9}$$

The time interval T_{G2} at the eNodeB is independent of the propagation delay. To support the case for which $T_p \rightarrow 0$ (i.e. a UE close to the eNodeB), T_{G2} needs to be dimensioned such that

$$T_{G2} = \max(T_{UE}^{TxRx}, T_{eNB}^{RxTx}) \tag{23.10}$$

In the case of HD-FDD, $T_{eNB}^{RxTx} = 0$, so T_{G2} is determined only by the time T_{UE}^{TxRx} . In order to support the case of low T_p , the (full duplex) eNodeB must still allow sufficient time for this UE switching delay if the uplink and downlink subframes surrounding the switching point are both active for a particular user. Hence in practice the uplink frame timing at the eNodeB should be advanced for the whole cell by an amount T_{UE}^{TxRx} relative to the downlink frame timing at the eNodeB even for a full-duplex eNodeB if it supports HD-FDD UEs in the cell. Full duplex FDD UEs communicating with the same eNodeB will likewise need to have their timing advanced to maintain uplink orthogonality with the HD-FDD UEs.

It is important to note that d_{max} may be significantly larger than the notional cell radius r_0 (i.e. half the ISD), due to propagation effects such as shadow fading. This effect is shown for

one example of a tri-sector deployment with frequency reuse factor 1 in Figure 23.9 (the shadow fading is assumed to be log-normal with standard deviation σ). It can be observed that in order to accommodate, for example, at least 98% of UE locations, the guard period should be dimensioned in accordance with $d_{\max} \geq \gamma r_0$ where γ is a factor between approximately 1.5 and 3 depending on the degree of shadow fading.




Figure 23.9: Cumulative distribution of serving cell propagation delays relative to nominal cell radius.

Overall, the total guard time T_G at a TDD eNodeB per uplink–downlink cycle is equal to the sum of T_{G1} and T_{G2} , as given by Equation (23.11). In the LTE specifications this is represented by a single guard period amalgamating both parts (with the uplink subframe timing advanced by an amount T_{G2} at the eNodeB with respect to the downlink timing). This is only a matter of representation, however, and the end result is essentially identical to the presence of the two separate guard periods:

$$T_G = T_{G1} + T_{G2} = 2T_{P_{\max}} + T_{UE}^{RxTx} + \max(T_{UE}^{TxRx}, T_{eNB}^{RxTx}) \quad (23.11)$$

By assuming reasonable values for the UE and eNodeB switching times (typically of the order of 10 to 20 μ s) the length of this amalgamated guard period can be dimensioned (in multiples of the OFDM⁴ symbol duration) for a particular deployment. Thus although

⁴Orthogonal Frequency Division Multiplexing.

guard periods represent an undesirable overhead for TDD, by allowing for a flexible and configurable guard period duration, systems can be tailored to the topology of the deployment whilst minimizing the spectral efficiency loss.

The LTE specifications [7] support a set of guard period durations ranging (non-contiguously) from 1 to 10 OFDM symbols for the normal CP (or from 1 to 8 OFDM symbols for the extended CP). A duration of 1 OFDM symbol should be sufficient for many of the anticipated cellular deployments of LTE (up to around 2 km nominal cell radius for $\gamma = 2$), whereas at the other end of the scale, guard period durations of the order of 700 μs support one-way propagation-path delays of the order of 100 km.

The guard period in LTE TDD is located within a mixed uplink/downlink subframe (known as a ‘special subframe’) as shown in Figure 6.2 and further discussed in Section 23.4.2.

23.4.2 Coexistence between Dissimilar Systems

There are scenarios in which an LTE TDD system could need to coexist with other (non-LTE) radio access technologies within the same frequency bands. As mentioned in Section 23.3.2, the ability to time-align the uplink and downlink transmissions between neighbouring systems can be used to avoid base-to-base and UE-to-UE interference paths and hence to alleviate the dependency on the other coexistence measures mentioned above. This also applies to the case of dissimilar neighbouring systems.

As shown in Figure 6.2, the LTE TDD system has a 10 ms radio frame supporting either one pair of switching points per 5 ms, or one pair of switching points per 10 ms, which we denote here as ‘5 ms switching’ and ‘10 ms switching’ respectively.

The provision for both 5 ms and 10 ms switching options in LTE TDD was introduced to enable switching point alignment between LTE and the UTRA TDD modes. Ideally, for maximum flexibility of alignment, the location of the downlink-to-uplink transition within the 5 ms or 10 ms cycle would be fully adjustable with symbol-level granularity, although certain practical considerations must also be taken into account. The overall intention is first to align the location of the uplink-to-downlink transitions between the dissimilar systems by means of a frame-timing offset, and then to adjust the position of the LTE downlink-to-uplink transition such that it is approximately aligned with that of the other system.

The TDD variant of Mobile WiMAX (based upon the IEEE 802.16e amendment) also utilizes 5 ms switching, and can itself accommodate a variable position of the switching point with symbol-level granularity (i.e. with adjustments of approximately $\pm 102.8 \mu\text{s}$). Hence switching point alignment between TDD WiMAX and TDD LTE is possible with sufficient resolution.

The presence of a switching point that is adjustable at the OFDM symbol level results in a mixed subframe, known as the ‘special subframe’, potentially containing both downlink and uplink regions (referred to as ‘DwPTS’ and ‘UpPTS’ respectively⁵) as shown in Figure 23.10 (see also Figure 6.2).

The special subframe is the most significant difference between the FDD and TDD physical layers in LTE, giving rise to a number of details which must be considered in the design. Uplink and downlink transmission durations in this irregular subframe are effectively

⁵DwPTS and UpPTS is terminology inherited from Time-Division Synchronous Code Division Multiple Access (TD-SCDMA) where the terms denoted Downlink Pilot Time Slot and Uplink Pilot Time Slot respectively. This, however, does not well reflect the usage of the fields bearing the same names in LTE.




Figure 23.10: Special subframe for downlink–uplink switching in TDD operation.

reduced compared to a normal subframe, implying that less forward error correction redundancy can be employed for a given transport block size. Alternatively, the transport block size itself can be reduced, although this smaller transport block size must then also be used for HARQ retransmissions which may occur in a regular 1 ms subframe. In general, the use of HARQ with incremental redundancy, as described in Section 10.3.2.5, limits the impact of the fact that the special subframe has less downlink transmission resource than a normal subframe.

For the downlink control signalling on the Physical Downlink Control CHannel (PDCCH), the special subframe does not have a significant effect, as the PDCCH is anyway contained within the first few OFDM symbols of a subframe. Hence PDCCH transmission in the DwPTS region is possible, with the exception that it is constrained in this case to a duration of a maximum of two OFDM symbols (instead of up to three in normal subframes, or even four in cases of narrow system bandwidths, as explained in Section 9.3).

The length of the UpPTS field in the mixed subframe is constrained in the LTE specifications to support only lengths of one and two SC-FDMA symbols. This field does not therefore support uplink data transmission and is instead used only for a shortened random access preamble (suitable only for small cells – see Section 17.6) and for transmission of uplink Sounding Reference Signals (SRSs) – see Section 15.6. One downside of the absence of support for UpPTS lengths other than one or two symbols is that it restricts the set of possible configurations for switching point alignment with other TDD systems. Nonetheless, Figure 23.11 shows some examples of uplink-downlink alignment that are possible between LTE TDD and other TDD systems.

23.4.3 HARQ and Control Signalling for TDD Operation

As explained in Section 10.3.2.5, transmission of downlink or uplink data with HARQ requires that an ACKnowledgement ACK or Negative ACK) be sent in the opposite direction to inform the transmitting side of the success or failure of the packet reception.

In the case of FDD operation, acknowledgement indicators related to data transmission in a subframe k are transmitted in the opposite direction during subframe $k+4$, such that a one-to-one synchronous mapping exists between the instant at which a transport block is transmitted and its corresponding acknowledgement. However, in the case of TDD operation, subframes are designated on a cell-specific basis as uplink or downlink (with the exception of the mixed subframe), thereby constraining the times at which resource grants, data transmissions, acknowledgements and retransmissions can be sent in their respective




Figure 23.11: Examples of switching point alignment between LTE and non-LTE TDD systems.

directions. The synchronous scheme used for FDD cannot therefore be directly reused for TDD operation.

For TDD in asymmetric uplink-downlink cases, it is inevitable that for one of the link directions there are insufficient subframes to support a synchronous one-to-one mapping between transmitted data and its acknowledgement. The LTE design for TDD therefore supports grouped ACK/NACK transmission to carry multiple acknowledgements within one subframe.

For uplink HARQ, the sending (in one downlink subframe) of multiple acknowledgements on the Physical Hybrid ARQ Indicator CHannel (PHICH) is not problematic since, when viewed from the eNodeB, this is not significantly different from the case in which single acknowledgements are sent to multiple simultaneous UEs. However, for downlink HARQ, if the asymmetry is downlink-biased, the uplink control signalling (PUCCH) formats of FDD are insufficient to carry the additional ACK/NACK information. This situation is evident from the example of Figure 23.12 where only 4 uplink sub-frames per radio frame are available to carry acknowledgements for the 6 subframes that may contain downlink data transport blocks. In this particular example, this is solved by means of uplink subframe numbers 2 and 7 each carrying acknowledgement data corresponding to two downlink subframes (numbers 5 and 6, and 0 and 1 respectively). Each of the TDD subframe configurations in LTE (shown in Figure 6.2) has its own such mapping predefined between downlink and uplink subframes for HARQ purposes (see [8, Section 7.3]), with the mapping being designed to achieve a balance between minimization of acknowledgement delay and an even distribution of ACK/NACKs across the available uplink subframes.

Two mechanisms are provided for grouping the acknowledgement information carried in the uplink in TDD operation, termed ‘ACK/NACK bundling’ and ‘ACK/NACK multiplexing’. Selection between these mechanisms is by higher-layer (Radio Resource Control (RRC)) configuration.




Figure 23.12: Grouped ACK/NACK transmission in TDD operation.

ACK/NACK bundling reuses the same 1- and 2-bit Physical Uplink Control Channel (PUCCH) formats ('1a' and '1b' – see Section 16.3) which are used for FDD. For each downlink codeword (up to two if downlink spatial multiplexing is used), only a single acknowledgement indicator is derived by performing a logical 'AND' operation of the acknowledgements across the group of downlink subframes associated with that uplink subframe; this indicates whether zero or more than zero transport blocks in the group were received in error.

For ACK/NACK multiplexing, a separate acknowledgement indicator is returned for each of the associated downlink subframes. However, to limit the amount of signalling information that this would generate, acknowledgements from multiple codewords on different spatial layers within a subframe are first condensed into a single acknowledgement, again by means of a logical 'AND' operation. This is known as 'spatial ACK/NACK bundling'. For the more extreme asymmetries, however, there remains a need to transmit more than two bits of ACK/NACK information in one uplink subframe. This is achieved using the normal 1- and 2-bit PUCCH formats augmented with a code selection scheme whereby the PUCCH code selected by the UE conveys the surplus information to the eNodeB (see [8, Section 10.1]).

A disadvantage of these lossy compression schemes for grouped acknowledgements is that the eNodeB does not know exactly which transport block(s) failed in decoding. In the event of a NACK, all transport blocks in the same group must be resent, increasing retransmission overheads and reducing link throughput. A more subtle impact is that the average HARQ round trip time (and hence latency) can be increased due to the fact that some blocks cannot be acknowledged until the remainder of the group have been received.

A further complication arises because the PDCCH control signalling is not 100% reliable and there is some possibility that the UE will miss some downlink resource assignments. This would introduce the possibility of HARQ protocol errors, including the erroneous transmission of ACK in the case when one or more downlink assignments were missed in a group of subframes. In order to help avoid this problem, a 'Downlink Assignment Index' (DAI) is included in the PDCCH to communicate to the UE the number of subframes in a group that actually contain a downlink transmission. In the case of ACK/NACK bundling, this helps the UE to detect missed downlink assignments and avoid returning ACK if one or more downlink assignments were missed, while in the case of ACK/NACK multiplexing the DAI

helps the UE to determine how many bits of ACK/NACK information should be returned. Nevertheless, this mechanism alone cannot safeguard against all possible error cases.

In Release 10, where even larger numbers of ACK/NACK bits may need to be transmitted in a single subframe due to carrier aggregation, new PUCCH mechanisms are provided, as explained in Section 28.3.2.

The presence of link asymmetry in TDD generally increases the overall HARQ round trip times (compared to the 8 ms of FDD as discussed in Section 10.3.2.5) as a result of the additional ‘waiting times’ needed for the appropriate link direction to become available for transmission of an acknowledgement, grant or retransmission. This, together with the degree of asymmetry, also means that the number of HARQ processes in TDD varies depending on the TDD uplink-downlink subframe configuration: in the downlink, the number of HARQ processes varies from 4 for configuration 0 to 15 for configuration 5, while in the uplink it varies from 1 for configuration 5 to 7 for configuration 0 – see [8, Tables 7-1 and 8-1].

The signalling of the granted uplink transmission resources also requires some specific attention for TDD. For FDD, the downlink subframe in which an uplink grant is sent implicitly also signals the specific uplink subframe that has been assigned (located four subframes later), whereas for TDD this relationship cannot always hold due to the various uplink-downlink configurations. An alternative linkage is therefore formulated for each specific uplink-downlink configuration to associate each uplink subframe with one preceding downlink subframe that controls it (maintaining the same four-subframe spacing as FDD wherever possible). For configuration 0 (shown in Figure 6.2), there are more uplink subframes per 10 ms than the number of subframes available for PDCCH, and here an additional ‘UL Index’ field in the uplink grant message is used to indicate the uplink subframe to which the grant relates. In addition, one value of the ‘UL Index’ field can be used to enable a single PDCCH message to grant uplink resources in two uplink subframes (see [8, Section 8.0]).⁶

23.4.4 Half-Duplex FDD (HD-FDD) Physical Layer Operation

In principle, one of two paths could be taken for the physical layer design of a HD-FDD system:

- a derivative of the FDD system, in which the scheduling is arranged such that the terminals are not required to simultaneously transmit and receive;
- a derivative of the TDD system, in which the uplink and downlink happen to reside on different carrier frequencies.

The scheme selected for HD-FDD operation in LTE is in accordance with the FDD derivative. This has helped to maintain the desirable commonality with FDD and has reduced the overall complexity of the solution. Unless otherwise informed, a UE in the FDD-derived HD-FDD scheme has no a priori knowledge of the uplink-downlink transmission pattern. Instead the UE checks any subframe which has not otherwise been pre-assigned to uplink transmission for the presence of PDCCH control signalling addressed towards it. The eNodeB is of course aware of any uplink transmission grants it has sent using the PDCCH, and does

⁶Note, however, that a UE is not required to receive more than one uplink grant message on the PDCCH in a single downlink subframe.

not therefore expect the UE to be able to receive downlink transmissions in the corresponding uplink subframe(s).

In this scheme, the fixed one-to-one association between downlink and uplink subframes (arising from the timing relationships between resource grants, data transmission and HARQ acknowledgements) is retained as in normal FDD operation. It is worth noting, however, that this approach results in the situation that, for a given user, no more than 50% of subframes may be used for any one link direction for a given UE.

Nonetheless, the perceived impacts to a user in anything other than an unloaded system are likely to remain small. The instantaneous (i.e. in one subframe) peak data rate is not affected and, in any normally loaded or even partially loaded system, the eNodeB is in any case unable to dedicate its full time resources to a user for any lasting period due to its need to service other users. Furthermore, at least for the downlink, scheduling strategies may be adapted to increase the amount of frequency resource allocated to a HD-FDD UE at each scheduling instant in order to alleviate reliance on frequent transmissions in the time domain. For the uplink this is also possible except for situations in which the UE is transmission-power-limited. In such cases, the application of wider bandwidths may not allow for increased instantaneous data rates and therefore is also unlikely to allow for a reduction in the fraction of time-domain resources needed to achieve a targeted aggregate uplink rate.

In order to allow sufficient time for downlink-to-uplink switching in HD-FDD operation, the UE is not expected to be able to receive the last symbol(s) of a downlink subframe that precedes an uplink subframe in which the UE is active. The length of the switching period is very similar to that previously given in Equation (23.11), with the exception that $T_{eNB}^{RxTx} = 0$ and the cell-specific maximum propagation delay value $T_{p,max}$ may be replaced with a value T_p that is applicable to the particular UE in question. Thus for half-duplex FDD:

$$T_G = T_{G1} + T_{G2} = 2T_p + T_{UE}^{RxTx} + T_{UE}^{TxRx} \quad (23.12)$$

During this time, the eNodeB may or may not continue to transmit the remainder of the data towards the UE. In order to simplify the system and to minimize the differences from full-duplex FDD the transport channel coding structure continues to output the same number of channel-coded bits as for a normal-length subframe. Correct reception of the data in spite of the UE's lack of reception during the switching period is therefore reliant on the presence of sufficient forward error correction. For larger switching times the effects on performance obviously become more pronounced due to the increasing code rate, and in these instances the scheduler could reduce the transport block size mapped to the downlink subframe containing the switching period.

Due to the fact that the eNodeB is full-duplex, cell-specific Reference Signals (RSs), common signalling and any other data directed towards other UEs may also continue to be transmitted by the eNodeB throughout a particular UE's switching period and uplink subframes; that is to say, the switching periods and designated uplink subframes are UE-specific in the case of HD-FDD, rather than cell-specific as is the case for TDD.

23.5 Reciprocity

Reciprocity is a general phenomenon encountered in wave propagation over linear media. The details of the underlying theory are not reiterated here, and the interested reader is

referred to the comprehensive treatment which can be found, for example in [9–12]. We cover here only the practical implications for cellular radio communication systems, and more specifically those for LTE TDD.

The efficiency of an antenna in converting between electrical and electromagnetic energy is equal in both directions at a given frequency. The result is that the directional sensitivity of an antenna (when acting as a receiver) is also equal to its far-field radiated pattern (when acting as a transmitter).

For a simple source and a simple receiver in a linear medium separated by sufficient distance such that one does not affect the load of the other, the reciprocity theorem states that the locations of the source and receiver may be interchanged yet the transfer function between them remains unchanged. Therefore when two transceivers A and B each utilize the same antenna for transmission and reception, the frequency-domain transfer function $H(f)$ (or equivalently the time-domain impulse response $h(t)$) observed at B when A is transmitting will be the same as the transfer function observed at A when B is transmitting. This is illustrated in Figure 23.13.




Figure 23.13: Reciprocity between a pair of transceivers.

This holds for any pair of antennas in each other’s far field, and hence is easily generalized to the MIMO case in which each transceiver has multiple antennas ($i = 1, \dots, I$ for transceiver A and $j = 1, \dots, J$ for transceiver B):

$$H_{ij}(f) = H_{ji}(f) \tag{23.13}$$

Considering for simplicity a particular frequency k , and denoting the MIMO channel matrix between A and B at frequency k as $H_{k,AB}$ where $H_{k,AB}(i, j) = H_{ij}(k)$, it is evident that

$$H_{k,AB} = H_{k,BA}^T \tag{23.14}$$

In order to exploit reciprocity in a radio communications link, the set of frequencies used for transmission and reception at each transceiver must usually overlap, as is the case for TDD systems. Short-term reciprocity cannot be exploited for FDD systems unless the nature of each constituent channel $H_{ij}(f)$ is frequency-invariant over a range of frequencies spanning both distinct transmit and receive bands, which is rarely the case. However, long-term reciprocity of the channel covariance matrix can sometimes still be exploited, as discussed in Section 23.5.2.7.

23.5.1 Conditions for Reciprocity

Reciprocity can be used to control the characteristics of an outbound transmission in anticipation of the channel response that the transmission will experience. In these instances, knowledge of the outbound channel is inferred at the transmitter from measurement of the inbound channel responses. For such techniques to be successful, however, one must first consider the following:

- the topology or configuration of the MIMO antenna system;
- the rate of change in the mobile radio channel;
- calibration of the transmitters and receivers involved.

23.5.1.1 Antenna Configuration

Equation (23.14) holds for any number of antennas at each end of the radio link, yet it may be the case that not every antenna element is used for both the transmit and receive functions at a given transceiver. For example, implementation of a transmit chain (and associated power amplifier) per receive antenna element may be feasible at an eNodeB whereas it may not be so in a UE for reasons of form-factor, cost and power consumption.

In these cases, knowledge of the Channel State Information (CSI) of the outbound channel is incomplete if obtained only using inbound RSs. Additional measures would then need to be adopted to determine the missing CSI components. These could include the use of supplementary and explicit feedback of the CSI from the remote receiver for any of its receive antennas that do not have transmit capability. Alternatively, the inbound MIMO channel may be sounded using RSs sent for example from only one transmit antenna at a time. This can then allow for a UE with only one transmit chain and power amplifier to sound the channel from all of its receive antennas by means of a time-controlled switch.

23.5.1.2 Time Variations in the Radio Channel

Mobile radio channels are of course time-varying in nature, as discussed in Chapter 8. For TDD systems, uplink and downlink transmissions occur at different times, thereby violating the reciprocity principle in a time-varying channel. However, the channels can still be considered as approximately reciprocal when viewed over a time span short enough to preclude any appreciable change. The duration between the instant at which the CSI is measured and the instant at which it is applied thus becomes of paramount importance.⁷

⁷Note that this applies also to systems using closed-loop feedback signalling, and not only to those attempting to exploit uplink and downlink channel reciprocity. It is simply the case that for control schemes reliant upon reciprocity the CSI is measured at the output of the inbound link rather than at the output of the outbound link.

The time over which a channel’s impulse response $h(t)$ remains relatively constant is termed its *coherence time* and is defined as the delay δt at which the squared magnitude of its autocorrelation first drops below a certain fraction of its value at zero delay (a fraction of 50% is often used to define the coherence time).

Clearly in order to benefit from channel reciprocity, the delay between the time at which the observation of the inbound channel is made and the time at which this knowledge is applied to the outbound signal, must be shorter than the coherence time of the channel. Typical relationships between mobile speed and coherence time (in this instance for a carrier frequency of 2 GHz) are shown in Figure 23.14.




Figure 23.14: Channel coherence times for flat Rayleigh fading at 2 GHz.

From a practical perspective, the processes to be executed within the channel coherence time include receiving the necessary RSs, estimating the CSI, computing the required action and encoding and processing the outbound signals. With present-day technology this turnaround time is unlikely to be much lower than 1 ms. Even if it were, the LTE TDD frame structure would need to support uplink–downlink transitions at this rate to ensure continual updating of the CSI on the inbound link (interspersed by transmission on the outbound link); however, due to the minimum switching period of 5 ms, this is not possible. The CSI cannot be applied earlier than the next occurrence of a subframe in the return direction. It may also be the case that a particular UE will not be scheduled in the first available subframe of the

return direction, potentially further increasing the delay between measurement of the CSI and its application.

The RSs commonly used for exploiting reciprocity in LTE are the cell-specific RSs in the downlink⁸ and the Sounding RSs (SRSs) in the uplink. The cell-specific downlink RSs are present in each downlink subframe as shown in Section 8.2.1, while the SRSs are located either in the short UpPTS field of the mixed subframe (see Figure 23.10) or on one symbol of a normal uplink subframe where configured.

Under the assumption that one would prefer to use UpPTS for sounding in TDD, then with a 5 ms switching cycle and a balanced DL/UL configuration (TDD configuration 1 in Figure 6.2), the shortest possible turnaround times to exploit reciprocity for the two downlink subframes and the DwPTS would be 2 ms, 3 ms and 4 ms respectively, as shown in Figure 23.15(a). These correspond to mobile speeds of up to around 24 km/h at a 50% coherence level at 2 GHz.

For the uplink, the ability of the system to take advantage of reciprocity is not limited by the occurrence interval of the wideband cell-specific RSs (as they are present in every downlink subframe), but primarily by the length of the contiguous uplink period within the frame and by the turnaround processing time at the UE. Figure 23.15(b) gives one example in which the processing delay at the UE is assumed to be of the order of 2 ms, resulting in similar turnaround latencies to the downlink case. With faster processing at the UE, these could be reduced a little further.

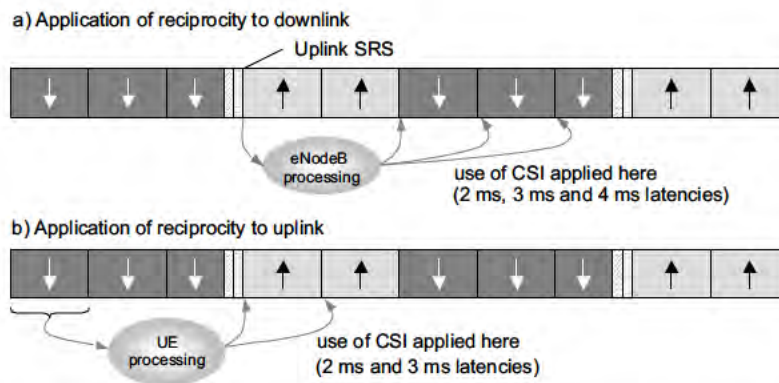


Figure 23.15: Example turnaround times for exploitation of reciprocity in LTE TDD:
(a) application of reciprocity to downlink; (b) application of reciprocity to uplink.

23.5.1.3 Calibration

Under conditions of channel reciprocity, the forward and reverse responses between the antennas at two transceivers A and B are the same, although this does not mean that the overall responses between the baseband processors at those transceivers are also equivalent

⁸In Release 10, the CSI-RS (see Section 29.1.2) may also be used.

in both directions. The responses of the individual transmitters and receivers themselves must also be taken into account, as shown for a 2×2 system transmitting from baseband A to B in Figure 23.16.




Figure 23.16: Transceiver elements in a reciprocal channel.

At a given frequency the responses of the multiple transceiver elements may be represented together in matrix form (assuming here that variations within the desired system bandwidth are small). Without coupling effects between elements, these are simply diagonal matrices containing the per-element complex gains, such that for I antennas at A and J antennas at B:

$$G_{Tx,A} = \begin{pmatrix} g'_{A,1} & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & g'_{A,I} \end{pmatrix}, G_{Rx,B} = \begin{pmatrix} g'_{B,1} & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & g'_{B,J} \end{pmatrix} \quad (23.15)$$

A similar formulation may be applied to $G_{Tx,B}$ and $G_{Rx,A}$ in the reverse direction assuming each antenna is capable of both transmit and receive. The composite transfer functions between A and B (and vice versa) at a given frequency then become

$$Z_{AB} = G_{Tx,A} H_{AB} G_{Rx,B} \quad \text{and} \quad Z_{BA} = G_{Tx,B} H_{BA} G_{Rx,A} \quad (23.16)$$

It is clear that Z_{AB} and Z_{BA} are not necessarily equal. Nonetheless, if the time between the measurements \hat{Z}_{AB} and \hat{Z}_{BA} of Z_{AB} and Z_{BA} respectively is sufficiently smaller than the coherence time, then using Equation (23.14) it can be deduced that (for invertible G matrices):

$$\hat{Z}_{AB} = G_{Tx,A} (G_{Rx,A}^T)^{-1} \hat{Z}_{BA}^T (G_{Tx,B}^T)^{-1} G_{Rx,B} \quad (23.17)$$

The general goal therefore of a system wishing to exploit reciprocity is to infer \hat{Z}_{AB} from the observation \hat{Z}_{BA} using knowledge of the relevant transceiver gains in G .

There are numerous methods by which this can, in theory, be achieved. One such method is via an explicit signalling exchange between transceivers of the measured composite channel responses (\hat{Z}_{AB} and \hat{Z}_{BA}), as observed within the coherence time of the channel. By analysing the relation between these, an active transmitter may then derive the necessary transfer function needed to correctly infer the composite outbound channel from future measurements of the composite inbound channels [13]. Note that the update rate of such feedback to track variations in the transceiver responses should be relatively slow – otherwise a pure feedback-based approach that is not reliant on reciprocity becomes a more natural choice. Fortunately, the variations in transceiver responses are predominantly driven by temperature changes, and

therefore a low update rate is usually possible (although other factors can also arise, such as interaction between the antenna and nearby objects such as human hands at the UE side). In theory, self-calibration methods are also possible whereby the transceivers make use of accurate reference transmitters or receivers, or of auxiliary calibration signals.

Calibration is not generally required for link control systems using measured CSI from the output of the outbound link (i.e. those not reliant upon reciprocity, including both feedback-based and certain feed-forward approaches⁹). In these cases, the aggregated transmitter and receiver responses for the link under control already form an integral part of the measured CSI and hence any variations are accommodated as a natural part of the ongoing link control process without the need for specific calibration procedures.

One noteworthy exception to the above, however, is the case of beamforming in which the ability of an array to steer energy towards (or collect energy from) a given direction is impaired if the phase between the elements is not adequately controlled. In this instance calibration would generally only be required at the transceiver attached to the array (such as at the eNodeB); an opposing transceiver without an array (such as a mobile terminal) would not require calibration.

In general therefore, it can be said that systems attempting to utilize channel reciprocity require calibration of their constituent transceivers in order to restore the overall reciprocity between them at baseband. Calibration is also required for antenna arrays reliant upon directionality, such as transmit and receive beamforming, although this applies irrespective of the presence of channel reciprocity.

23.5.2 Applications of Reciprocity

For the purposes of this discussion, we define the ‘outbound link’ as the link which is being actively controlled in response to the channel.

Active control techniques may be generally classified as feed-forward (open-loop) or feed-back (closed-loop). Note that both may involve feedback of information from the receiving side to the transmitting side, the distinction lying in the particular nature of this information. Receivers in closed-loop feedback systems return information regarding how well the transmitting side is performing in relation to a target. Conversely, feed-forward transmitters do not know how well their transmissions are performing, and simply use whatever information is available (via feed-back information or otherwise) to try to pre-empt one or more aspects of the radio channel.

There also exists a class of passive open-loop systems in which the transmitting side takes no pre-emptive action, and exploitation of the channel is confined solely to the receiving side. Figure 23.17 shows this classification of link control methods.

For active feed-forward control systems, the information required by the transmitter often takes the form of CSI which supplies information regarding the response between transmitter and receiver. It is really only this class of system for which reciprocity may be used – the outbound link CSI being inferred from the inbound link CSI given that the conditions discussed in Section 23.5.1 are met and that inbound RSs covering the desired range of frequencies are available for this purpose.

⁹Feed-back and feed-forward control mechanisms are discussed further in Section 23.5.2.




Figure 23.17: Feed-forward, feed-back and open-loop exploitation of the radio channel.

However, for wireless communication, it is often not only the channel between the controlling transmitter and the receiver that is of interest; information relating to interference caused by other transmitters is also relevant and cannot be obtained via reciprocity.

Some link control methods make use of only amplitude information from the channel (e.g. open-loop uplink power control and frequency-domain scheduling), whereas others, such as MIMO precoding, also need phase information. Even in a TDD deployment, it cannot always be assumed that reciprocity can be freely exploited to achieve this. In this regard, we discuss below some link control examples of relevance in a TDD LTE system.

23.5.2.1 Downlink Frequency-Domain Scheduling

The eNodeB scheduler can exploit the frequency-selective nature of the channel by attempting to schedule downlink data to a user using frequency resources that are currently experiencing higher than average SINR (see Chapter 12). Although reciprocity allows for the frequency response of the downlink radio channel to be inferred via the observed response of the uplink channel, it does not provide information concerning the downlink interference and noise levels observed at the UE. Hence knowledge of the channel provides only partial knowledge of the SINR.

23.5.2.2 Uplink Frequency-Domain Scheduling

As for the downlink case, this is also an interference-responsive mechanism. Contrary to the downlink, however, the eNodeB in this instance does have some knowledge of the interference and noise levels affecting the uplink transmission at its own receiver. Together

with the required CSI of the uplink channel gleaned from the SRS transmissions, this provides the eNodeB with all the information required. However, at no point is reciprocity required in order for this process to function, due to the fact that all involved signals exist only in one link direction.

23.5.2.3 Downlink Power Control and Link Adaptation

This function requires interference feedback from the UE if it is to operate effectively. Reciprocity can be used by the eNodeB to infer the downlink CSI, and thereby to provide partial information, but pure feedback-based techniques tend to be better suited to this application.

23.5.2.4 Uplink Power Control and Link Adaptation

In theory, reciprocity could be used in the feed-forward sense to adjust the UE's transmit power and coding/modulation. In this instance, the fact that the interference levels may be controlled by the eNodeB means that the received SINR can be better correlated with the channel response, thereby improving the usefulness of reciprocity.

For uplink power control, the LTE system only makes use of long-term reciprocity via the downlink path-loss estimation (see Section 18.3.2.1). For uplink link adaptation, autonomous selection of the modulation and coding rate by the UE is deliberately not permitted, in order to enable the eNodeB to exercise full control of the uplink and its optimization.

23.5.2.5 Phase-Passive MIMO

Phase-passive MIMO schemes are those in which the transmitter does not attempt to adapt to the phase of the channel and responds only to amplitude information. Two techniques in LTE that fall into this category are open-loop downlink MIMO (i.e. transmission mode 3 not using Precoding Matrix Indicator (PMI) feedback – see Section 9.2.2.1) with per-stream Modulation and Coding Scheme (MCS) selection, and switched transmit antenna diversity for the uplink. The former essentially comprises parallel instances of downlink rate control and hence the discussion of Section 23.5.2.3 is applicable.

LTE supports switched antenna diversity for the uplink in one of two configurable modes. In the first mode, the eNodeB selects the transmit antenna that the UE should use as described in Section 16.6.1 and signals this decision to the UE. This mode does not require reciprocity in order to function. In the second mode, the UE is allowed to select the transmit antenna for uplink transmission, and in TDD operation this can therefore be based on reciprocity using for example the cell-specific downlink RSs to provide the UE with the necessary information on which to base its decision.

23.5.2.6 Short-Term MIMO Precoding

As discussed in Chapter 11, precoding refers to the application of a set of antenna weights per transmitted spatial layer (the precoding matrix) which, when applied to a forthcoming transmission, interact beneficially with the radio channel in terms of improved reception or separation of the layers at the intended receiver(s). Both feed-back and feed-forward techniques are possible in terms of precoding matrix selection strategy. Feed-back approaches

(using PMI in LTE) do not rely on reciprocity, only on the feedback loop delay being shorter than the channel coherence time. For feed-forward approaches, the necessary CSI can, in theory, be obtained in one of two ways:

1. using ‘Direct Channel Feed-Back’ (DCFB), in which the CSI observed by the receiver is explicitly signalled back to the transmitter;
2. using channel sounding RSs (e.g. SRS on uplink) to sound the link inbound to the precoding transmitter.

DCFB obviates the need for reciprocity, whereas the RS-assisted method is reliant upon it. LTE does not support DCFB, and hence reciprocity and TDD are prerequisites for feed-forward-based short-term precoding, whereas the feedback-based (PMI) scheme is applicable to all duplex modes.

As described in Chapter 11, explicit control signalling is required for the PMI-based scheme, whereas for a reciprocity-based scheme the role of control signalling is undertaken by UE-specific RSs in the downlink (see Section 8.2.2) and SRS in the uplink.

A reciprocity-based scheme carries the advantage that the applied precoding matrix is not constrained to a finite set of possibilities¹⁰ (a codebook). However, the support for UE-specific RSs was limited in the first release of LTE (Release 8) to allow only a single spatial layer,¹¹ thereby precluding the use of spatial multiplexing using non-codebook-based precoding. In Release 9, the UE-specific RS design was extended to support two layers (see Sections 8.2.3, 11.2.2.3) thereby also enabling dual-layer short-term precoding using techniques based on reciprocity. In Release 10, the use of UE-specific RSs is further extended to up to eight layers, as explained in Section 29.1.1.

23.5.2.7 Long-Term MIMO Precoding (Beamforming)

At a beamforming transmitter, energy is focused in the directions which are more commonly observed in the inbound channel than others, assuming that a similar angular distribution also applies in the mean sense to the outbound channel. Beamforming is reliant, therefore, only on longer-term statistics of the radio channel and does not require short-term correlation between the inbound and outbound channels in order to function. As such, this scheme is applicable to both FDD and TDD; short-term channel reciprocity in the sense described above is not required.

23.5.3 Summary of Reciprocity Considerations

Exploitation of channel reciprocity for LTE TDD is possible in certain areas. These include its application to uplink transmit antenna selection (using only amplitude information) and to non-codebook-based downlink precoding (also using phase information). Support for the latter has been greatly improved since the first release of LTE via the introduction of multi-layer support for UE-specific RSs. The benefits arising from the use of reciprocity are generally dependent upon the intended deployment scenario, and in some cases pure feedback-based approaches are sufficient.

¹⁰Referred to as non-codebook-based precoding, this can potentially improve performance by removing mismatch between the best-fit precoding matrix in the codebook and the precoding matrix optimal for the channel.

¹¹The initially intended application of UE-specific RSs was to support single-layer beamforming.

In order to realize a true state of reciprocity between a base station and a UE, the feed-forward loop delay must be within the coherence time of the channel. Calibration of the transceivers is also required (notably when exploiting phase information), and the antenna elements involved must generally be shared by both transmit and receive functions.

References¹²

- [1] Ofcom, ‘Spectrum Framework Review: Implementation Plan’, Ofcom consultation document, www.ofcom.org.uk, January 2005.
- [2] 3GPP Technical Report 36.942, ‘Evolved Universal Terrestrial Radio Access (E-UTRA); Radio Frequency (RF) System Scenarios (Release 8)’, www.3gpp.org.
- [3] T. Wilkinson and P. Howard, ‘The Practical Realities of UTRA TDD and FDD Co-Existence and their Impact on the Future Spectrum’, in *Proc. 15th IEEE International Symposium Personal and Indoor Mobile Radio Communications (PIMRC 2004)* (Barcelona, Spain), September 2004.
- [4] 3GPP Technical Report 25.814, ‘Physical Layer Aspects for Evolved Universal Terrestrial Radio Access (UTRA)’, www.3gpp.org.
- [5] J. H. Winters, ‘Optimum Combining in Digital Mobile Radio with Cochannel Interference’. *IEEE Journal on Selected Areas in Communications*, Vol. 2, pp. 528–539, July 1984.
- [6] E. G. Larsson, ‘Model-Averaged Interference Rejection Combining’. *IEEE Trans. on Communications*, Vol. 55, pp. 271–274, February 2007.
- [7] 3GPP Technical Specification 36.211, ‘Evolved Universal Terrestrial Radio Access (E-UTRA); Physical Channels and Modulation’, www.3gpp.org.
- [8] 3GPP Technical Specification 36.213, ‘Evolved Universal Terrestrial Radio Access (E-UTRA); Physical Layer Procedures’, www.3gpp.org.
- [9] J. W. S. Rayleigh, *The Theory of Sound* [Macmillan & Co., 1877 and 1878]. Reprinted by Dover, New York, 1945.
- [10] H. A. Lorentz, ‘Het theorema van Poynting over de energie in het electromagnetisch veld en een paar algemeene stellingen over de voortplanting van licht’. *Verhandelingen en bijdragen uitgegeven door de Afdeling Natuurkunde, Koninklijke Nederlandse Akademie van Wetenschappen te Amsterdam*, Vol. 4, pp. 176–187, 1895–1896.
- [11] J. R. Carson, ‘Reciprocal Theorems in Radio Communication’ in *Proc. Institute of Radio Engineers*, June 1929.
- [12] S. Ballantine, ‘Reciprocity in Electromagnetic, Mechanical, Acoustical, and Interconnected Systems’, in *Proc. Institute of Radio Engineers*, June 1929.
- [13] M. Guillaud, D. Slock and R. Knopp, ‘A Practical Method for Wireless Channel Reciprocity Exploitation through Relative Calibration’, in *Proc. Eighth International Symposium on Signal Processing and Its Applications*, August 2005.

¹²All web sites confirmed 1st March 2011.

Picocells, Femtocells and Home eNodeBs

Philippe Godin and Nick Whinnett

24.1 Introduction

The use of small cells is becoming increasingly important due to their ability to provide increased system capacity compared to a homogeneous network of macrocells.¹ Small cells can generally be characterized as either *picocells* (also referred to as hotzone cells), controlled by a pico eNodeB, or *femtocells*, controlled by a Home eNB (HeNB).

The definitions of picocells and femtocells are somewhat variable, but typically the main differentiating features can be summarized as follows:

- A pico eNodeB usually controls multiple small cells which are planned by the network operator in a similar way as the macrocells. Picocells are often mounted at low elevation and operate at lower power than the macrocells. A HeNB controls only one cell and is deployed by the customer (the registered owner), usually without planning.
- Femtocells are typically Closed Subscriber Group (CSG) cells accessible only to a limited group of users; in contrast, picocells are usually open to all users (Open Subscriber Group (OSG)) but may offer preferential treatment to some users, e.g. to the staff of a particular establishment.
- The transmission power of femtocells is lower, being designed typically to cover a house or apartment. Picocells usually operate with a higher transmission power to cover an enterprise, mall or other hotzone, or simply to extend macrocellular coverage.

¹A combined network of macrocells and small cells is often known as a 'heterogeneous network'.

- Femtocells do not necessarily have the same network interfaces as macro eNodeBs, whereas pico eNodeBs follow the same logical architecture principles as macro eNodeBs as described in Chapter 2.

Small cells may be deployed on the same carrier as the macrocells or on a dedicated carrier. If a dedicated carrier is used it would typically be at a higher frequency to deliver high data rates within a limited area without causing excessive interference. If the small cells share a carrier with the macrocells, careful attention needs to be given to the interference that may arise between the small cells and the macrocells. This is discussed in Section 24.3.

This chapter first addresses the architectural aspects specific to femtocells/HeNBs. Then the potential interference issues are analysed and some possible interference management schemes are discussed. Finally some of the Radio Frequency (RF) challenges of both HeNBs and pico eNodeBs are discussed.

24.2 Home eNodeB Architecture

The main architectural aspects of HeNBs are based on the general concepts applicable to all types of LTE cell as described in Chapter 2. The following sections describe the aspects specific to HeNBs.

24.2.1 Architecture Overview

The architecture of a network involving HeNBs has been designed to be flexible and scalable with respect to the number of HeNBs. In particular, if the network contains a large number of HeNBs, a HeNB Gateway (HeNB GTW) can optionally be deployed to manage the HeNBs from the perspective of the Evolved Packet Core (EPC). The capacity of a HeNB GTW can be up to several tens of thousands of HeNBs. A HeNB GTW can be deployed on the Control Plane (HeNB GTW CP) only, or it may also include a User Plane (HeNB GTW UP) part. The logical architecture is shown in Figure 24.1.




Figure 24.1: E-UTRAN HeNB logical architecture. Reproduced by permission of © 3GPP.

When a HeNB GTW is deployed, it serves as a concentrator for the CP S1-MME interface (see Section 2.3). Each HeNB therefore has only one Stream Control Transmission Protocol

(SCTP, see [1] and Section 2.5.1) association towards the HeNB GTW and each Mobility Management Entity (MME) likewise has only one SCTP association towards the HeNB GTW. This significantly reduces the overall number of SCTP associations needed in a pool area.² For example if the number of MME nodes is M in the pool area and the number of HeNBs is N , then only $N + M$ SCTP associations will be necessary instead of NM (assuming only one HeNB GTW is used in the pool area).

The S1-U interface may also optionally be terminated in the HeNB GTW UP. In this case the UP connection for one E-UTRAN Radio Access Bearer (E-RAB)³ of one UE consists of two GPRS Tunnelling Protocol (GTP) tunnels instead of one – i.e. one from the Serving-GW (S-GW) to the HeNB GTW UP and one from the HeNB GTW UP down to the HeNB. This option may, for example, be used to avoid the S-GW having to manage paths to tens of thousands of HeNBs. Figure 24.1 also shows the Security Gateway (SeGW) which is optional and is not necessarily part of the HeNB GTW.

The HeNB GTW is assumed to be *transparent* for the S1 interface – i.e. the HeNB sees the HeNB GTW exactly as if it were an MME and conversely the MME sees the HeNB GTW as a regular eNodeB. In terms of protocol this also means that the same S1-MME interface (and the same CP protocol, S1 Application Protocol – S1-AP) is used between the HeNB GTW and the EPC, between the HeNB and the HeNB GTW, between the HeNB and the EPC, and between the eNB and the EPC.

Despite its transparency, the HeNB GTW when deployed ensures the essential function of the S1-flex⁴ in a pool area on behalf of the HeNBs. This means that a particular HeNB only sees one MME (in fact the HeNB GTW) and does not need to implement the S1-flex, while the HeNB GTW is in charge of selecting, on a per-UE basis when a UE attaches⁵ to the network, which one of the MMEs in the pool will handle the UE's requests. This also means that when a HeNB is connected to an HeNB GTW, it will not simultaneously connect to another HeNB GTW or another MME. However, when no HeNB GTW is deployed, the HeNB can still use the S1-flex like any regular eNodeB.

24.2.2 Functionalities

When a HeNB is directly connected to the EPC, it supports exactly the same functions as a normal eNodeB. However, an HeNB is limited to the support of a single cell and it may support an X2 interface only from Release 10 onwards. When the HeNB is connected via the intermediate HeNB GTW it cannot support the Non-Access Stratum (NAS) node selection function used in S1-flex for the reasons explained above, but it does support some additional functions such as initial discovery of a suitable HeNB GTW.

HeNBs can in principle operate in any of three different access modes:

- **Closed access mode:** The HeNB belongs to one or more specific Closed Subscriber Groups (CSGs) identified by CSG Identifiers (CSG IDs). Only UEs which have the corresponding CSG ID included in their CSG subscription list are allowed access. This is the most common mode for HeNBs.

²The pool area is the common area served by a set of MME/S-GWs – see Section 2.2.2.

³See Section 2.4.

⁴See Section 2.2.2.

⁵See Section 2.2.1.1.

- **Hybrid access mode:** Like for closed access mode, the HeNB belongs to a particular CSG. It provides service to all UEs but gives preferential treatment to UEs which include the corresponding CSG ID in their CSG subscription list; such UEs are called ‘CSG members’.
- **Open access mode:** The HeNB behaves as a regular eNodeB.

When a HeNB operates in closed or hybrid access modes, during incoming handovers it has to check the validity of the CSG ID which has been used by the MME for the Access Control or Membership Verification of the user.

The HeNB GTW hosts two main functions:

- It transparently relays the UE-associated S1-AP messages between the MME and the HeNB which serves the UE. All protocol functions associated with a UE-dedicated S1 procedure resides within the HeNB and the MME only.
- It runs the non-UE associated S1-AP procedures between it and the HeNB on one side and between it and the EPC on the other side, and interfaces between them.

In addition, the HeNB GTW can perform some optional functions such as terminating the S1-U interface (see Sections 24.2.1 and 2.2.2) and hence relaying the user plane data between the HeNB and the S-GW, and paging optimization based on a list of CSG IDs provided by the MME in the PAGING message (avoiding paging a UE in closed access mode HeNB cells whose CSG ID is not in the list).

The MME performs Access Control for UEs accessing or handing over to a closed access HeNB, and Membership Verification in the case of hybrid access HeNBs – i.e. the MME checks that the CSG ID of the closed or hybrid HeNB is in the list of allowed CSG IDs for the UE; for a closed access HeNB, the MME rejects the UE if this check is unsuccessful while for a hybrid access HeNB it indicates to the HeNB whether the user is a ‘CSG member’ or not in order to allow any special treatment to be triggered. The MME also controls the routing of handover messages towards HeNBs, either using dedicated HeNB addresses or, if the HeNB is connected through an HeNB GTW, using a ‘Tracking Area Identity’ (TAI) (see [2, Section 19.4.2.3]). The MME may also optionally support paging optimization.

24.2.3 Mobility

Several kinds of mobility are supported via the MME regardless of whether the HeNB is connected via a HeNB GTW:

- Mobility from HeNB to an eNodeB, i.e. outbound mobility, supported from Release 8 onwards;
- Mobility from eNodeB to HeNB, i.e. inbound mobility from macrocell, via S1 handover, supported from Release 9 onwards;
- Mobility from HeNB to HeNB, i.e. inbound mobility from an HeNB cell, via S1 handover, supported from Release 9 onwards;
- Optimized mobility from HeNB to HeNB, i.e. inbound mobility from an HeNB cell via X2 handover, supported from Release 10 onwards.

The HeNB GTW should preferably connect to the EPC in such a way that inbound and outbound mobility to cells served by the HeNB GTW do not necessarily require inter-MME handovers. An example of the general call flow for an inbound mobility via S1 handover is shown in Figure 24.2. The messages may pass through the HeNB GTW (if present) without any functional change.




Figure 24.2: A typical call flow for mobility towards a HeNB. Reproduced by permission of © 3GPP.

The source node for the handover indicates to the MME the ECGI,⁶ CSG ID and access mode of the target HeNB. These elements are generally reported by UEs which are able to read the System Information (SI) of the target cell, but may be otherwise inferred by the source node if it can resolve any Physical Cell Identifier (PCI) confusion. The MME performs the Access Control or Membership Verification based on the received information and the CSG subscription data of the UE. If the UE is not a member of the CSG and the target HeNB is in closed access mode, the MME ends the handover procedure. Otherwise the MME sends the ‘HANDOVER REQUEST’ message (see Section 2.6) to the target HeNB including the value of the CSG ID. If the target HeNB is configured for hybrid access, the MME also indicates in the ‘HANDOVER REQUEST’ message whether the UE is a CSG member or not. The target HeNB further checks whether the received CSG ID is correct and allocates the appropriate resources. The HeNB sends back the ‘HANDOVER REQUEST ACKNOWLEDGE’ message to complete the handover.

If the received CSG ID is found to be incorrect by the target HeNB, the incoming handover will fail if the HeNB is in closed access mode, while in the case of hybrid access mode the HeNB will accept the incoming handover but report back to the MME (in the ‘HANDOVER REQUEST ACKNOWLEDGE’ message) the valid CSG ID so that the MME can apply accurate charging.

⁶E-UTRAN Cell Global Identifier. See Section 22.2.2 for reporting requirements.

In Release 10, mobility between two HeNBs does not necessarily need to use S1 handover via the MME; it can instead use X2 handover directly. This optimization reduces backhaul traffic and delay, even though in practice this X2 connection may still need to be routed via a central SeGW (if present) if the environment is not secure.

X2 handover between two HeNBs is only allowed for cases where the MME does not need to perform access control – i.e. when the source and target HeNBs are in closed or hybrid access mode and have the same CSG ID, or when the target HeNB is in open access mode.

The call flow for this optimized mobility between two HeNBs via X2 handover is the same as the general call flow for an X2 handover between two eNodeBs (see Section 2.6.3.1), even though the PATH SWITCH REQUEST/ACKNOWLEDGE messages may pass (transparently) through the HeNB GTW if present. This optimized mobility works regardless of the type of S1 connectivity employed by the source and target HeNBs (they may be directly connected to an MME or to a HeNB GW, or even each be connected to different HeNB GWs).

24.2.4 Local IP Access Support

In Release 10, a HeNB may operate in Local IP Access (LIPA) mode. LIPA enables an IP-capable UE connected via a HeNB to access other IP-capable devices in the same residential/enterprise IP network without the user plane traversing the mobile operator's network, as shown in Figure 24.3.




Figure 24.3: HeNB operating in LIPA mode in Release 10.

Reproduced by permission of © 3GPP.

When a HeNB operates in LIPA mode, it supports an S5 interface towards the S-GW and an SGi interface towards the IP network to direct the LIPA user plane traffic. The HeNB may reuse the S1-interface IP address for the S5 interface in order to be able to reuse the S1

IPSec tunnel, or it may use a different IP address which would result in the establishment of a different IPSec tunnel. A LIPA connection is always released on outbound handover.

In order to be able to operate in LIPA mode, HeNBs support the following additional functions, regardless of the presence of a HeNB GTW:

- Transfer of the co-located LIPA GateWay (L-GW) IP address of the HeNB via the S1-MME interface to the EPC on every idle-active transition and on every uplink NAS PDU transfer in connected mode;
- Transfer over the S5 and S1-MME interfaces of the co-located L-GW uplink Tunneling End IDs (TEIDs, see Section 2.3.1) used for the LIPA bearers for correlation purposes between the co-located L-GW function and the HeNB;
- Support of basic Packet Data Network GateWay (P-GW) functions in the co-located L-GW, such as support of the SGi interface corresponding to LIPA;
- Additional support of first packet sending, buffering of subsequent packets and internal direct L-GW-HeNB user path management;
- Support of a restricted set of S5 procedures relevant to LIPA.

24.3 Interference Management for Femtocell Deployment

Interference management is a key issue for heterogeneous network deployments of macrocells and small cells in LTE. The problem is most acute for femtocells operating in closed access mode on the same carrier frequency as the macrocells. Interference is typically more manageable in systems with picocells or other open- or hybrid-access small cells because UEs which are causing or suffering from interference can be handed over freely between the macro and small cells. For example, if a Macro UE (MUE) (i.e. a UE connected to a macrocell) is near the edge of the macrocell and is also close to a CSG femtocell, then it will be transmitting at high power and may interfere with the uplink of the small cell. If the MUE can be handed over to the small cell, the interference to the small cell would be eliminated. Similarly in the downlink, a MUE near the macrocell edge might suffer interference from a nearby small cell, which would be eliminated if the MUE could be handed over to the femtocell. The most severe interference occurs when the small cells and macrocells are deployed on the same carrier frequency; otherwise, the interference depends upon the Spectrum Emission Mask (SEM), and, in the case of adjacent channel operation, on the Adjacent Channel Leakage Ratio (ACLR) and Adjacent Channel Selectivity (ACS).

Therefore in this section we focus on the case of co-channel deployments of femtocells and macrocells. Section 31.2 addresses co-channel deployments of picocells and macrocells.

Femtocells are usually installed by the consumer in an ad-hoc fashion, rather than being part of a planned deployment. They are therefore designed to be self-configuring, in that they are required to sense their environment (e.g. to detect and measure neighbouring macrocells and femtocells) and adapt their operation accordingly (see, for example, Section 24.4.1 regarding adjacent channel protection. If the femtocell can choose its carrier frequency, then interference may be controlled by the appropriate carrier selection (for example based on measured Reference Signal Received Power (RSRP) and cell reselection priority information).

Note that co-channel interference between femtocells situated close to each other may also occur.

24.3.1 Interference Scenarios

Here we outline the possible interference scenarios of relevance. (Note that these scenarios apply for both FDD and TDD systems; in the case of TDD it is assumed that synchronization of uplink/downlink switching points is achieved between the macrocell and small cells.)

Downlink transmissions from the femtocell suffer interference from transmissions from a macrocell as shown by interference path A in Figure 24.4.

Macro to Femto, Downlink. UEs connected to femtocells (known as Femto UEs or FUEs) are more susceptible to this interference when they (and their associated HeNB) are closer to the macrocell, since the transmit power of the macrocell is much higher than that of the femtocell and thus the received interference power will be higher. FUEs are also more susceptible when located far from the serving HeNB, especially if they are outside the house or apartment the femtocell is designed to cover.




Figure 24.4: Macro/femto downlink interference scenarios.

Potential interference mitigation approaches for this scenario may include:

- Control channel protection (Physical Downlink Control Channel (PDCCH), synchronization signals (PSS/SSS), Physical Broadcast Channel (PBCH)) by arranging control channels to be orthogonal in time and/or frequency, e.g. by applying a subframe boundary offset in the femtocells relative to the macrocells.

- Data channel protection (PDSCH). If partial frequency re-use (see Section 12.5) is employed in the macrocell then a HeNB can schedule data on the RBs with low transmission power from the macrocell (e.g. RBs being used for cell centre UEs by the macrocell). The HeNB could ascertain the frequency partition information of the MeNB by various means, e.g. by configuration or by monitoring the macrocell transmissions.

Femto to Macro, Downlink. Downlink transmissions from a macrocell suffer interference from transmissions from a HeNB as shown by interference path B in Figure 24.4. This can cause a ‘deadzone’ around a HeNB within which a MUE is unable to receive transmissions from the macrocell. Such deadzones are larger for HeNBs near the edge of the macrocell, where the signal received from the macrocell is weakest, or for MUEs located indoors in the coverage of a CSG HeNB. Potential interference mitigation approaches for this scenario may include:

- Enabling hybrid or open access if possible. In the case of hybrid access, the power settings of the HeNB could be adapted differently to the closed access case, taking into account the total system performance (macro + hybrid cell) and the resources consumed by ‘visiting’ non-CSG UEs.
- Downlink power setting. The HeNB can limit the maximum downlink power (or power per RB) according to its environment. In the case of co-channel deployment of macrocells and femtocells with closed access, there are a number of possible ways this could be achieved [3], such as setting the power to achieve a trade-off between coverage and interference, based on the estimated path-loss between the HeNB and the victim MUEs, and the coverage requirements of the femtocell. Such a solution could also involve detecting the presence of nearby victim MUEs and correspondingly reducing the downlink transmission power; this could be done by detecting uplink transmissions at the HeNB, or by means of measurement reports from victim MUEs to the serving MeNB if it is possible to then signal this information to the HeNB.
- Time-domain coordination making use of Almost Blank Subframes (ABS), a concept introduced in Release 10 (see Section 31.2.2). ABSs contain only certain essential transmissions, leading to a reduction in interference to victim UEs. Typically, an aggressor HeNB would set up a pattern of ABSs resulting in reduced interference to victim MUEs. In Release 10 there is no X2 interface between HeNBs and macro eNodeBs, and therefore ABS patterns at a HeNB would need to be configured either by Operation and Maintenance (O&M) or autonomously by the HeNB; for example, an ABS pattern could be set up at an HeNB to protect subframes containing PSS/SSS/PBCH and paging occasions at a macro eNodeB, assuming the HeNB has obtained time synchronization with the macro eNodeB. Additional mitigation of the residual interference due to the essential transmissions is possible e.g. by arranging the frequency positions of CRS to be different in victim and aggressor cells.

System simulation results [3] are given in Table 24.1 showing the effect of introducing HeNBs, setting the HeNB power appropriately and enabling hybrid access. The system bandwidth is 5 MHz with a co-channel deployment of macrocells (with 1 km inter-site distance) and HeNBs, with 22 UEs per macrocell, 12 of which are FUEs.

Table 24.1: Effect of various small cell configurations on system throughput.

	Outage Probability (SNR < -6 dB)	Worst 20% UE throughput (kbps)	Median throughput (kbps)
No HeNB	12.7%	35	150
CSG HeNB, Fixed Power 8 dBm	18.9%	100	5600
CSG HeNB, Power Control (-10 , +10 dBm)	9.8%	250	3300
Hybrid HeNB, Fixed Power 8 dBm	2%	900	5100
Hybrid HeNB, Power Control (-10 , +10 dBm)	3%	400	3400

Macro to Femto, Uplink. Uplink transmissions from an FUE suffer interference from transmissions from a MUE as shown by interference path C in Figure 24.5.




Figure 24.5: Macro/femto uplink interference scenarios.

Femtocells are more susceptible to this interference when they are located close to the edge of the macrocell edge since MUEs in the vicinity of the HeNB will be transmitting at higher power in this case. System simulations [3] have suggested that this scenario does not create a significant interference impact because in most cases MUEs cannot operate very

close to the HeNB, and therefore the FUEs will typically be closer to the HeNB than the MUEs are. Nevertheless, the interference may be severe if the MUE and HeNB are both situated indoors; in this case the following potential interference mitigation approaches may be relevant:

- Uplink power control: the HeNB can control the power of its FUEs to overcome the interference from neighbouring MUEs;
- Control channel protection (PUCCH): if the PUCCH resources are over-provisioned at the HeNB then the RBs used for PUCCH at the HeNB can be moved away from the band edges (see Section 16.3.1) such that they do not overlap with the RBs used for PUCCH by the macrocell.

Femto to Macro, Uplink. Uplink transmissions from an MUE suffer interference from transmissions from an FUE as shown by interference path D in Figure 24.5. The impact of this interference increases as the density of femtocells within the macrocell coverage area increases. The severity of the impact also depends on whether the FUEs are operating outdoors (in which case the FUEs will tend to be transmitting at higher power and the path-loss to the macrocell will be lower). Potential interference mitigation approaches for this scenario include:

- Uplink power control: the HeNB can control the power of its FUEs to limit the interference to neighbouring macrocells. This can, for example, be based on the estimated path-loss between the FUEs and the macrocells or the estimated path-loss between the FUEs and the HeNB [3].
- Control channel protection (PUCCH): as described above for mitigating macro-to-femto uplink interference.

Femto to Femto, Downlink Downlink transmissions from one femtocell suffer interference from transmissions from another femtocell as shown by interference paths E and F in Figure 24.6. This figure assumes an apartment scenario, with just two apartments for simplicity; however, it should be borne in mind that in practice interference can occur between multiple apartments on the same and different floors.

Potential interference mitigation approaches for this scenario include setting up ABS patterns and/or frequency re-use schemes whereby each HeNB determines its neighbours (from measurements made at the HeNB or FUE) and the associated path-losses. This neighbour information is then used to construct orthogonal patterns of RBs or ABSs to be used by neighboring HeNBs. Both distributed approaches (with or without information exchange between HeNBs [4,5]) and centralized approaches (e.g. at the HeNB GW [3]) have been proposed for constructing the orthogonal sets. Direct X2 connectivity between HeNBs was added in Release 10 for the support of mobility, and this also allows coordination of ABS patterns between HeNBs.

Femto to Femto, Uplink. Uplink transmissions from one FUE suffer interference from transmissions from another FUE connected to another HeNB as shown by interference paths G and H in Figure 24.6. Uplink power control is one potential interference mitigation technique for this scenario.




Figure 24.6: Femto/femto interference scenarios.

24.3.2 Network Listen Mode

Many of the interference mitigation techniques discussed above require the HeNB to make measurements of surrounding macrocells and femtocells. HeNBs are therefore commonly designed to have a Network Listen Mode (NLM) of operation which involves making measurements and decoding system information from neighbouring eNodeBs. This may be done at initial system setup and periodically thereafter. Examples of measurements made in NLM include:

- Uplink interference power, to infer the presence of nearby MUEs;
- Determination of cell IDs and CSG status/ID, by decoding the SI of the neighbour cells;
- Co-channel RSRP and RS transmission power to estimate the path-loss to neighbour cells, which in turn may be used for uplink and downlink power control at the HeNB;
- Reference Signal Received Quality (RSRQ) which can be used together with RSRP to determine the reliability of coverage of a macrocell, to help determine a suitable power setting for a femtocell operating in hybrid access mode.

In addition, TDD femtocells may obtain time synchronization from neighbouring macrocells.

24.4 RF Requirements for Small Cells

The unique characteristics of small cells, and in particular the potential interference scenarios, give rise to specific Radio Frequency (RF) challenges and requirements. These are explained in this section and compared to those defined for macrocells (see Chapter 21). In Release 8 of the LTE specifications, RF requirements were defined only for the Wide Area eNodeB class which covers macrocell applications, while Release 9 saw the introduction of the Local Area eNodeB class (covering picocell applications) and the Home eNodeB class (covering femtocell applications).

A major factor influencing the RF requirements for pico and Home eNodeBs is the lower minimum separation distance (and hence lower coupling loss) between UEs and the pico or Home eNodeBs when compared to the Wide Area eNodeB case. As explained in

Section 24.1, picocells tend to be small cells planned by the network operator, with the pico eNodeBs mounted at low elevation (e.g. interior walls), in contrast to femtocells for which the Home eNodeBs are typically self-installed in a home environment by the end user. Both picocells and femtocells will have lower Minimum Coupling Loss (MCL) to UEs than macrocells. Typically an MCL of 70 dB is assumed for macrocells compared to 45 dB for picocells and femtocells [6, 7].

Other factors influencing the differences from the Wide Area eNodeB requirements include lower transmit power, degraded receiver sensitivity (reflecting the lower cost), and for Home eNodeBs, the lack of a planned deployment. For these reasons in many cases the Home eNodeB class generally has the most stringent RF requirements, with the Local Area class requirements being intermediate between those of the Wide Area and the Home classes.

24.4.1 Transmitter Specifications

A comparison of transmission requirements between Home, Local Area and Wide Area eNodeBs is provided in Table 24.2.

Table 24.2: Summary of transmitter requirements.

Requirement	Wide Area	Local Area	Home
Maximum Output Power ^a (MOP) (dBm)	46	24	20
Power adjustment for adjacent channel	No	No	Yes
Adjacent Channel Leakage Ratio (ACLR) lower limit (dBm)	-15	-32	-50
Spurious emissions (dBm/100 kHz) ^b	-62	-62	-71
Frequency Error (ppm)	0.05	0.10	0.25

^a Summed over transmit antennas.

^b For co-existence with UTRA or E-UTRA downlinks in other bands.

Maximum Output Power (MOP). For the Local Area class, the MOP was selected to provide acceptable performance loss to overlaid Wide Area macrocellular systems on the same or adjacent carriers, and also to ensure safety from electromagnetic radiation [8, 9]. For the Home eNodeB class, the adjacent channel selectivity of UEs was also taken into account, such that UEs are still able to operate when a HeNB at 45 dB MCL is transmitting in a channel adjacent to that being received by the UE [7].

An HeNB is required to detect if the downlink adjacent channel is being used by another operator for either UTRA or E-UTRA, and to set its MOP according to the measured power of the adjacent channel base station and the total measured noise plus interference in its channel, I_{oh} [6]. This aims to reduce the impact of adjacent channel leakage on neighbouring operators. For a UTRA victim system the requirements are the same as for a UMTS Home NodeB (HNB) [10]. For an LTE victim system, different MOP levels can be set depending on the value of the cell-specific Reference Signal (CRS) Received Power (RSRP) per Resource Element (RE) of the adjacent-channel eNodeB, $\hat{E}_{c,CRS}$ [6]. If $\hat{E}_{c,CRS} < -127$ dBm, then there are unlikely to be UEs connected to this adjacent channel eNodeB in the vicinity of the

HeNB, since the signal is weak, so no power reduction by the HeNB is required. If $\hat{E}_{c,CRS}$ is less than a threshold, then it is considered that the power measurement on the adjacent channel is unreliable due to adjacent channel interference from the uplink channel being used by the HeNB, and the MOP is set to an intermediate value of 10 dBm. Otherwise, the MOP is set between 8 and 20 dBm depending on the value of $\hat{E}_{c,CRS}$ [6].

Adjacent Channel Leakage Ratio. For all eNodeB classes the ACLR is defined by a relative value (45 dBc), subject to an absolute lower limit [6]. For the HeNB class this lower limit is lower than Wide Area eNodeB class due to the lower MCL [11].

Spurious emissions. Spurious emissions limits are defined for co-existence with HeNBs operating in other frequency bands. In deriving the HeNB requirement it was assumed that the HeNB is in an adjacent apartment to a neighbouring HNB or HeNB operating in a different band, and that an MCL of 47 dB applies in this case. Moreover, a 0.8 dB desensitization criterion was assumed [12]. Spurious emission requirements for the protection of the receiver band (Home and Local Area FDD classes), and for co-located base stations (Local area class) are all relaxed by 8 dB, corresponding to the relaxation in the sensitivity requirement relative to the Wide area eNodeB.

Frequency error. Table 24.2 also shows the requirements for frequency error. The total error seen by a UE is the sum of the errors due to frequency error at the eNodeB and Doppler shift due to mobility. For Home and Local Area eNodeBs the mobility is assumed to be restricted to 30 km/h and 50 km/h respectively, compared to a maximum speed of 350 km/h for Wide Area eNodeBs. Therefore for the same total error, the component arising from frequency error can be larger for the Home and Local Area eNodeB classes. In [12] it is shown that 0.25 ppm is a sufficient accuracy considering handover measurement performance, demodulation performance and maintenance of timing synchronization for TDD HeNBs. For the Local Area eNodeB class the frequency error is the same as that already supported in GSM and UTRA Local Area base stations, i.e. 0.1 ppm [13].

Spectrum Emission Mask (SEM). The SEM for the Home and Local Area eNodeB classes differs from the Wide Area class in both the absolute value and the level relative to the carrier power. For the Home class, the mask at high frequency offsets is a function of the transmitted power (summed over all antennas), subject to a minimum of 2 dBm and a maximum of 20 dBm. This is shown in Figure 24.7 for different system bandwidths.

24.4.2 Receiver Specifications

This section introduces the receiver specifications for the Local Area and Home eNodeB classes. Table 24.3 summarizes the wanted and interfering signal levels for these classes for the case of a 10 MHz LTE system.

Reference sensitivity and noise figure. For Local Area and Home eNodeBs the maximum path-loss between a served UE and the eNodeB is considerably less than for a Wide Area eNodeB. In addition, lower implementation cost is important, especially for HeNBs. For

these reasons the assumed noise figure for Local Area and Home eNodeBs is 13 dB, 8 dB higher than for Wide Area eNodeBs.




Figure 24.7: Spectrum Emission Mask for Home eNodeB class.

Table 24.3: Wanted and interfering signal power level requirements for 10 MHz system bandwidth.

	Wide Area		Local Area		Home Area	
	Wanted Signal (dBm)	Interference (dBm)	Wanted Signal (dBm)	Interference (dBm)	Wanted Signal (dBm)	Interference (dBm)
Reference sensitivity	-101.5	-	-93.5	-	-93.5	-
Dynamic range	-70.2	-79.5	-62.2	-71.5	-25.7	-35.0
In-channel selectivity	-98.5	-77.0	-90.5	-69.0	-90.5	-69.0
Narrowband blocking	-95.5	-49.0	-87.5	-41.0	-79.5	-33.0
ACS	-95.5	-52.0	-87.5	-44.0	-71.5	-28.0
Receiver IM	-95.5	-52.0	-87.5	-44.0	-79.5	-36.0
Out-of-band blocking	-95.5	-15.0	-87.5	-15.0	-79.5	-15.0
In-band blocking	-95.5	-43.0	-87.5	-35.0	-79.5	-27.0
Co-located blocking*	-95.5	+16.0	-87.5	-6.0	--	--

* For the Local Area class these requirements are adopted in the case of co-location with UTRA or E-UTRA (however, the interfering signal level is slightly different in the case of co-location with GSM/DCS – see [6, Table 7.6.2.1-2]).

This increase of the noise figure has an impact on several other receiver requirements. Specifically, both the wanted signal level and interfering signal level are increased by 8 dB for the in-channel selectivity (for both Local Area and Home classes), in-band blocking (for the Local Area class), narrowband blocking (for the Local Area class) and Adjacent Channel Selectivity (ACS) (for the Local area class).

Dynamic range. The required dynamic range in an eNodeB depends not only on the ability to control the transmit power of the uplink transmissions from served UEs, but also on the interference levels seen from UEs in neighbouring cells, in particular from co-channel UEs served by the macrocellular network which could be transmitting with high power. For the Home eNodeB class operating in closed access mode, the fact that the deployment is unplanned means that such co-channel MUEs could come close to the HeNB. However, as discussed in Section 24.3.1, in such cases the HeNB will create a ‘deadzone’ around itself due to the interference caused to MUEs in the downlink. As a result of this deadzone, a coupling loss higher than the MCL of 45 dB was assumed when deriving the dynamic range requirement.

Blocking, Narrowband Blocking, Adjacent Channel Selectivity (ACS) and Receiver Intermodulation. The eNodeB blocking requirements consist of three components as shown in Figure 24.8. The first component is Out-Of-Band (OOB) blocking, which is defined over a wide frequency range,⁷ excluding the operating band.

The second component is the in-band blocking, for which the interfering signal is an LTE signal at a specified frequency separation from the wanted signal. For the HeNB class, the interfering signal level is increased relative to the other classes to take account of the unplanned nature of HeNB deployments. This is also the case for the narrowband blocking and ACS requirements. For these cases, system simulations were used to define the interference powers [12].

The third component is the co-located blocking, which is defined for co-located (e)NodeBs operating in other frequency bands. No requirements are currently defined for the HeNB class. For the Local Area class, requirements are defined which assume the same 30 dB coupling loss as for the Wide Area class and a maximum transmit power of 24 dBm from interfering Local Area E-UTRA eNodeBs.

Due to the lower probability of two large interfering signals being present simultaneously, the power level of the interfering signals for the inter-modulation requirement is lower compared to the blocking requirement [12].

24.4.3 Demodulation Performance Requirements

For the Home eNodeB class, the maximum served UE speed is assumed to be 30 km/h which corresponds to a maximum Doppler of about 70 Hz at a 2.5 GHz carrier frequency. For the Local Area class, the maximum served UE speed is assumed to be 50 km/h. Furthermore, due to the small cell size in the indoor environment, for both these classes the multipath delays tend to be significantly smaller than for the Wide Area eNodeB class. Therefore only a subset of the Wide Area demodulation performance requirements are specified, consisting of the

⁷For most operating bands, it is defined for a frequency range of 20 MHz either side of the operating band.




Figure 24.8: Blocking requirements.

EPA and EVA propagation models⁸ with a maximum Doppler frequency no larger than 70 Hz. In these cases the performance requirements are the same as for the Wide Area eNodeB class. In addition for the Home class, new requirements are defined for the decoding of Hybrid Automatic Repeat reQuest ACKnowledgements (HARQ-ACKs) and channel quality information feedback.

24.4.4 Time Synchronization for TDD Operation

Time synchronization in TDD systems is an important consideration for the avoidance of interference between uplink and downlink transmissions on neighbouring eNodeBs (see Section 23.3). For the HeNBs this is particularly important in view of the potential interference scenarios and unplanned deployment. Therefore a mandatory accuracy requirement is specified for HeNBs, even though, in common with all eNodeBs, no specific method is mandated. The difference in radio frame start timing, measured at the transmit antenna connectors, between the HeNB and any other HeNB or eNodeB which has overlapping coverage [12] is required to be less than 3 μ s, except in cases where timing is obtained by monitoring another eNodeB which is more than 500 m away. In this case the requirement is relaxed in line with the additional one-way propagation delay beyond 500 m.

The requirement is designed to ensure that the combination of synchronization error, propagation delay to a victim, and multipath delay spread remains less than the smallest Cyclic Prefix (CP) length, taking into account that multipath delay spreads tend to be small in femtocells; if this condition is satisfied, interference at the uplink/downlink switching points within the TDD radio frame will be avoided.

⁸Extended Pedestrian A and Extended Vehicular A – see Chapter 20.

24.5 Summary

LTE Release 9 saw the introduction of new classes of base stations: Home eNodeBs or femtocells to cover homes and apartments, usually on a closed subscriber group basis, and pico eNodeBs to cover hotzones and other local areas.

While the basic architectural principles remain the same as for LTE macro eNodeBs, some new features were introduced specifically to handle the challenges posed by small cells – most notably the introduction of the Home eNodeB Gateway to address the needs of very large densities of femtocells, and, in Release 10, direct X2-interface handover between Home eNodeBs.

Heterogeneous deployments of small cells and macrocells bring new problems in terms of interference; some potential interference management schemes are described in this Chapter.

Finally, in the last part of the chapter, some of the RF challenges and requirements are discussed for both pico and Home eNodeBs.

References⁹

- [1] Request for Comments 4960 The Internet Engineering Task Force (IETF), Network Working Group, ‘Stream Control Transmission Protocol’, www.ietf.org.
- [2] 3GPP Technical Specification 23.003, ‘Numbering, addressing and identification (Release 9)’, www.3gpp.org.
- [3] 3GPP Technical Report 36.921, ‘FDD Home eNodeB (HeNB) RF requirements analysis (Release 9)’, www.3gpp.org.
- [4] Qualcomm, ‘R4-091906: Frequency reuse results with full buffer’, 3GPP TSG RAN WG4, meeting 51, San Francisco, USA, May 2009.
- [5] CMCC, ‘R4-092872: Downlink interference coordination between HeNBs’, 3GPP TSG RAN WG4, meeting 52, Shenzhen, China, August 2009.
- [6] 3GPP Technical Specification 36.104, ‘Evolved Universal Terrestrial Radio Access (E-UTRA); Base Station (BS) Radio Transmission and Reception (Release 9)’, www.3gpp.org.
- [7] CATT, ‘R4-094074: Home eNode B Maximum output power’, 3GPP TSG RAN WG4, meeting 52bis, Miyazaki, Japan, October 2009.
- [8] CATT, ‘R4-092852: Proposal of maximum output power for Pico eNB’, 3GPP TSG RAN WG4, meeting 52, Shenzhen, China, August 2009.
- [9] Huawei, ‘R4-092810: LTE Pico NodeB maximum output power’, 3GPP TSG RAN WG4, meeting 52, Shenzhen, China, August 2009.
- [10] 3GPP Technical Specification 25.104, ‘Base Station (BS) Radio Transmission and Reception’, www.3gpp.org.
- [11] CMCC, ‘R4-091789: Analysis of absolute ACLR1 requirements for LTE TDD HeNB’, 3GPP TSG RAN WG4, meeting 51, San Francisco, USA, May 2009.
- [12] 3GPP Technical Report 36.922, ‘TDD Home eNodeB (HeNB) RF requirements analysis (Release 9)’, www.3gpp.org.
- [13] CATT and Huawei, ‘R4-092809: TP on frequency error for Pico eNodeB’, 3GPP TSG RAN WG4, meeting 52, Shenzhen, China, August 2009.

⁹All web sites confirmed 1st March 2011.

Self-Optimizing Networks

Philippe Godin

25.1 Introduction

The provision of Self-Optimizing Network (SON) functions is one of the key differentiators of LTE compared to previous generations of cellular systems such as UMTS and GSM. Self-optimization of the network is a tool to derive the best performance in a cost-effective manner, especially in changing radio environments. It allows the network operator to automate key aspects of the network configuration processes, and thus reduces the need for centralized planning and human intervention. For these reasons, this feature has been given a high priority and was a cornerstone around which the LTE radio, S1 and X2 procedures were designed. This makes SON functionality particularly efficient in the LTE system. The involvement of the User Equipment (UE) in the SON functionality of LTE is another key contributor to its success.

This chapter starts with an explanation of the Automatic Neighbour Relation (ANR) Function, the functionality for self-configuration of the eNodeB and MME¹ and the automatic Physical Cell Identity (PCI) configuration as natively implemented within the basic S1 and X2 interface procedures.² These three SON functions, which were included in Release 8, are of particular relevance for the initial deployment of an LTE network.

The chapter then explores other SON functions developed in Release 9 which are designed to optimize deployed LTE networks, such as the Mobility Load Balancing (MLB), Mobility Robustness Optimization (MRO) and Random Access CHannel (RACH) optimization functions. The chapter also includes the latest Release 10 SON enhancements which further optimize advanced LTE networks. Finally, as SON is a continuously evolving area, the

¹Mobility Management Entity.

²For details about the S1 interface see Section 2.5 and [1]; for X2 see Section 2.6 and [2].

chapter concludes with examples of other new SON functions which are envisioned to complement the SON family in the near future.

Further details of the specified SON techniques can be found in [3,4].

25.2 Automatic Neighbour Relation Function (ANRF)

The ANR Function (ANRF) is an example of a SON function which exploits both the design of the LTE radio interface and the UE to enable the network to optimize itself. The purpose of the ANRF is to relieve the network operator from the burden of manually managing relations between neighbouring cells.

25.2.1 Intra-LTE ANRF

The ANRF relies on the cells broadcasting their globally unique identity, termed the E-UTRAN Cell Global Identifier (ECGI). The ECGI is composed of 3 bytes carrying the Public Land Mobile Network (PLMN) ID and 28 bits to identify the cell within that PLMN. The function involves User Equipments (UEs), when requested by their serving eNodeB, reading and reporting the ECGI broadcast by a neighbouring cell that has been detected previously by that UE or another UE.

When an eNodeB receives from a UE a Physical Cell Identity (PCI) of a neighbour cell as part of a normal measurement report, and the eNodeB does not recognize the PCI, the eNodeB can instruct the UE to execute a new dedicated reporting procedure which uses the newly discovered PCI as a parameter. Through this procedure, the UE reads and reports to the requesting eNodeB some system information of the detected neighbouring cell, including the ECGI, the Tracking Area Code (TAC) and all available PLMN IDs. An example of this procedure is illustrated in Figure 25.1.




Figure 25.1: The intra-LTE ANR Function. Reproduced by permission of © 3GPP.

By means of the ANRF, each eNodeB can thus automatically populate a Neighbour Relation Table (NRT) for each cell it controls, containing all the Neighbour cell Relations (NRs) of the cell. An existing NR is defined as a unidirectional cell-to-cell relation from one

source cell controlled by the eNodeB to a target cell for which the eNodeB controlling the source cell:

- knows the ECGI and PCI;
- has an entry in the NRT for the source cell that identifies the target cell;
- has the attributes (detailed in the following subsection) in this NRT entry defined either by Operation and Maintenance (O&M) or set to default values.

25.2.2 Automatic Neighbour Relation Table

From the discussion above, it can be seen that the ANRF is a cell-related function managed by the eNodeB. In the NRT, each NR can have a number of attributes, including:

- **No Remove flag:** The eNodeB shall not remove the NR from the NRT if this flag is set. This is typically used if the NR is certain, for example because it has been configured by O&M.
- **No Handover flag:** The NR shall not be used by the eNodeB for handover purposes if this flag is checked. This may, for example, be used if handovers are not useful from the source cell to the target cell, but the NR is nevertheless useful for the purpose of exchanging interference information.
- **No X2 flag:** The NR shall not use an X2 interface to initiate procedures towards the eNodeB that controls the target cell. This may, for example, be due to the source and target cell belonging to different PLMNs.

The ANRF incorporates three main functions for the management of NRs in the NRT, as shown in Figure 25.2:

- The ‘NRT management function’ which manages the table, including modifying NR attributes;
- The ‘neighbour detection function’ which finds new neighbour cells and adds them to the NRT;
- The ‘neighbour removal function’ which removes outdated NRs, for example after expiry of a timer.

The ANRF also allows the NRTs to be managed by O&M. O&M functionality can add or delete NRs or change the attributes of an NR. Conversely, O&M is also informed about changes in an NRT.

25.2.3 Inter-RAT or Inter-Frequency ANRF

The ANRF also exists between LTE and other Radio Access Technologies (RATs) or towards other frequencies. For inter-RAT and inter-frequency ANR, each cell is assigned an inter-frequency search list containing all frequencies that should be searched. The target RATs for inter-RAT ANRF are UTRAN³, GERAN⁴ and CDMA2000. The inter-RAT/inter-frequency ANR function can be divided into five steps:

³UMTS Terrestrial Radio Access Network.

⁴GSM EDGE Radio Access Network.




Figure 25.2: The ANR function and the ANR table. Reproduced by permission of © 3GPP.

1. In connected state, the eNodeB instructs a UE to search for neighbour cells in the target RATs/frequencies (it may need to schedule appropriate idle periods for this).
2. The UE reports the PCIs of the detected cells in the target RATs/frequencies (note that each RAT has its own specific format of PCI).
3. The eNodeB instructs the UE, using the newly discovered PCI as a parameter, to read the key RAT-specific cell identification parameters from the broadcast channel of the cell (e.g. the CGI and Routing Area Code (RAC) for GERAN; CGI, Local Area Code (LAC) and RAC for UTRAN; CGI for CDMA2000).
4. The UE reports these key RAT-specific cell identification parameters to the eNodeB of the serving cell.
5. The eNodeB updates its inter-RAT/inter-frequency NRT. The eNodeB can then make use of the NR, for example when triggering subsequent handovers.

25.3 Self-Configuration of eNodeB and MME

Self-configuration of the eNodeB/MME is a SON function which is implemented in the basic S1 and X2 interface procedures [1, 2].

25.3.1 Self-Configuration of eNodeB/MME over S1

With the native support of the S1-flex function in LTE (see Section 2.5), an eNodeB must set up an S1 interface towards each MME of the pool area to which it belongs. The list of MME nodes of the pool area together with an initial corresponding remote Internet Protocol (IP) address can be directly configured in the eNodeB at deployment (although other means may also be used). Once the eNodeB has initiated a Stream Control Transmission Protocol (SCTP, see Section 2.5.1.1) association with each MME of the pool area using that IP address, they can exchange, via the 'S1 Setup procedure' (see Section 2.5), some application-level configuration data which is essential for the system operation. This automatic configuration process thus saves some manual configuration effort for the network operator, together with the associated risk of human error.

Examples of such application-level configuration data exchanged between the eNodeB and the MMEs include Tracking Area (TA) identities,⁵ lists of PLMNs of different operators who may be sharing the network so that all the PLMN IDs can be broadcast over the air for their respective UEs, and Closed Subscriber Group (CSG) IDs to allow auto-configuration of the Home eNodeB (HeNB) gateway when it connects to thousands of HeNBs.⁶ Once all the data to be broadcast over the radio interface have been configured within each and every eNodeB, they are sent automatically to all the relevant MME nodes of the pool area via the S1 Setup procedure.⁷

The eNodeB can later update the configuration data it had previously sent to the MME by sending an 'eNB CONFIGURATION UPDATE' message. In this case it only sends the updated configuration data, and the data that is not included is interpreted by the MME as being unchanged. Conversely, the MME can also send updates of its data to the eNodeBs by means of an 'MME CONFIGURATION UPDATE'. The updated configuration data is assumed to be stored in both the eNodeB and the MME for the duration of the SCTP association or until any further update occurs.

25.3.2 Self-Configuration of IP address and X2 interface

Similarly to the S1 interface, self-configuration of IP addresses and of the X2 interface is implemented in the basic X2 and S1 interface procedures. The X2 interface may be established between one eNodeB and one of its neighbour eNodeBs when they need to exchange load, interference or handover related information (see Section 2.6). The automatic initialization of the X2 interface consists of three steps:

1. The eNodeB identifies a suitable neighbour;
2. The eNodeB retrieves a suitable IP address for this neighbour if not already available and sets up an SCTP association with it;
3. The two eNodeBs exchange configuration data.

⁵TAs correspond to the zones in which UEs are paged, and their mapping to eNodeBs must remain consistent between the E-UTRAN and the Evolved Packet Core (EPC).

⁶This enables the paging optimization feature in the HeNB gateway. Further details can be found in Chapter 24.

⁷Note that in case of CSG, the S1 Setup messages are exchanged via the Home eNodeB gateway to the MME (see Section 24.2).

The first step can be achieved either by configuration or by using the ANRF described in Section 25.2. If the ANRF is used, this step basically consists of the eNodeB being made aware of the ECGI and Tracking Area Identity (TAI) of the detected neighbour.

For the second step, the eNodeB needs to know the IP address of the neighbour in order to set up an SCTP association. This IP address may again be either configured or retrieved via the network (the latter being used if the ANRF is used during the first step). Auto-configuration of the IP address is achieved by the requesting eNodeB sending over the S1 interface a dedicated ‘eNB CONFIGURATION TRANSFER’ message that includes both routing information (such as the ECGI of the detected target cell) and the nature of the information that is requested – in this case an IP address for the purpose of X2 initiation. The requesting eNodeB also includes its own ECGI to be used for routing back the answer.

If the receiving eNodeB agrees, it returns one or more IP addresses which can be used for the establishment of an X2 interface. When this procedure is complete, the requesting eNodeB can set up the SCTP association with its neighbour by sending an SCTP INIT message. In Release 10, in order to protect the eNodeBs from malicious SCTP INIT requests from unauthorized parties, this auto-configuration process has been enhanced to enable the use of an Access Control List (ACL)⁸ of authorized source IP addresses in the receiving eNodeB. For this purpose, the ‘eNB CONFIGURATION TRANSFER’ message has been enhanced with the possibility for the requesting eNodeB to include one or several IP addresses that the receiving eNodeB can store in its ACL. Thus, whenever a further SCTP INIT message is received to set up an X2 interface, the receiving eNodeB can first check that the source IP address corresponds to one notified earlier. The protocol also allows the requesting eNodeB to provide IP addresses for an IPsec⁹ transport endpoint for scenarios where IPsec is expected to be used (e.g. routing via a security gateway). Finally, the requesting eNodeB can also include the IP addresses that it intends to use for the data-forwarding GPRS Tunnelling Protocol (GTP) tunnels that it will later establish. This would also allow for checking of the user plane traffic at the receiving eNodeB. The reciprocal behaviour is also supported: the receiving eNodeB may similarly provide in the ‘eNB CONFIGURATION TRANSFER’ message to the requesting eNodeB all the IP addresses it intends to use for its control plane SCTP endpoint, user plane GTP endpoints and/or IPsec endpoint.

Once an SCTP association exists between these two neighbour eNodeBs, the third step can be started, i.e. exchanging configuration data. This consists of application-level data similar to the data exchanged during the self-configuration of the S1 interface (see Section 25.3.1). In this case, the ‘X2 Setup’ procedure is used for the exchange. For example, an eNodeB can report, via the ‘X2 SETUP REQUEST’ message to a neighbour eNodeB, information about each cell it manages, such as the cell’s PCI, frequency band, TAI and/or associated PLMNs. More detailed radio parameters can also be included, such as the cyclic prefix length (see Section 5.4.1), the transmission bandwidth or the uplink-downlink subframe configuration (see Section 6.2) for Time Division Duplex (TDD) cells.

An eNodeB can also exchange the list of pool areas to which it belongs with a neighbour eNodeB. The neighbour eNodeB can thus automatically learn if it shares a pool area in common and therefore whether it will need to use the S1 or the X2 handover procedure

⁸A receiving network node where ACL functionality is applied may only accept connections from other peer network nodes once the source addresses of the sending network node are known in the receiving node.

⁹IP Security – a collection of protocols and algorithms for IP security, including key management.

to transfer UEs. Indeed, if the eNodeBs do not share a pool area in common, the MME associated with the UE must be relocated on handover and the S1 handover has to be used (see Sections 2.5.6 and 2.6.3 respectively).

25.4 Automatic Configuration of Physical Cell Identity

The application-level configuration data exchanged during the X2 setup procedure is also the core of another SON feature: automatic self-configuration and self-optimization of the Physical Cell Identities (PCIs). This helps the eNodeBs to select PCIs that avoid collisions and hence cell confusion. Cell confusion arises when two neighbouring cells broadcast the same PCI so that a UE cannot discriminate between the two cells when it reports measurements. As a consequence the serving eNodeB of that UE cannot determine which one of these two cells should be the handover target for the UE. Increased inter-cell interference may also arise. The possibility for cell confusion stems from the fact that only about five hundred PCI values are available. PCI collision can be avoided by careful configuration on the part of the network operator, i.e. by selecting the PCIs allocated to each cell so that they are unique within clusters of adjacent neighbouring cells; however, this is a laborious operation, and moreover there remains an ongoing risk of further PCI collisions due to the start up of new eNodeBs as the network is densified.

The SON solution to this problem relies on the exchange of PCI values between neighbour eNodeBs during the X2 Setup procedure. In both the 'X2 SETUP REQUEST' and the 'X2 SETUP RESPONSE' messages, an eNodeB can include the list of PCI values used not only by its own cells but also by the 'direct neighbours' of its own cells. A direct neighbour of a cell is defined as any cell controlled by an eNodeB that is a neighbour of the eNodeB controlling the first cell (even if that cell has not yet been reported by any UE). This exchange of direct neighbour PCI values over the X2 interface enables an eNodeB to become quickly aware of the set of PCI values that are being used in the cluster to which it belongs. In particular the eNodeB can easily identify any collision in this cluster and can decide to change the PCI of one of its cells if needed; if it does so, it can signal the change to its neighbours in an 'eNB CONFIGURATION UPDATE' message. However, the exact PCI change algorithm supported by eNodeBs is not standardized and remains up to vendor implementations. An example of self configuration of PCIs is illustrated in Figure 25.3.

Via O&M, the network operator can use a variety of degrees of self-configuration at start up of the network: the operator could assign no PCI values and let each eNodeB select PCIs fully autonomously, or alternatively a range of possible PCI values can be assigned in order to assist the convergence of the self-configuration algorithms.

25.5 Mobility Load Balancing Optimization

Release 9 incorporates SON load balancing functionality, the objective of which is to counteract local traffic load imbalance between neighbouring cells with the aim of improving the overall system capacity and reducing congestion. The feature functions by first detecting any traffic imbalance and then applying solutions such as adjusting the cell reselection/handover parameters (such as handover thresholds). These parameters can be autonomously changed




Figure 25.3: Illustration of PCI allocation. Reproduced by permission of © 3GPP.

and directly communicated between neighbouring cells by means of the X2 Parameter Negotiation procedure as explained in Section 25.5.2.

25.5.1 Intra-LTE Load Exchange

In order to detect an imbalance, it is first necessary to exchange load information between neighbouring eNodeBs over X2 for comparison. A client-server mechanism is used for this purpose: a requesting eNodeB (client) sends a ‘RESOURCE STATUS REQUEST’ message to request a load report from some of its neighbours. The ‘RESOURCE STATUS REQUEST’ message can simultaneously request multiple types of load report and may also be directed at multiple cells of the receiving eNodeB (server). The neighbours that receive the request report the requested load information over the X2 interface via the ‘RESOURCE STATUS RESPONSE/UPDATE’ message. The reporting of the load is periodic with period indicated in the ‘RESOURCE STATUS REQUEST’ message.

The reported information can indicate any of four different types of cell load information:

- Current usage of Physical Resource Blocks (PRBs), possibly partitioned into real-time and non-real-time traffic;
- Current hardware load;
- Current S1 transport load;
- Available composite load.

The first three measurements represent a global view of the current load situation in the node that reports them. The ‘available composite load’ indicator represents the amount of overall resources that the reporting node is ready to accept. ‘Composite’ means that the reporting node takes into account multiple internal resource criteria via a proprietary evaluation to build up its report. The ‘available’ characteristic is interesting as an estimate

of the amount of non-GBR¹⁰ traffic that can be handed over into the cells controlled by the eNodeB. For example, in a case when one UE was using all PRBs, the ‘current usage of PRBs’ would typically indicate that all PRBs were fully utilized whereas those PRBs would in fact still be available to be shared by the traffic of a second UE. Two formats may be used by the operator for the ‘available composite load’ indicator:

- A simple percentage of the total E-UTRAN resources available (i.e. of the total cell uplink or downlink bandwidth known from the X2 Setup procedure);
- A percentage weighted according to a cell capacity class value.¹¹

25.5.2 Intra-LTE Handover Parameter Optimization

As a result of the exchange of load information and detection of local load imbalance, eNodeBs may take immediate action such as deciding to handover some UEs to cells which are less loaded. However, longer-term actions can also be taken to combat the imbalance more efficiently. For example when an imbalance is detected between two specific cells it may be desirable to shift the handover trigger threshold by, for example, 0.5 dB. By so doing, UEs served by an overloaded cell may find that a less-loaded neighbour cell becomes a relatively attractive handover target, thus allowing some UEs from the overloaded cell to be served reliably by the less-loaded cell. In order to avoid a ping-pong effect, it is often desirable that the lightly loaded cell shift its corresponding threshold by a similar amount as the overloaded cell but in the opposite direction. For this to happen, the lightly loaded cell must be made aware of the change applied by the overloaded cell, or alternatively the overloaded cell could first request the less-loaded cell to change its handover threshold and wait for a positive response before initiating the change.

This is the object of the ‘Handover Parameter Negotiation’ procedure. If the two cells in question belong to the same eNodeB, this negotiation happens within the eNodeB node itself via a proprietary algorithm; otherwise, the negotiation procedure is conducted over X2 interface via the ‘Mobility Settings Change’ procedure. This procedure enables one eNodeB to send a ‘MOBILITY CHANGE REQUEST’ message to another, to indicate the handover trigger parameter shift that the first eNodeB sees as necessary for one of the cells controlled by the receiving eNodeB. The same message can optionally indicate if the first eNodeB has already performed any configuration change of this parameter for one of its own cells. If the second eNodeB accepts the proposed handover trigger parameter modification and is able to complete it, then it replies with a ‘MOBILITY CHANGE ACKNOWLEDGE’ message; otherwise it responds with a ‘MOBILITY CHANGE FAILURE’ message which can include an ‘allowed parameter modification range’ if the reason for the failure was that the change proposed by the first eNodeB was outside the possible range.

¹⁰Guaranteed Bit-Rate – see Section 2.4.

¹¹The cell capacity class value is a parameter between 0 and 100 configured by the operator to classify one particular cell capacity with respect to the other cells available in its network. The use of this parameter is mandatory when the ‘available composite load’ indicator is exchanged inter-RAT, e.g. towards UMTS.

25.5.3 Inter-RAT Load Exchange

In order to perform load information exchange with neighbour cells of a RAT other than LTE, an eNodeB can trigger a ‘Cell Load Reporting Request/Response’ procedure towards a neighbouring Radio Network Controller (RNC) (for WCDMA), Base Station Controller (BSC) (for GSM) or evolved High Speed Packet Access (HSPA) NodeB (for HSPA). This new inter-RAT procedure was purposely designed to be independent of the handover procedure in order that it can be triggered at any time by a requesting eNodeB even in the absence of any mobility event. The procedure is implemented as a generic ‘SON Transfer’ container carried on top of the RAN Information Management (RIM) protocol.¹² The new inter-RAT cell load reporting procedure allows an eNodeB to evaluate the load situation of GSM or UMTS/HSPA neighbour cells before triggering any load-related action such as handing traffic over to those neighbours or switching off a carrier. The new procedure may be used in all directions, i.e. from any LTE/UMTS/GSM source node to any LTE/UMTS/GSM target node of another RAT.

25.5.4 Enhanced Inter-RAT Load Exchange

One limitation of the Release 9 inter-RAT load exchange procedure described in Section 25.5.3 is that the RIM messaging which carries it passes through the core network nodes, which may become overloaded if load exchange occurs too frequently. On the other hand, variations in the radio conditions mean that it is important for the RAN nodes to have up-to-date load values from their peers in order to assess the load situation across different RATs. In order to solve these two opposing requirements, the inter-RAT load exchange procedure was enhanced in Release 10 by the introduction of two new types of reporting:

Event-triggered reporting enables a source RAN node to request a target RAN node to report only when a pre-defined cell load-level threshold is crossed in the target node. The source node signals the granularity with which it wishes to receive indications of load variation: for example, if it requests a granularity of five levels, the target RAN node will divide the cell-load scale by five, evenly distributed on a linear scale below the overload threshold and will report the target cell load each time the load changes from one reporting level to another or enters or exits the overload state. Other parameters (such as hysteresis and averaging factors) can be configured by O&M at the target nodes.

Multiple-cell reporting enables a source RAN node to request the a target RAN node to report the load of multiple cells within a single report. For example, in the case of UMTS, an eNodeB will only need to send one request to the RNC instead of individual requests for all the cells controlled by the RNC.

¹²RIM is a well known protocol dedicated to information exchange between two RAN nodes of different RATs. An example is given in Section 25.6.7.

25.6 Mobility Robustness Optimization

The Mobility Robustness Optimization (MRO) SON feature aims at detecting and preventing connection failures that occur as a result of mobility. These failures are of four types:

- Too-late handover, leading to a connection failure in the serving cell;
- Coverage hole, leading to a connection failure in the serving cell;
- Too-early handover, leading to a connection failure in the target cell;
- Handover to an inappropriate cell, leading to a connection failure in the wrong target cell.

Connection failure cases generally correspond to Radio Link Failures (RLFs), including handover failure. These scenarios and their respective SON solutions are described in more detail below.

25.6.1 Too-Late Handover

When a handover is executed too late, a connection failure may occur in the source cell and the UE may try to re-establish the radio link in a different cell controlled by a different eNodeB, as illustrated in Figure 25.4. In order to enable eNodeB of the source cell to improve the situation, an ‘RLF INDICATION’ message can be used to report the event from the eNodeB under which the UE re-establishes the connection to the original source eNodeB. The message includes the PCI of the cell in which the connection failed, the CGI of the cell where the radio link was re-established and the identity (C-RNTI¹³) that the UE had in the cell where the connection failed. The message may also include the short MAC-I bits¹⁴ to allow verification of UE identity at the source eNodeB: the source eNodeB can recalculate these bits using the security configuration of the source cell. This mechanism makes it possible to eliminate false RLF detections that could happen due to PCI collision, for example. If a PCI collision occurs, an eNodeB would typically send the ‘RLF INDICATION’ message to multiple neighbour eNodeBs which all control cells with the same PCI, but the use of the short MAC-I bits enables only the source cell in which the RLF had really occurred to concern itself about the too-late handover.

25.6.2 Coverage Hole Detection

Receiving the ‘RLF INDICATION’ message at a source eNodeB might be caused either by a connection failure due to a too-late handover from the source to the target cell or by a coverage hole at the edge of the source cell. In order to improve the operation of the ‘too-late handover’ counter in the source eNodeB, it is useful to discriminate between these two scenarios. Hence, after completing the re-establishment procedure in the eNodeB where the connection is re-established, the UE may indicate to this eNodeB whether it can report the last cell measurements it has performed before the connection failed. If this is possible, this eNodeB may retrieve an ‘R9 UE RLF Report’ measurement and add it into the ‘RLF INDICATION’ message sent to the source eNodeB. By evaluating the measurements

¹³Cell-Radio Network Temporary Identifier.

¹⁴The 16 least-significant bits of the Message Authentication Code for Integrity (MAC-I, see Section 4.2.3).




Figure 25.4: Example of signal strength fluctuation in the source and target cells during handover.

contained in this report, the source eNodeB should be able to determine if the connection failure was actually due to a coverage hole or not.

25.6.3 Too-Early Handover

When a handover is executed too early, the UE may experience a connection failure in the target cell due to fluctuations of the radio signal in cell B as illustrated in Figure 25.4, and will try to re-establish the radio link in the original source cell. Since the source eNodeB would have deleted all contexts related to that UE when the handover was completed, the re-establishment in the source cell would result in the source eNodeB sending an 'RLF INDICATION' message to the target eNodeB, exactly as if a too-late handover had taken place in the opposite direction. However, the target eNodeB can still discriminate the too-early handover from a too-late handover in the other direction if it had previously started a timer T_{early} on the incoming handover of the UE and notices that the 'RLF INDICATION' message is received from the same source cell and for the same UE before T_{early} expires. The implementation of the SON remedy for too-early handover therefore requires both a new timer to be run in the target eNodeB and a new message, 'HANDOVER REPORT', to be sent from the target eNodeB B to the source eNodeB in response to the 'RLF INDICATION' message, to signal to the source eNodeB that the source cell had a too-early handover problem towards the target cell rather than a 'too-late' handover problem.

These examples also show that implementations of SON remedies for too-late and too-early handover are best implemented in conjunction with each other as a comprehensive solution in order to avoid wrong MRO verdicts.

25.6.4 Handover to an Inappropriate Cell

Handover to an inappropriate cell shares some similarities with the too-early handover described above. Consider an example where a UE moves from a cell A of eNodeB A to

a cell B of eNodeB B, but traverses a small overlap area in which a cell C of eNodeB C presents the strongest signal. This scenario is illustrated in Figure 25.5.




Figure 25.5: Handover to an inappropriate cell.

In this scenario, eNodeB A may be tempted to trigger the handover to cell C, but if the period for which cell C is the strongest is too short this could result in the UE experiencing a connection failure in cell C and re-establishing the radio link in the real suitable target cell B. However, the problem can again easily be reported to the originating eNodeB A: When the UE re-establishes the radio link in cell B, eNodeB B will send an 'RLF INDICATION' message to eNodeB C for that UE before the timer T_{early} expires in cell C, indicating the originating cell A. Then eNodeB C can send a 'HANDOVER REPORT' message to eNodeB A to inform it about the inappropriate handover to cell C. eNodeB A can then take appropriate corrective action if it receives multiple similar reports. The 'HANDOVER REPORT' message can also be sent from eNodeB C to eNodeB A over the X2 interface, even when both cells B and C belong to the same eNodeB and the 'RLF INDICATION' was thus internal to that eNodeB.

It can be noticed that the same 'HANDOVER REPORT' message is used to report the two MRO verdicts of too-early handover and handover to inappropriate cells, an information element 'type of handover problem' is used within the message to discriminate between these two cases.

25.6.5 MRO Verdict Improvement

As mentioned above, discrimination between a 'too-late handover' and a 'coverage hole' MRO verdict requires that the eNodeB where the UE re-establishes the connection after

the connection failure wait for the completion of the re-establishment and retrieve some measurements from the UE (the ‘R9 UE RLF Report’) before it can send the ‘RLF INDICATION’ message. As a result, this message may arrive late and hence confuse the MRO verdict between too-late, too-early and wrong cell handover. In order to remove this dependency on the timing of the network delivery of the ‘RLF INDICATION’ message, the ‘R9 UE RLF Report’ was extended in Release 10 to include the ‘Time elapsed since the last handover initialization until connection failure’. Since this timer runs in the UE, its value is independent of network delivery conditions and therefore more accurate. When receiving the ‘RLF INDICATION’ message, the source eNodeB will compare this new timer (instead of the Release 9 timer T_{early} described above) to an internal threshold and thus obtain a more accurate verdict between the three MRO types (too-late, too-early and wrong cell).

25.6.6 Handover to an Unprepared Cell

Release 10 further extends the scope of MRO to cover the case where, after a connection failure, the UE is in idle mode when it re-connects to the network, either because the RLF took too long or simply because the UE first tried to re-establish the connection on a target cell that was unprepared and hence the re-establishment failed. In such scenarios, a network context cannot be relied upon, and therefore the Release 9 MRO detection scheme cannot be used. Therefore Release 10 includes an equivalent context-less MRO solution:

1. A Release 10 UE starts a log when connection failure occurs, containing the following items:
 - The information contained in the Release 9 UE RLF report;
 - The new timer defined for the MRO verdict improvement (see Section 25.6.5), which measures the time elapsed since the last handover initialization until the connection failure;
 - The identity of the last cell that served the UE;
 - The identity of the cell which received the first (unsuccessful) re-establishment attempt;
 - The identity of the cell that served the UE at the last handover initialization.
2. The UE starts this log at connection failure detection and delivers it when it re-connects to the network from idle mode over the newly established Radio Resource Control (RRC) connection. By analysing the UE’s log, the network can obtain all the elements to make a verdict between too-late handover, too-early handover, wrong-cell handover or coverage hole, despite the fact that the UE had transitioned through idle mode.

25.6.7 Unnecessary Inter-RAT Handovers

MRO was also extended in Release 10 to cover inter-RAT scenarios. The unnecessary inter-RAT handover scenario addresses the case where a UE is handed over from LTE to another RAT (e.g. UMTS or GSM) despite the fact the quality of the LTE coverage was actually good enough and would have allowed the UE to remain on LTE. This can be seen as a too-early handover to another RAT with no connection failure; it may affect the user experience if the

other RAT offers lower performance and the UE cannot easily reconnect to LTE; it may also lead to a non-optimal use of network resources. It is typically due to the handover threshold for the LTE serving cell being set too high.

Such unnecessary inter-RAT handovers are detected by the target RAN node when triggered by the source eNodeB. The source eNodeB first requests the target RAN node during the inter-RAT Handover Preparation procedure to instruct the UE to continue measuring the LTE cells. It also provides the duration for which these measurements should continue and a minimum reference quality above which the target RAN node is allowed to report. It may also provide a list of selected LTE frequencies, for example to minimize the impact of compressed mode while in UMTS. Then, whenever the target RAN node receives a UE measurement above this threshold for one LTE cell for a sufficient period of time, it marks this LTE cell as to be reported. The reporting phase is performed by the target RAN node: when the measurement duration ends, the target RAN node sends a report (using a generic SON transfer container on top of the RIM protocol) to the originating LTE cell containing the results of the measurements – i.e. the list of LTE cells whose quality the UE measured to be above the configured threshold. Upon reception of the report, the originating eNodeB can take further corrective action, such as deciding if and how its parameters (e.g. the threshold to trigger inter-RAT handover) should be adjusted.

25.6.8 Potential Remedies for Identified Mobility Problems

Whenever the mechanisms described above are used in a network to detect such connection failure scenarios between cells, the source eNodeB may take some appropriate actions to combat them. For instance, in the case of repeated too-late or too-early handovers, one possible action could be to change the threshold at which the handovers are triggered. The source eNodeB can then use the same ‘Mobility Settings Change’ procedure as for MBL scenarios to inform the neighbour eNodeB of the change of threshold it has performed. In this case the ‘MOBILITY CHANGE REQUEST’ message is sent with a cause value ‘handover optimization’ and indicates the change (in dB) of the handover trigger parameter change performed in the source cell. For instance, in the case of repeated handover to an inappropriate cell, the source eNodeB could modify the way it builds its list of candidate target cells.

25.7 Random Access Channel (RACH) Self-Optimization

Optimization of the RACH configuration in cells is a Release 9 SON feature that is key to optimizing the system performance of a mobile network. A poorly configured RACH may result in higher call setup and handover delays due to frequent RACH collisions, or low preamble-detection probability and limited coverage. The amount of uplink resource reserved for RACH also affects the system capacity. Therefore a network operator should carefully monitor that the RACH parameters are appropriately set, taking into account factors such as the RACH load, the uplink interference, the traffic patterns and the population under the cell coverage. The task becomes more complicated given that these factors may change dynamically. For example, if the antenna tilt is changed in a cell, it will affect the rates of call arrival and handover in this cell and the surrounding cells, and therefore the RACH load per

preamble in all those cells. A change in transmission power settings or handover thresholds may have similar effects.

Whenever such a network configuration change happens, the RACH self-optimization feature should automatically make appropriate measurements of the RACH performance and usage in all the affected cells and determine any necessary updates of the RACH parameters. Some useful measurements are UE reports of the number of RACH attempts needed to obtain access, or time elapsed from the first attempt until access is finally granted. The reports may use RRC signalling or MAC¹⁵ control elements depending on how much data is to be transmitted and on reliability and timing requirements. RACH parameters that can then be adjusted are typically the split of RACH preambles between contention-free access, contention-based access with high payload and contention-based access with low payload (see Section 17.3), the RACH back-off parameter value or the RACH transmission power ramping parameters.

In addition, the RACH self-optimization feature facilitates automatic configuration of the Physical RACH (PRACH) parameters (including the PRACH resource configuration, preamble root sequence and cyclic shift configuration – see Sections 17.4.2.5 and 17.4.3) to avoid preamble collisions with neighbouring cells. The principle of this automatic configuration is similar to the automatic PCI configuration SON feature described in Section 25.4: the PRACH configuration information is included in the ‘X2 Setup’ and ‘eNB Configuration Update’ procedures. Therefore, whenever a new eNodeB is initialized and learns about its neighbours via the ANR function, it can at the same time learn the neighbouring PRACH configurations. It can then select its own PRACH configuration to avoid conflicts with the neighbouring ones.

Whenever a conflict is identified, one of the cells should change its configuration, but the algorithm for selecting which cell should change and in what manner is not specified. The network operator can also combine PRACH self-optimization with manual configuration if necessary, but this is typically more prone to errors and more time consuming than automatic RACH optimization.

25.8 Energy Saving

Energy saving is a feature that aims at cost savings, as well as reduced environmental impact. Energy savings in the RAN typically focus on automatically adapting the capacity offered by the network to the actual traffic demand at a given point in time. Examples include radio solutions, such as not scheduling transmissions in certain subframes, and network solutions, such as switching off cells, transceivers or antennas.

Such a radio solution was introduced in Release 9 for LTE heterogeneous networks deploying LTE capacity-booster cells on top of cells that give wide-area coverage. The function enables the capacity-booster cells to be switched off when their capacity is no longer needed, and to be re-activated from a dormant state on a per-need basis.

The eNodeBs managing the capacity-booster cells operate an algorithm by which they decide autonomously whether to switch off a cell or not. The algorithm may take into account

¹⁵Medium Access Control.

various implementation-dependent criteria, such as some knowledge of the traffic load in the overlaid macrocell.¹⁶

Before switching off a capacity-booster cell, the eNodeB may initiate handover actions in order to offload any users in the cell; in this case it indicates the reason for handover to facilitate the target eNodeB in taking appropriate subsequent action, such as not selecting the switched-off cell as a target for subsequent handovers.

After switch-off, all peer eNodeBs are informed via the X2 interface by means of the eNB Configuration Update procedure, and they are then assumed to maintain the switched-off cell's configuration data while it is in the dormant state.

These peer eNodeBs may request a re-activation of the switched-off cell via the X2 interface if they need additional capacity. The algorithm that determines whether and when to request re-activation is implementation-dependent; it would typically take into account the current load condition in the macrocell(s).¹⁷ The eNodeB managing the dormant capacity-booster cell would normally be expected to obey such a request.

After switch-on, all peer eNodeBs are again informed of the change of status of the capacity-booster cell via the X2 interface.

Importantly, the energy-saving solution described here is assumed to be under close control of the network operator, as it should have no impact on the basic coverage. Therefore it is assumed that the operator would first configure the eNodeBs that are allowed to decide to switch cells off, together with associated policies, and then configure the eNodeBs that are allowed to request re-activation of a list of dormant cells.

25.9 Emerging New SON Use Cases

Further enrichment of the LTE SON functionality is expected to include the following aspects:

- Further improvements related to mobility load balancing and mobility robustness;
- SON features for Home eNodeBs;
- SON features for Relay Nodes.

For example, further improvements of Mobility Load Balancing and Mobility Robustness in Release 11 may comprise:

- detection and prevention of intra-LTE rapid handovers – i.e. too-early handovers that do not lead to connection failures;
- detection and prevention of inter-RAT too-late handovers (which lead to connection failures);
- exchange of mobility parameters between RATs to optimize inter-RAT handover (e.g. to combat unnecessary inter-RAT handovers).

¹⁶The decision to switch off a cell can also be taken by O&M.

¹⁷Like the decision to switch off, the decision to re-activate may also be taken by O&M.

References¹⁸

- [1] 3GPP Technical Specification 36.413, 'Evolved Universal Terrestrial Radio Access Network (E-UTRAN); S1 Application Protocol (S1AP)', www.3gpp.org.
- [2] 3GPP Technical Specification 36.423, 'Evolved Universal Terrestrial Radio Access Network (E-UTRAN); X2 Application Protocol (X2AP)', www.3gpp.org.
- [3] 3GPP Technical Specification 32.511, 'Automatic Neighbour Relation (ANR) management; Concepts and Requirements', www.3gpp.org.
- [4] 3GPP Technical Specification 32.521, 'Self-Organizing Networks (SON) Policy Network Resource Model (NRM) Integration Reference Point IRP; Requirements', www.3gpp.org.

¹⁸All web sites confirmed 1st March 2011.

LTE System Performance

Tetsushi Abe

26.1 Introduction

The system performance requirements set for LTE Release 8 in [1] (summarized in Table 1.1) demanded substantial improvements over Release 6 of UMTS.¹ For the downlink, the average cell spectral efficiency was required to be three to four times higher than HSDPA,² and the cell edge spectral efficiency two to three times higher. For the uplink, average and cell edge spectral efficiencies two to three times higher than those of HSUPA³ were required.

This chapter first summarizes the main technical features of LTE Release 8 that deliver these substantial system capacity enhancements. Evaluation results are then presented for spectral efficiency and coverage for typical LTE macrocell deployment scenarios, known as ‘Case 1’ and ‘Case 3’ [2] with inter-site distances of 500 m and 1732 m respectively, together with system performance evaluations in test environments defined by the International Telecommunication Union Radiocommunication Sector (ITU-R) [3].

26.2 Factors Contributing to LTE System Capacity

The preceding chapters have extensively explained the state-of-the-art technical features that significantly improve the system performance of LTE compared to legacy systems. The main difference in both downlink and uplink between LTE Release 8 and UMTS Release 6 (HSDPA/HSUPA) is that the LTE system provides orthogonal resource allocation in the frequency domain, which enables frequency-domain multi-user diversity gain to be exploited. In addition, the LTE downlink supports transmission with up to two or four spatial layers via

¹Universal Mobile Telecommunications System.

²High Speed Downlink Packet Access.

³High Speed Uplink Packet Access.

LTE – The UMTS Long Term Evolution: From Theory to Practice, Second Edition.

Stefania Sesia, Issam Toufik and Matthew Baker.

© 2011 John Wiley & Sons, Ltd. Published 2011 by John Wiley & Sons, Ltd.

multiple antennas, which enhances both the peak data rate and the cell average and cell edge spectral efficiencies. The key features are discussed in more detail in the following sections.

26.2.1 Multiple Access Techniques

Downlink OFDMA. As explained in Chapter 5, the LTE downlink is based on Orthogonal Frequency Division Multiple Access (OFDMA) which enables flexible channel-dependent multi-user resource allocation in both the frequency and time domains as illustrated in Figure 5.12. This leads to improved multi-user diversity gain. Inter-Symbol Interference (ISI) reduction by means of the Cyclic Prefix (CP) leads to a simplified receiver structure, which is well suited to Multiple-Input Multiple-Output (MIMO) transmission.

Uplink SC-FDMA. The Single-Carrier Frequency Division Multiple Access (SC-FDMA) scheme used for the LTE uplink (as explained in Chapter 14) achieves frequency-domain intra-cell orthogonality among User Equipment (UEs) while also maintaining a low Peak-to-Average Power Ratio (PAPR) which is important for maximizing data rates at the cell edge. In addition, the Sounding Reference Signals (SRSSs) supported by the LTE uplink (see Section 15.6) facilitate multi-user scheduling and rate adaptation strategies to enhance spectral efficiency.

26.2.2 Frequency Reuse and Interference Management

Similarly to WCDMA,⁴ LTE is designed to operate with a frequency reuse factor of one to maximize the spectral efficiency. In such a system, however, data and control channels can experience a significant level of interference from neighbour cells, which reduces the achievable spectral efficiency, especially at the cell edge. LTE therefore supports various techniques to manage and mitigate inter-cell interference.

In the downlink, these include:

- A cell-specific frequency-shift is applied to the mapping of cell-specific Reference Signals (RSs) to subcarriers to avoid inter-cell RS collisions, as described in Section 8.2.1.
- With respect to the downlink control channels, a cell-specific frequency offset is applied to the PCFICH and PHICH⁵ positions, as described in Sections 9.3.3 and 9.3.4; for the Physical Downlink Control CHannel (PDCCH), interleaving provides frequency diversity and enhances the robustness against inter-cell interference (see Section 9.3.5.1).
- For the data, Inter-Cell Interference Coordination (ICIC) techniques can be applied by utilizing the Relative Narrowband Transmit Power (RNTP) messages that can be exchanged among eNodeBs over the X2 interface as explained in Section 12.5.

Uplink interference mitigation techniques in LTE include the following:

- Fractional power control (see Section 18.3.2.1) is supported to improve the throughput near the eNodeB and mitigate inter-cell interference at the cell edge. Power control

⁴Wideband Code Division Multiple Access.

⁵Physical Control Format Indicator CHannel and Physical HARQ Indicator CHannel.

can be performed jointly with frequency-domain resource allocation, whereby cell-centre UEs are allocated more Resource Blocks (RBs) to enhance the data rate, while cell-edge UEs are allocated fewer RBs for coverage extension. When devising power control strategies for interference management, it is important to control the ratio of the Interference over Thermal noise (IoT) below a target level.

- Various means are provided to avoid inter-cell RS collisions, including cyclic shift hopping, sequence-group hopping and planning, as explained in Sections 15.3 and 15.4.
- For the Physical Uplink Control CHannel (PUCCH), cell-specific symbol-level cyclic shift hopping is applied for inter-cell interference randomization, as explained in Section 16.3.3. Each PUCCH RB pair is mapped to both edges of the system bandwidth to achieve frequency diversity. In addition, the fact that the PUCCH is mapped to different RBs in the frequency domain from those of the Physical Uplink Shared CHannel (PUSCH, see Section 16.3.1) means that independent interference management techniques (such as power control) can be applied for the control and data channels.
- Uplink ICIC techniques can be applied using the Overload Indicator (OI) and High Interference Indicator (HII) messages that can be exchanged between eNodeBs over the X2 interface as explained in Section 12.5.

Scrambling is applied to the data and control channels, and to the downlink RSs, to randomize inter-cell interference.

26.2.3 Multiple Antenna Techniques

Downlink Spatial Multiplexing and Diversity. The multiple antenna schemes supported by LTE contribute much to the overall spectral efficiency gain with respect to Release 6 of UMTS. The performance of the open-loop and closed-loop spatial multiplexing modes supported in LTE is shown in Section 11.2.4.

Uplink Multi-User MIMO. In the LTE uplink, as explained in Section 16.6.2, orthogonal demodulation RSs can be assigned to multiple UEs, thus enabling the eNodeB receiver to estimate the channel of multiple UEs scheduled simultaneously to transmit in the same set of RBs.

26.2.4 Semi-Persistent Scheduling

The Semi-Persistent Scheduling (SPS) provided in LTE (see Section 4.4.2.1) can alleviate pressure on the limited downlink control channel capacity by replacing dynamic scheduling signalling with semi-static signalling. This allows a larger number of UEs to be scheduled, which is especially beneficial for services such as Voice over IP (VoIP).

26.2.5 Short Subframe Duration and Low HARQ Round Trip Time

LTE has a subframe duration of 1 ms for both uplink and downlink – shorter than the 2 ms subframe duration of UMTS. This leads to reduced latency (with a shorter HARQ⁶ Round Trip Time (RTT)) and more flexible multi-user scheduling in the time domain.

26.2.6 Advanced Receivers

Advanced receivers provide an implementation method to enhance further the capacity of the LTE system. A typical example is the Linear Minimum Mean Squared Error (LMMSE) receiver with Interference Rejection Combining (IRC) [4]. Suitable for both uplink and downlink, such receivers compute the signal combining weights by exploiting statistical knowledge, such as the covariance matrix, of the inter-cell interference (unlike Maximum Ratio Combining (MRC) receivers which do not take the spatial characteristics of the interference into account). The ability of an IRC receiver to suppress interference is a function of many factors including the number and strength of the interfering signals and the number of receive antennas. For a single dominant source of interference, the mean Signal-to-Interference Ratio (SIR) gain that IRC is able to offer relative to an MRC receiver improves with the interference-to-thermal-noise ratio and not with the wanted Signal-to-Noise Ratio (SNR). However, if significant interference arrives from more sources or directions, an IRC receiver with only a small number of antennas is limited in the amount of interference suppression it can provide, especially if the multiple interference sources are received with similar powers.

Other more advanced receiver structures may also be considered, such as MMSE receivers with Successive Interference Cancellation (SIC) (particularly for downlink Single-User MIMO (SU-MIMO)) and Maximum Likelihood Detection (MLD) (see Section 11.1.3.3).

26.2.7 Layer 1 and Layer 2 Overhead

Any part of the time-frequency transmission resources that are not used directly for data transmission constitutes an overhead when considering the overall spectral efficiency. One design criterion for LTE was to minimize these overheads while achieving high system performance and flexibility. Table 26.1 summarizes the major sources of Layer 1 and Layer 2 overhead in the LTE downlink transmissions in a 10 MHz Frequency Division Duplex (FDD) deployment, as a percentage of the total transmission resources over the duration of a 10 ms radio frame.⁷ The main contributors are guard bands, the OFDM CP, RSs, and control channels. It can be seen that the percentage overhead increases with the number of transmit antennas, due to the higher RS overhead for the larger number of antenna ports. The gain from MIMO therefore has to more than offset this increased overhead if it is to be worthwhile.

⁶Hybrid Automatic Retransmission reQuest.

⁷Note that transmission mode 7 (see Section 9.2.2.1) is not considered in these examples, and therefore the overhead due to UE-specific RSs is not taken into account. A worst-case downlink control channel duration of three OFDM symbols per subframe is used here, although in practice LTE supports dynamic resource allocation for the downlink control channels, so the average control channel duration would depend on the deployment scenario.

Table 26.1: Examples of percentage overhead in the LTE FDD downlink (calculated over a 10 ms radio frame for a 10 MHz system bandwidth).

Source of overhead	1 antenna port	2 antenna ports	4 antenna ports	Illustration
Guard bands (1 MHz)	10.0	10.0	10.0	Figure 21.1
OFDM CP	6.0	6.0	6.0	Figure 5.13
Cell-specific RSs	4.0	8.0	12.0	Figures 8.2 & 8.3
Control channels (3 symbols)	17.0	16.0	14.0	Figure 9.5
Synchronization signals	0.29	0.29	0.29	Figure 7.4
PBCH	0.28	0.26	0.24	Figure 9.1
Total (%)	37.6	40.6	42.6	

Table 26.2 similarly summarizes the major sources of Layer 1 and Layer 2 overhead in the LTE FDD uplink. Guard bands, CP, demodulation RSs and PUCCH constitute the main sources of uplink overhead.⁸

Table 26.2: Examples of percentage overhead in the LTE FDD uplink (calculated over a 10 ms radio frame for a 10 MHz system bandwidth).

Source of overhead	Overhead (%)	Illustration
Guard bands (1 MHz)	10.0	
SC-FDMA CP	5.9	Table 14.1
PUSCH demodulation RSs (2 symbols per subframe)	6.0	Figure 15.7
PUCCH (4 RBs)	6.7	Figure 16.2
RACH (6 RBs, 10 ms period)	1.1	Figure 17.5
SRS (48 RBs bandwidth, 10 ms period)	0.55	Table 15.1
Total (%)	30.3	

26.3 LTE Capacity Evaluation


In this section the capacity of LTE Release 8 is evaluated for various deployment scenarios. Table 26.3 shows the key parameters for the deployment scenarios designated ‘Case 1’ and ‘Case 3’ according to [2, 5].

Key parameters for the models defined by ITU-R are given in Table 20.8. Four ITU-R test environments are defined, each of which may be represented by one or more typical deployment scenarios as described in Section 20.3.6.1. The evaluations presented in this chapter focus on the Indoor Hotspot (‘InH’), Urban Microcell (‘UMi’), Urban Macrocell (‘UMa’) and Rural Macrocell (‘RMa’) deployment scenarios [3].


⁸The overhead due to the Random Access CHannel (RACH) depends on its configured transmission period, and that of the SRSs depends on their configured period and bandwidths.

Table 26.3: Key parameters for 3GPP Case 1 and Case 3 models.

	Case 1	Case 3
Deployment scenario	Macrocell	Macrocell
Network layout	Hexagonal grid	Hexagonal grid
Channel model	3GPP Spatial Channel Model [6]	3GPP Spatial Channel Model [6]
Inter-Site Distance (ISD)	500 m	1732 m
Carrier frequency	2 GHz	2 GHz
eNodeB antenna height	32 m	32 m
eNodeB transmit power	43 dBm (1.25 MHz, 5 MHz) 46 or 49 dBm (10 MHz, 20 MHz)	43 dBm (1.25 MHz, 5 MHz) 46 or 49 dBm (10 MHz, 20 MHz)
UE transmit power	23 or 24 dBm	23 or 24 dBm
UE speed	3 km/h	3 km/h



(a) Average spectral efficiency



(b) Cell-edge spectral efficiency

Figure 26.1: Ratio of Release 8 LTE FDD and HSDPA downlink spectral efficiencies for the Case 1 and Case 3 deployment scenarios.

26.3.1 Downlink and Uplink Spectral Efficiency

Figure 26.1 shows the average and cell-edge spectral efficiencies of the Release 8 LTE FDD downlink, as a ratio relative to Release 6 HSDPA for the 3GPP Case 1 and Case 3 deployment scenarios. Here, the cell-edge user throughput is defined as the 5-percentile user throughput. The LTE performance assumes 2x2, 4x2 and 4x4 antenna configurations with closed-loop MIMO, while 1x2 is assumed for HSDPA. Antenna separations at the eNodeB and UE are assumed to be ten wavelengths and half a wavelength respectively; eNodeB antenna tilting is not assumed.

The results show that the 2x2 configuration satisfies the LTE performance requirement of three to four times higher average spectral efficiency and two to three times higher cell-edge spectral efficiency than HSDPA. This substantial gain is mainly attributable to frequency-domain multi-user scheduling and dual-stream MIMO transmission. The additional gain from a 4x2 antenna configuration can also be seen, justifying the increased RS overhead.

Figure 26.2 shows the average and cell-edge spectral efficiencies of the Release 8 LTE FDD uplink compared to Release 6 HSUPA. For LTE a 1x2 or 1x4 antenna configuration is assumed, and 1x2 for HSUPA.




Figure 26.2: Ratio of Release 8 LTE FDD and HSUPA uplink spectral efficiencies for the Case 1 and Case 3 deployment scenarios.

The results show that the 1×2 LTE configuration satisfies the performance requirement of two to three times higher average and cell-edge spectral efficiencies than HSUPA. This is mainly attributable to intra-cell orthogonality and frequency-domain multi-user scheduling of LTE. The figures clearly show the effect of increasing the number of receive antennas.

Figure 26.3 shows the downlink average and cell-edge spectral efficiencies for Release 8 LTE FDD 4×2 downlink SU-MIMO transmission in the ITU-R deployment scenarios.⁹ Vertical antenna tilting is assumed at the eNodeB, with tilt angles of 12, 12 and 6 degrees in the urban microcell, urban macrocell, and rural macrocell scenarios respectively [5]. The antenna separation is assumed to be four wavelengths at the eNodeB and half a wavelength at the UE.




Figure 26.3: Downlink spectral efficiency of a Release 8 LTE FDD system with 4×2 SU-MIMO in the ITU-R deployment scenarios.

Figure 26.4 shows an example of the Cumulative Distribution Function (CDF) of average received downlink Signal-to-Interference-plus-Noise Ratio (SINR) for each of the ITU-R deployment scenarios. The higher performance in the indoor hotspot scenario compared with the others can be attributed to the isolated cell environment with low inter-cell interference; this is especially advantageous for high-order MIMO spatial multiplexing.

Figure 26.5 shows the uplink average and cell-edge spectral efficiencies for Release 8 LTE FDD with 1×4 Single-Input Multiple-Output (SIMO) and Multi-User MIMO (MU-MIMO) transmissions in the ITU-R deployment scenarios. These results are partially published in [7].

⁹A downlink control channel duration of three OFDM symbols is assumed. These results are partially published in [7].




Figure 26.4: Downlink received SINR distribution for the ITU-R deployment scenarios.

The same antenna separations and eNodeB vertical antenna tilt angles are assumed as in the downlink evaluation.




Figure 26.5: Uplink spectral efficiency of a Release 8 LTE FDD system with 1x4 SIMO and MU-MIMO in the ITU-R deployment scenarios.

Figure 26.6 shows an example of the CDF of average received uplink SINR for each deployment scenario. Similarly to the downlink, these results show that the indoor scenario provides higher throughput performance compared to the other scenarios, due to the higher SINR distribution. The potential gain of MU-MIMO can also be seen.




Figure 26.6: Uplink received SINR distribution for the ITU-R deployment scenarios.

26.3.2 VoIP Capacity

Figure 26.7 shows the VoIP capacity for the 3GPP and ITU-R deployment scenarios.

The VoIP capacity is typically evaluated in terms of the number of users per cell. For the ITU-R scenarios, VoIP capacity is defined as the minimum of the downlink and uplink capacities. The results here assume the 12.2 kbps codec with a 50% voice activity factor. A VoIP user is considered to be in outage if the 98-percentile radio interface latency is greater than 50 ms. For the 3GPP Case 1 and 3 evaluations, 2×2 and 1×2 antenna configurations are assumed for downlink and uplink respectively, while for the ITU-R scenarios 4×2 or 1×2 antenna configuration are assumed for the downlink and 1×4 for the uplink. The eNodeB vertical antenna tilting assumptions are the same as in the spectral efficiency evaluations above. Application of SPS to VoIP helps to avoid any limitation on the PDCCH capacity, thus enabling a high number of VoIP users to be supported in both uplink and downlink.

26.4 LTE Coverage and Link Budget

Figure 26.8 shows the uplink and downlink coverage for the 3GPP and ITU-R deployment scenarios.¹⁰

The PDCCH coverage in the downlink is estimated on the basis of a 44-bit Downlink Control Information (DCI) message with 8 Control Channel Elements (CCEs) (see Sections 9.3.5 and 9.3.5.1), while the PUCCH coverage in the uplink is based on a 4-bit CQI report (see

¹⁰The values for the ITU-R deployment scenarios are published in [7, Annex C2].




Figure 26.7: VoIP capacity of LTE Release 8 FDD.




Figure 26.8: LTE coverage (10 MHz, FDD).

Sections 10.2.1 and 16.3). A 2×2 antenna configuration is assumed in the downlink and 1×4 in the uplink. The eNodeB antenna gain is 17 dBi, and the total cable loss at the transmitter and receiver is 4 dB for both uplink and downlink. Noise figures of 7 dB and 5 dB are assumed for the downlink and uplink respectively, and a receiver implementation margin of 2 dB. The thermal noise density is -174 dBm (see Section 21.4.4.2). Propagation models for each deployment scenario are discussed in Section 20.3.6.1, and full details can be found in [5]. ‘Coverage’ is defined here as 95% successful reception for the control channels and 90% for the data channels. The target Block Error Rates (BLERs) are 1% for the control channels and 10% for the data channels. A 0.5 dB HARQ gain is assumed for the data channels. In the ITU-R deployment scenarios, the coverage exceeds the cell radius in each case. In 3GPP Case 3, PCFICH power boosting and narrowing the PUSCH transmission bandwidth would further improve the coverage of the PDCCH and PUSCH respectively.

Table 26.4 shows the link budget calculation in detail for Case 3. A UE transmit and receive antenna gain of 0 dBi is assumed, penetration loss of 20 dB, downlink receiver interference over thermal noise ratio of 1 dB for both data and control channels, uplink receiver interference over thermal noise ratio of 1.5 dB for data and 4.5 dB for control channels, and log-normal shadow fading margins of 10.5 dB for control and 6.7 dB for data. The same SINR threshold values as in the ITU-R rural macrocell deployment scenario are assumed. Full details of the link budget calculations for the ITU-R deployment scenarios can be found in [7, Annex C2].

Table 26.4: LTE link budget for Case 3 scenario with Non-Line-of-Sight (NLOS) propagation channel.

	Downlink		Uplink	
	PDCCH (42 kbps)	PDSCH (2 Mbps)	PUCCH (4 kbps)	PUSCH (187 kbps)
Bandwidth (RBs)	50	50	1	5
Coverage (m)	945	1053	898	668
Tx power (dBm)	46	46	24	24
Rx sensitivity* (dBm)	-98.7	-96.6	-119.8	-111.2
Interference + Noise (dBm/Hz)	-166	-166	-165	-168
Required SINR (dB)	-4.2	-1.7	-9.9	-4.8

* See Chapter 21.

26.5 Summary

This chapter reviews the main technical features of the LTE downlink and uplink that influence the system capacity. In particular, the role of the orthogonal multiple access schemes OFDMA and SC-FDMA is highlighted, together with downlink MIMO, in achieving significant capacity improvements over the Release 6 HSDPA and HSUPA.

Extensive performance evaluations for LTE have been carried out in 3GPP, and key results are presented here for cell average and cell-edge spectral efficiencies, VoIP capacity, and

coverage. Considering a range of deployment scenarios defined by both 3GPP and the ITU-R, it is shown that the performance requirements set in [1] and summarized in Table 1.1 for the first version of LTE (i.e. Release 8) are satisfied.

The LTE system performance is further enhanced in Release 10 for LTE-Advanced, as explained in Chapters 27 to 32.

References¹¹

- [1] 3GPP Technical Report 25.913, 'Requirements for Evolved UTRA (E-UTRA) and Evolved UTRAN (E-UTRAN) (Release 7)', www.3gpp.org.
- [2] 3GPP Technical Report 25.814, 'Physical Layer Aspects for Evolved Universal Terrestrial Radio Access (UTRA) (Release 7)', www.3gpp.org.
- [3] ITU-R Report M.2135, 'Guidelines for Evaluation of Radio Interface Technologies for IMT-Advanced', www.itu.int/itu-r.
- [4] J. H. Winters, 'Optimum Combining in Digital Mobile Radio with Cochannel Interference'. *IEEE Journal on Selected Areas in Communications*, Vol. 2, pp. 528–539, July 1984.
- [5] 3GPP Technical Report 36.814, 'Further Advancements for E-UTRA Physical Layer Aspects (Release 9)', www.3gpp.org.
- [6] 3GPP Technical Report 25.996, 'Spatial Channel Model for Multiple Input Multiple Output (MIMO) simulations (Release 8)', www.3gpp.org.
- [7] 3GPP Technical Report 36.912, 'Feasibility Study for Further Advancements for E-UTRA (LTE-Advanced) (Release 9)', www.3gpp.org.

¹¹All web sites confirmed 1st March 2011.

Part V

LTE-Advanced

Introduction to LTE-Advanced

Dirk Gerstenberger

27.1 Introduction and Requirements

With the completion of LTE Release 8, 3GPP started to look into ways to further evolve LTE for the future, in order to build upon the existing LTE technology and to ensure that LTE remains the leading global standard for mobile broadband.

Enhanced performance can in principle be achieved in two ways – by using more radio spectrum, and by using the available spectrum more efficiently. The International Telecommunication Union (ITU) has taken steps to ensure that more radio spectrum will be available, globally whenever possible, for systems beyond the 3rd Generation. The World Radiocommunication Conference (WRC) 2007 resulted in some new spectrum bands being earmarked for mobile services. In order to satisfy the perceived needs and to ensure that effective use is made of spectrum allocations, in March 2008 the ITU Radiocommunication Sector (ITU-R) issued a ‘circular letter’ [1] calling for submission of candidate Radio Interface Technologies (RITs). Successful proposals would fulfil the ITU-R’s requirements for *IMT-Advanced*¹ [2].

Key features of IMT-Advanced set out in the circular letter are:

- A high degree of commonality of functionality worldwide, while retaining the flexibility to support a wide range of services and applications in a cost-efficient manner;
- Compatibility of services within IMT and with fixed networks;
- Capability of interworking with other radio access systems;
- High quality mobile services, user equipment suitable for worldwide use, user-friendly applications, services and equipment, and worldwide roaming capability;

¹International Mobile Telecommunications - Advanced (following on from the IMT-2000 family of systems).

- Enhanced peak data rates to support advanced services and applications (100 Mbps for high mobility and 1 Gbps for low mobility were established as targets for research).

The key radio access requirements set by ITU-R for IMT-Advanced for different deployment scenarios are summarized in Table 27.1. Note that the spectral efficiency requirements in downlink and uplink are defined on a per-cell basis, and no explicit peak data rate requirements are defined.

Table 27.1: Key radio access requirements of IMT-Advanced for different deployment scenarios (see Figure 20.8 for details of the deployment scenarios).

Parameter	Downlink	Uplink
Maximum Bandwidth	At least 40 MHz	
Peak spectral efficiency (bps/Hz)	15	6.75
Average spectral efficiency (bps/Hz/Cell)	3.0 (Indoor Hotspot) 2.6 (Urban Micro) 2.2 (Urban Macro) 1.1 (Rural Macro)	2.25 (Indoor Hotspot) 1.8 (Urban Micro) 1.4 (Urban Macro) 0.7 (Rural Macro)
Cell-edge user spectral efficiency (bps/Hz)	0.1 (Indoor Hotspot) 0.075 (Urban Micro) 0.06 (Urban Macro) 0.04 (Rural Macro)	0.07 (Indoor Hotspot) 0.05 (Urban Micro) 0.03 (Urban Macro) 0.015 (Rural Macro)
VoIP capacity (user/cell/MHz)	50 (Indoor Hotspot) / 40 (Urban Micro and Urban Macro) / 30 (Rural Macro)	
User plane latency (ms)	10	
Control plane latency (ms)	100	

In response to the call for proposals from ITU-R, a workshop of 3GPP TSG RAN² took place in April 2008 to identify targets and potential techniques for further advancement of LTE, which led to a set of requirements being approved in May 2008 [3].

These targeted advancements became *LTE-Advanced*, specified as LTE Release 10 and beyond. LTE-Advanced is the 3GPP candidate RIT for IMT-Advanced. It is designed to meet the requirements of mobile network operators for the evolution of LTE, and to exceed the IMT-Advanced requirements.

3GPP's key radio-access targets for LTE-Advanced are outlined in Table 27.2. In addition, 3GPP set requirements on backward compatibility with earlier releases of LTE, as discussed in more detail in Section 27.3. This allows network operators to continue serving existing LTE customers while their network equipment is progressively upgraded. Requirements on spectrum deployment and flexibility, coexistence with legacy Radio Access Technologies (RATs), and complexity and service support were also defined [3].

²Technical Specification Group Radio Access Network – see Section 1.1.3.

Table 27.2: Key radio access targets for LTE-Advanced as set by 3GPP [4].

Parameter	Downlink	Uplink
Maximum Bandwidth	Up to 100 MHz	
Peak data rate (Mbps)	1000	500
Peak spectral efficiency (bps/Hz)	30	15
Average spectral efficiency (bps/Hz/Cell)	2.6 for 'Case 1'	2 for 'Case 1'
Cell-edge user spectral efficiency (bps/Hz)	0.09 for 'Case 1'	0.07 for 'Case 1'
VoIP capacity (user/cell/MHz)	Exceeding LTE Release 8	
User plane latency (ms)	10	
Control plane latency (ms)	50 (Idle to Active), 10 (Dormant to Active)	

3GPP's targets for LTE-Advanced were set independently from the IMT-Advanced requirements; it can be seen that some of the 3GPP targets exceed the IMT-Advanced requirements, such as the peak spectral efficiency and the control plane latency targets. This is related both to the fact that LTE Release 8 already fulfils many of the IMT-Advanced requirements (see Chapter 26), and to the fact that LTE-Advanced is not limited to LTE Release 10 but will also include new features in subsequent LTE releases. The ITU-R process called for complete descriptions of the IMT-Advanced candidates to be submitted by June 2009, with submission of the final details including a performance evaluation following by October 2009. 3GPP documented its submission in [5].

The 3GPP submission to ITU-R included an FDD³ and a TDD⁴ RIT component, which were developed with the goal of maximizing their commonality. Together, the FDD and TDD RITs comprise a 'Set of RITs' (SRIT). 3GPP's submission to ITU-R [6, 7] included detailed technology characteristics, link budget analysis and information about supported services, spectrum and technical performance.

An evaluation of LTE-Advanced was carried out by 18 companies in 3GPP, showing that LTE-Advanced completely satisfies the criteria set by the ITU-R for IMT-Advanced. The results of the evaluation are included in [7]. As a result, LTE-Advanced was accepted by the ITU as an IMT-Advanced technology in October 2010.

³Frequency Division Duplex.

⁴Time Division Duplex.

27.2 Overview of the Main Features of LTE-Advanced

The main components of LTE-Advanced that are added to LTE in Release 10 are:

- Carrier aggregation;
- Enhanced downlink multiple antenna transmission;
- Uplink multiple antenna transmission;
- Relaying;
- Support for heterogeneous network deployments.

Data rates of the order of 1 Gbps might theoretically be achieved using contiguous bandwidths of 40 MHz or more. However, competition for spectrum and fragmentation of the available spectrum makes it unrealistic to expect such large contiguous bandwidths in most cases. LTE-Advanced therefore makes use of *carrier aggregation* (see Chapter 28) to support such large bandwidths. This also has the advantages of limiting the cost of equipment and enabling much of the technology developed for LTE Release 8 to be reused. Each ‘component carrier’ within an aggregation is designed to be fundamentally similar to an LTE Release 8 carrier so that they can be configured in a backward-compatible way and used by legacy UEs if desired. Up to five component carriers with a bandwidth of up to 20 MHz each can be aggregated in LTE-Advanced to make efficient use of the available spectrum and achieve the desired total bandwidth and peak data rate. LTE-Advanced enables a variety of different arrangements of component carriers to be aggregated, including component carriers of the same or different bandwidths, adjacent or non-adjacent component carriers in the same frequency band, and component carriers in different frequency bands. The physical layer mechanisms for carrier aggregation are largely independent of the frequency location of the component carriers, but in order to minimize the number of configurations that have to be supported and avoid unnecessary terminal implementation complexity, the set of supported scenarios is carefully prioritized in 3GPP. This is discussed in more detail in Section 28.4.3.

LTE-Advanced can also make use of carrier aggregation to support deployments of *heterogeneous networks* consisting of a layer of macrocells and a layer of small cells coexisting with at least one carrier being common between them. In such a deployment, transmissions from one cell can interfere strongly with the control channels of another, thus impeding scheduling and signalling. LTE-Advanced supports *cross-carrier scheduling* (see Section 28.3.1) to enable control signalling to be transmitted on one component carrier corresponding to data transmissions on another; in this way, control channel interference between macrocells and small cells can be avoided.

Although the use of larger bandwidths by means of carrier aggregation allows higher peak data rates to be achieved, it does not increase the spectral efficiency as is required by the peak spectral efficiency targets shown in Table 27.2. LTE-Advanced therefore supports enhanced downlink MIMO transmission, by increasing the number of antennas at the eNodeB and UE, and hence the maximum number of spatial transmission layers for Single-User MIMO (SU-MIMO), from four in LTE Release 8 to eight. This may increase the multiplexing gain by a factor of two depending on the level of decorrelation between the antennas, and thus helps to achieve the spectral efficiency target of 30 bps/Hz. This is discussed in detail in Chapter 29.

Similarly to the downlink, the number of spatial layers supported in the uplink for SU-MIMO is increased from one to four in Release 10 in order to meet the peak spectral

efficiency target of 15 bps/Hz. In addition, transmit diversity is introduced for the uplink control signalling.

In order to further improve the spectral efficiency, especially at the cell edge, a later release of LTE-Advanced may incorporate enhanced support for *Coordinated MultiPoint* (CoMP) schemes. CoMP transmission in the downlink entails the coordination of transmissions from multiple cells. This may take the form of coordinated scheduling to one or more UEs to reduce interference or to achieve spatial multiplexing gain by benefiting from the macrodiversity that results from the low correlation between geographically diverse base station sites. With an even higher degree of coordination, multisite beamforming approaches may be considered. Some further details can be found in Section 29.5. Release 10 supports enhanced reference signals to facilitate multicell measurements by the UEs. In the uplink, CoMP reception at different cells is already possible as part of the network implementation in LTE Release 8.

In order to support deployments of LTE in parts of the network where a wired backhaul is not available or is very expensive, *Relay Nodes* (RNs) are supported by LTE-Advanced (see Chapter 30). An LTE-Advanced RN appears to the UEs as a Release 8 cell with its own cell ID. A UE receives and transmits all its control and data signals from and to the RN, while the RN separately uses LTE-Advanced technology to transfer control and data to and from a donor cell. The main characteristics and challenges of relaying are explained in Chapter 30.

27.3 Backward Compatibility

LTE-Advanced is defined as an evolution of LTE which can also be deployed on new bands. Hence, one of the design targets for LTE-Advanced was backward compatibility between LTE Release 8 and LTE Release 10 and beyond. This is reflected in [3] by requiring that a Release 8 LTE UE can work in a Release 10 LTE-Advanced network, and that an LTE-Advanced UE can work in a Release 8 LTE network.

This is an important requirement to give operators confidence to deploy LTE and to build upon it, as it means that LTE operators upgrading their network to LTE-Advanced will be able to do so without swapping their existing UE base; Release 8 LTE UEs will be able to enjoy service continuity in an LTE-Advanced network. Backward compatibility is also of key significance for UE and network complexity, as well as for the cost of implementation and verification, since it enables implementation reuse on both the UE and the network sides, and minimization of interoperability testing.

The requirement for backward compatibility does not, in general, prevent the introduction of new features, thereby providing a degree of future-proofness for LTE. New functionality can be configured by the network on a per-UE basis without affecting legacy UEs, for example by scheduling the Physical Downlink Shared CHannel (PDSCH) in a UE-specific way. It is more difficult to introduce new cell-wide features, which have to be compatible with both Release 8 and Release 10 terminals. LTE-Advanced will therefore be visible in the 3GPP specifications simply as 'LTE Release 10 and beyond', including the base functionality of Release 8 LTE. This reflects the nature of LTE-Advanced as the further evolution of LTE.

27.4 Deployment Aspects

The existence of internationally identified common frequency bands is a key factor for significant economies of scale in the development and production of terminals [8].

A key outcome of the WRC 2007 was that a total of 136 MHz of new global spectrum was allocated for use by IMT-designated radio technologies:

- 450–470 MHz;
- 790–806 MHz;
- 2300–2400 MHz.

Other region-specific bands were also allocated:

- 790–862 MHz for ITU Region 1 (EMEA⁵) and ITU Region 3 (all other Asia Pacific);
- 698–806 MHz for ITU Region 2 (North and South America) and ITU Region 3 (nine countries, including Japan, China and India);
- 3400–3600 MHz allocated to mobile use on a primary basis for ITU Region 1 (EMEA in 82 countries), ITU Region 2 (Americas in 14 countries, except US/Canada) and Region 3.

All new bands identified by the WRC 2007 are valid generically for IMT technologies – i.e. they are not specific to IMT-2000 or IMT-Advanced only [9]. The deployment of the frequency bands for the different regions is illustrated in Figure 27.1.

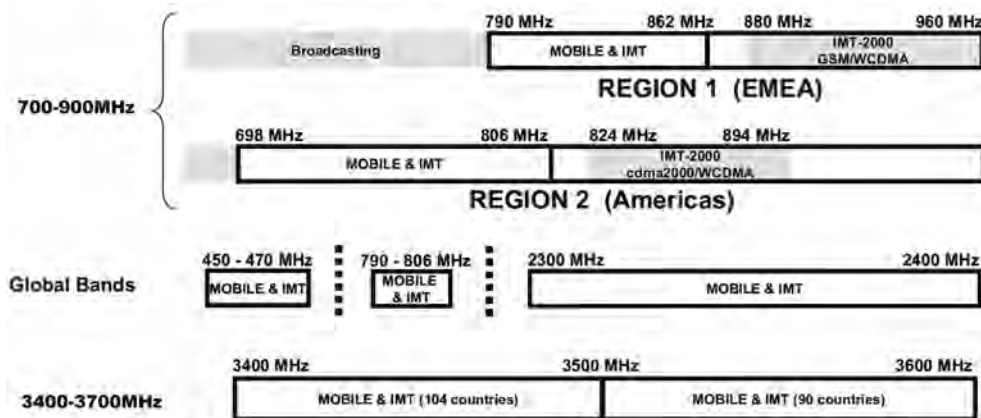


Figure 27.1: Allocation of new global spectrum resulting from WRC-07.

Carrier aggregation allows operators to gather spectrum for deployment of LTE from different parts of one band or from different bands. In order to limit the number of bands and combinations of bands to be supported by an LTE-Advanced UE, care has to be taken

⁵Europe, Middle East and Africa.

in focusing on the spectrum deployment scenarios with highest relevance for the mobile industry on a global level. 3GPP therefore prioritizes the development of requirements for different band combinations based on operators' requirements reflecting the different regions. Details of the first prioritized band combinations are given in Section 28.4.3.

27.5 UE Categories for LTE-Advanced

The definition of Release 10 UE categories builds upon the principles used in Releases 8 and 9 (see Section 1.3.4), where the number of UE categories is limited to avoid fragmentation of UE implementations and excessive variants in the market.

Three new Release 10 UE categories (6, 7 and 8) are specified [10], as shown in Table 27.3, defined in terms of their peak rates which reach about 3 Gbps in the downlink and 1.5 Gbps in the uplink. The highest UE category combines the aggregation of five 20 MHz component carriers with eight MIMO layers in the downlink and four in the uplink.

The peak data rate of categories 6 and 7 can be achieved by different means – for example, it is possible to achieve 300 Mbps either by supporting two MIMO layers together with the aggregation of 40 MHz or by four MIMO layers with a single 20 MHz carrier.⁶

Table 27.3: UE categories supported in Release 10.

	UE category		
	6	7	8
Approximate supported downlink data rate (Mbps)	300	300	3000
Approximate supported uplink data rate (Mbps)	50	100	1500
Number of downlink MIMO layers supported	2 or 4	2 or 4	8
Number of uplink MIMO layers supported	1, 2 or 4	1, 2 or 4	4
Support for 64QAM modulation in downlink	✓	✓	✓
Support for 64QAM modulation in uplink	✗	✗	✓
Relative memory requirement for downlink HARQ processing (normalized to category 1 level)	14.6	14.6	144

The Release 8 and 9 UE categories are reused and can in Release 10 support carrier aggregation; for example, a category 3 Release 10 UE may support the aggregation of two component carriers of up to 10 MHz bandwidth each. A category 6, 7 or 8 UE will also signal a category in the range 1–5, to allow backward compatibility in networks which do not yet support Release 10.

Additional UE categories are expected to be defined in the future, depending on market requirements.

⁶Note that the number of supported MIMO layers for each band combination is signalled by the UE.

References⁷

- [1] Circular Letter 5/LCCE/2, 'Invitation for submission of proposals for candidate radio interface technologies for the terrestrial components of the radio interface(s) for IMT-Advanced and invitation to participate in their subsequent evaluation', Radiocommunication Bureau, International Telecommunication Union (ITU), www.itu.int.
- [2] ITU-R Report M.2134, 'Requirements related to technical performance for IMT-Advanced radio interface(s)', www.itu.int/itu-r, 2008.
- [3] 3GPP Technical Report 36.913, 'Requirements for Further Advancements for Evolved Universal Terrestrial Radio Access (E-UTRA) (LTE-Advanced)', www.3gpp.org.
- [4] 3GPP Technical Report 25.814, 'Physical Layer Aspects for Evolved UTRA', www.3gpp.org.
- [5] 3GPP Technical Report 36.912, 'Feasibility study for Further Advancements for E-UTRA (LTE-Advanced)', www.3gpp.org.
- [6] 3GPP ITU-R Ad Hoc, 'RP-090736: Proposed Cover page for the October submission of "LTE Release 10 & beyond (LTE-Advanced)"', www.3gpp.org, 3GPP TSG RAN, meeting 45, Seville, Spain, September 2009.
- [7] LTE-Advanced Rapporteur, 'RP-090939: 3GPP Submission Package for IMT-Advanced', www.3gpp.org, 3GPP TSG RAN, meeting 45, Seville, Spain, September 2009.
- [8] GSMA, 'The advantage of common frequency bands for mobile handset production', www.gsmworld.com.
- [9] ITU World Radio Conference 2007 (WRC-07), October–November, 2007, Geneva.
- [10] 3GPP Technical Specification 36.306, 'Evolved Universal Terrestrial Radio Access (E-UTRA); User Equipment (UE) Radio Access Capabilities (Release 10)', www.3gpp.org.

⁷All web sites confirmed 1st March 2011.

Carrier Aggregation

Juan Montojo and Jelena Damnjanovic

28.1 Introduction

As discussed in Chapter 27, LTE-Advanced aims to support peak data rates of 1 Gbps in the downlink and 500 Mbps in the uplink [1]. In order to fulfil such requirements, a transmission bandwidth of up to 100 MHz is required; however, since the availability of such large portions of contiguous spectrum is rare in practice, LTE-Advanced uses carrier aggregation of multiple Component Carriers (CCs) to achieve high-bandwidth transmission. Release 8 LTE carriers have a maximum bandwidth of 20 MHz, so LTE-Advanced supports aggregation of up to five 20 MHz CCs.

A second motivation for carrier aggregation is to facilitate efficient use of fragmented spectrum, irrespective of the peak data rate. Carrier aggregation in LTE-Advanced is designed to support aggregation of a variety of different arrangements of CCs, including CCs of the same or different bandwidths, adjacent or non-adjacent CCs in the same frequency band, and CCs in different frequency bands. Each CC can take any of the transmission bandwidths supported by LTE Release 8, namely 6, 15, 25, 50, 75 or 100 Resource Blocks (RBs), corresponding to channel bandwidths of 1.4, 3, 5, 10, 15 and 20 MHz respectively. For Frequency Division Duplex (FDD) operation, the number of aggregated carriers in uplink and downlink may be different (although Release 10 focuses on the case where the number of downlink CCs is not less than the number of uplink CCs).¹ This flexibility enables a large variety of fragmented spectrum arrangements of relevance to network operators to be supported.

A third motivation for carrier aggregation is support of *heterogeneous networks*. A heterogeneous network deployment typically consists of a layer of high-power macrocells

¹For TDD deployment, the number of CCs and the bandwidth of each CC in uplink and downlink is expected to be the same.

LTE – The UMTS Long Term Evolution: From Theory to Practice, Second Edition.

Stefania Sesia, Issam Toufik and Matthew Baker.

© 2011 John Wiley & Sons, Ltd. Published 2011 by John Wiley & Sons, Ltd.

and a layer of low-power small cells (e.g. picocells, Closed Subscriber Group (CSG) femtocells or relay nodes – see Chapters 24 and 30) with at least one carrier being used by both layers. In such a deployment, transmissions from one cell can interfere strongly with the control channels of another, thus impeding scheduling and signalling. Rather than simply using separate carriers for the two layers, which would result in inefficient spectrum usage, carrier aggregation enables multiple carriers to be used for a given layer, while interference can be avoided by means of *cross-carrier scheduling*. Cross-carrier scheduling allows the Physical Downlink Control Channel (PDCCH) on the CC of one serving cell to schedule transmission resources on a CC of another serving cell, as explained in detail in Section 28.3.1.

All CCs in Release 10 are designed to be *backward-compatible*. This means that it is possible to configure each CC such that it is fully accessible to Release 8 User Equipment (UEs). Therefore, essential Release 8 channels and signals such as Primary and Secondary Synchronization Signals (PSS and SSS) and System Information (SI) specific to each CC are transmitted on the respective CC. Backward-compatibility also has the advantage that the technology developed for LTE Release 8 can be reused on aggregated Release 10 CCs. From the higher-layer perspective, each CC appears as a separate cell with its own Cell ID. A UE that is configured for carrier aggregation connects to one *Primary Serving Cell* (known as the ‘PCell’) and up to four *Secondary Serving Cells* (known as ‘SCells’). The PCell is defined as the cell that is initially configured during connection establishment; it plays an essential role with respect to security, NAS² mobility information, SI for configured cells, and some lower-layer functions. An SCell is a cell that may be configured after connection establishment, merely to provide additional radio resources. The term *Serving Cell* can refer to either a PCell or an SCell. The same frame structure is used in all aggregated serving cells, and, for TDD carrier aggregation, the uplink-downlink configuration (see Section 6.2) across all serving cells is the same.

The CCs corresponding to the PCell are referred to as the Downlink and Uplink Primary Component Carriers (PCCs), while the CCs corresponding to an SCell are referred to as Downlink and Uplink Secondary Component Carriers (SCCs). In a given geographic cell, all CCs that may be aggregated are assumed to be synchronized and belong to the same eNodeB. A default linkage between downlink and uplink CCs is signalled in System Information Block 2 (SIB2) on each downlink CC.

A UE’s identity (C-RNTI³) is the same in the PCell and its configured SCells.

28.2 Protocols for Carrier Aggregation

28.2.1 Initial Acquisition, Connection Establishment and CC Management

As noted above, the PSS and SSS are transmitted on all CCs to facilitate cell search.⁴ A UE establishes a connection to a cell by following the usual Release 8 and 9 procedures. After the initial security activation procedure, E-UTRAN may configure a UE supporting carrier aggregation with one or more SCells in addition to the PCell that is initially configured

²Non-Access Stratum.

³Cell Radio Network Temporary Identifier.

⁴See Chapter 7 for details of the cell search procedures defined in Release 8.

during connection establishment. The configured set of serving cells for a UE always contains one PCell and may also contain one or more SCells. The number of serving cells that can be configured depends on the aggregation capability of a UE. For each SCell, the usage of uplink resources by the UE in addition to the downlink ones is configurable – the number of downlink SCCs configured for a UE is therefore always greater than or equal to the number of uplink SCCs, and no SCell can be configured for usage of uplink resources only. From a UE viewpoint, each uplink resource belongs to only one serving cell.

The PCell provides the security inputs, NAS mobility information and SI for serving cells. A single Radio Resource Control (RRC) connection is established with the PCell, which controls all the CCs configured for a UE.

After RRC connection establishment to the PCell, reconfiguration, addition and removal of SCells can be performed by RRC. When adding a new SCell, dedicated RRC signalling is used to send all the required SI for the new SCell. While in connected mode, changes of SI for an SCell are handled by release and addition of the affected SCell, and this may be done with a single RRC reconfiguration message.

In RRC_CONNECTED state, as the radio conditions for a UE change on different CCs or the load on different CCs changes, the network may decide to change the PCell for a UE. The Release 8 signalling specifications already enable this, via the handover procedure (i.e. with security key change and the random access procedure – see Section 3.2.3.4), which is the only means by which the PCell can be changed. The detailed flow chart for PCell change is shown in Figure 28.1.

In the case of intra-LTE handover, RRC can add, remove, or reconfigure SCells for the target PCell. This enables a UE to begin immediately to use the assigned CCs after handover signalling is complete. The source PCell passes all necessary information to the target PCell (e.g. E-UTRAN Radio Access Bearer (E-RAB) attributes and RRC context). In addition, to enable SCell selection in the target PCell, the source PCell can provide a list of the best cells in decreasing order of radio quality. The target PCell decides which SCells are configured for use after handover, which may include cells other than the ones indicated by the source PCell. The UE does not autonomously release any SCell configuration at handover; as usual, control is network-based.

28.2.2 Measurements and Mobility

For the purposes of mobility, a UE sees a CC in the same way as any other carrier frequency, and a measurement object (see Section 3.2.5.1) has to be set up for each CC in order for the UE to measure it. Inter-frequency neighbour cell measurements encompass all carrier frequencies which are not configured as CCs. Release 8 measurement events (see Section 3.2.5.2) are applicable for UEs configured with carrier aggregation, and the following rules apply:

- There is at most one serving cell (PCell or SCell) per measurement identity;
- For measurement events A1 and A2, the serving cell of the event is the configured serving cell (PCell or SCell) corresponding to the measurement object (i.e. the eNodeB may configure separate events A1 and A2 for each serving cell).

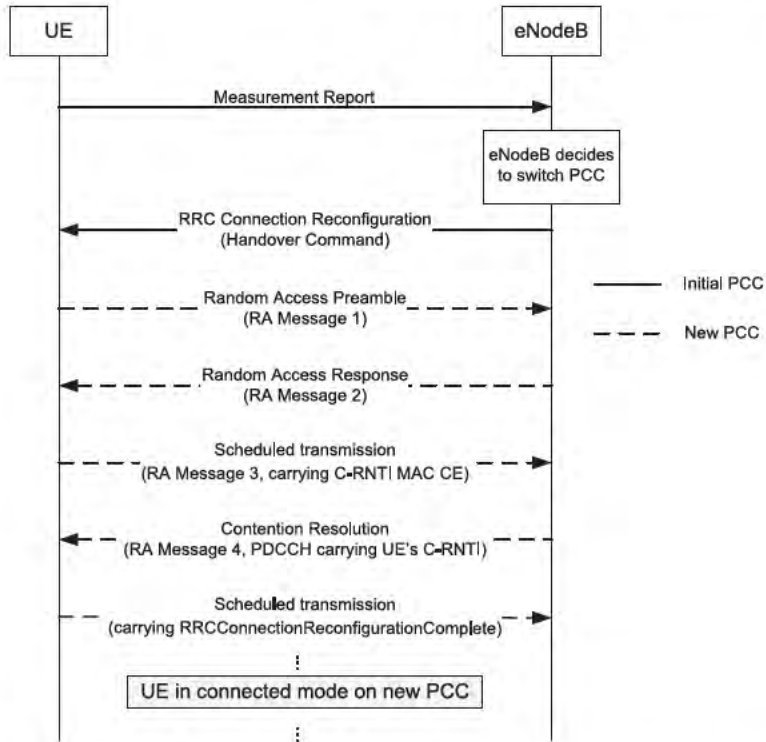


Figure 28.1: Procedure for PCell change with contention-based handover.

- For measurement events A3, A5 and B2, the serving cell used as a reference is the PCell. The measurement object linked to an A3 or A5 event can be any frequency, and, if an SCC is the target object, the corresponding SCell is included in the comparison.

In addition, a new measurement event A6 is introduced for carrier aggregation. Measurement event A6 is defined as ‘intra-frequency neighbour becomes better than an offset relative to an SCell’ and is intended for intra-frequency measurement events on SCCs. For this measurement, the neighbour cells on an SCC are compared to the SCell of that SCC. An example of the relationship between A3 and A6 is given in Figure 28.2.

Measurements on activated CCs can be done without measurement gaps. Measurement gaps are UE specific; UE capability signalling is used to inform the eNodeB about the need for measurement gaps independently for each supported measured band.

Measurements on all activated cells follow the Release 8 procedures and requirements (see Section 22.3). Measurement periods of deactivated SCells are configurable by RRC signalling (with a range of values from 160 ms to 1280 ms, with a default value of 320 ms). Measurement accuracy requirements are the same for all cells.

The quality threshold for cell selection (the S-criterion – see Section 3.3.3) applies to the PCell and controls all non-serving-cell measurements. In other words, when the PCell




Figure 28.2: Measurement Events A3 and A6.

Reference Signal Received Power (RSRP) (after Layer 3 filtering) is higher than the S-criterion, all measurements other than those that are only on the PCell or only on an SCell can be disabled.

28.2.2.1 Radio Link Failure

Detection of a Radio Link Failure (RLF) (see Section 22.7) on the PCell triggers an RRC connection re-establishment. Triggers for RRC connection re-establishment include:

- Failure of the PCell according to the same criteria as are used for RLF detection in Release 8, based on N310/N311/T310;
- Random access failure in the PCell, as in Release 8;
- An indication from Radio Link Control (RLC) that the maximum number of retransmissions has been reached, as in LTE Release 8.

The UE does not perform radio link monitoring for downlink SCCs. The eNodeB can detect poor link quality on an SCC from CQI reports and/or existing RRM measurement reports for SCells.

28.2.2.2 Idle Mode Procedures

The same mobility procedures as defined in LTE Release 8 apply to a Release 10 UE in a network which deploys carrier aggregation. A UE in RRC_IDLE therefore always behaves as a single-carrier UE, without the possibility of having multiple aggregated CCs.

RRM requirements are defined for carrier aggregation in both idle and connected modes [2]. This is to ensure that good mobility performance is met in all cases including low- and high-mobility scenarios. The mobility performance in the different scenarios can be optimized using different network-controlled parameters for cell reselection in idle mode and for handover in active mode.

28.2.3 User Plane Protocols

From the perspective of the NAS (see Section 2.2.1), the UE is connected to the PCell, which provides the security keys at handover and the tracking area for Tracking Area Updates (TAUs). Other CCs are simply considered as additional transmission resources.

The multiple CCs of carrier aggregation are not visible to the Packet Data Convergence Protocol (PDCP) and Radio Link Control (RLC) layers, and these protocols are therefore unchanged from LTE Release 8 except to enable them to support data rates up to 1 Gbps.

At the MAC layer, each CC has its own independent Hybrid Automatic Repeat reQuest (HARQ) entity. From the perspective of the UE, the characteristics of the HARQ procedures are unchanged with respect to those defined for Release 8 (see Section 4.4). One transport block and an independent HARQ entity are scheduled per CC in the absence of spatial multiplexing, and up to two when spatial multiplexing is configured. The User Plane structures for the downlink and the uplink are shown in Figures 28.3 and 28.4 respectively.

28.2.3.1 Scheduling

It can be seen from Figure 28.3 that a single scheduler entity covers all the UEs and all their corresponding CCs. Elementary queuing theory indicates that a globally optimized scheduler will achieve better performance than a per-CC scheduler. In practice, however, independent schedulers may be utilized for each CC, depending on eNodeB implementation.

Dynamic scheduling is performed every subframe by means of grants transmitted on PDCCH. The grants may be transmitted on the same carrier as the assigned data resources or on a different carrier if cross-carrier scheduling is configured (see Section 28.3.1).

Semi-Persistent Scheduling (SPS) (see Section 4.4.2.1) can only be configured for the PCell, and only PDCCH allocations for the PCell can override an SPS resource allocation.

28.2.3.2 Random Access Procedure

As in LTE Release 8, the MAC layer is responsible for controlling the random access procedure (see Section 4.4.2.3). In the case of carrier aggregation, a UE performs the random access procedure (see Chapter 17) on the uplink CC associated with the PCell. No more than one random access procedure is ongoing at any time, irrespective of the carrier aggregation capability or configuration of the UE.

When carrier aggregation is configured, the first three steps of the contention-based random access procedure (see Section 17.3.1) occur on the PCell, while cross-carrier scheduling from the PCell (see Section 28.3.1) can be used for the contention resolution (step 4). In the non-contention-based random access procedure (see Section 17.3.2), the Random Access Preamble assignment via PDCCH of step 0, step 1 and step 2 occur on the PCell.

28.2.3.3 Discontinuous Reception Procedure

The discontinuous reception (DRX) procedures defined in Release 8 (see Section 4.4.2.5) remain applicable in Release 10. If one or more SCells are configured for a UE in addition to the PCell, the same DRX operation applies to all the serving cells. This means that the active times for PDCCH monitoring are identical across all downlink CCs.




Figure 28.3: Downlink Layer 2 protocol structure for carrier aggregation.
 Reproduced by permission of © 3GPP.




Figure 28.4: Uplink Layer 2 protocol structure for carrier aggregation.
 Reproduced by permission of © 3GPP.

28.2.3.4 SCell Activation and Deactivation

In addition to the DRX operation, some UE power saving may be achieved by fast activation and deactivation of individual SCells (the PCell cannot be deactivated).

When an SCell is deactivated, the UE does not have to receive data transmissions or monitor the PDCCH for that SCell. The UE is also not required to perform measurements for Channel State Information (CSI) reporting. Deactivated SCells can, however, be used as the path-loss reference for the measurements for uplink power control (see Section 28.3.5.1). It is assumed that these measurements would be less frequent while the SCell is deactivated, in order to obtain power savings at the UE. When the downlink CC of an SCell is activated or deactivated, the SIB2-linked uplink CC follows suit. Deactivation of the uplink CC includes ceasing Sounding Reference Signal (SRS) transmissions and all PUSCH transmissions (including any pending retransmissions).

Activation and deactivation of SCells is under eNodeB control. The activation and deactivation is executed by means of MAC Control Elements (see Section 4.4.2.7), which can activate or deactivate one or more SCells indicated by an 8-bit bitmap. A timer may also be used for automatic deactivation if no data or PDCCH messages are received on a CC for a certain period; this is the only case in which deactivation can be executed autonomously by the UE.⁵ Even so, the duration of the timer is configured by the eNodeB and can take the value ‘infinity’, effectively disabling timer-based deactivation.

The timing of activation and deactivation is carefully defined in order to ensure that there is a common understanding between the eNodeB and the UE. If a MAC control element activating an SCell is received in subframe n , then the SCell has to be ready for operation in subframe $n + 8$. Hence, from subframe $n + 8$, the UE is required to monitor the PDCCH for both uplink grants and downlink assignments corresponding to the newly activated SCell. SRS transmissions can also be started in subframe $n + 8$. CSI reports are commenced in subframe $n + 8$, and CSI measurements in subframe $n + 8$ at the latest. If there is no CSI measurement available for the SCell when the UE first starts reporting CSI, the UE is expected to report the value ‘out of range’. Power headroom reporting (see Section 28.3.5.3) also starts in subframe $n + 8$, and the SCell deactivation timer is started. If a MAC control element deactivating an SCell is received (or the deactivation timer expires) in subframe n , CSI reports cease from subframe $n + 8$.⁶

When an RRC reconfiguration occurs that includes mobility control information (i.e. a handover), all SCells are deactivated. If mobility control information is not included, SCells that are added to the set of serving cells are initially in the deactivated state, while any SCells that remain in the set of serving cells (either unchanged or reconfigured) do not change their activation status.

⁵Downlink SCell quality is never used to cause a UE to cease transmissions in an SCell.

⁶If the UE misses a PDCCH message, there may be a temporary misalignment between the UE’s and eNodeB’s understandings of the activation status of an SCell. This may affect the rate-matching (and hence decoding) of the PUSCH (see Section 16.4) due to uncertainty as to the presence of a CSI report; the eNodeB may be able to mitigate any consequent effects on uplink throughput by making another decoding attempt of the PUSCH with a different rate-matching assumption.

28.2.3.5 Buffer Status Reporting

As in Release 8, there can only be one Buffer Status Report (BSR) per transport block. However, there can be several BSRs in a subframe:

- Zero or one Regular or Periodic BSRs;
- Zero, one or more Padding BSRs of possibly different kinds, but all following the Release 8 rules (see Section 4.4.2.2).

All BSRs transmitted in a subframe reflect the buffered data that remains after all the MAC Protocol Data Units (PDUs) have been built for the CCs that are scheduled in the subframe. When more than one serving cell allows a Regular or Periodic BSR to be sent in a subframe, the UE can choose the serving cell in which the Regular or Periodic BSR is transmitted.

The amount of data that may have to be indicated by a BSR when multiple CCs are aggregated is much higher than in Release 8, due to the higher data rates supported.⁷ An additional table of BSR values is therefore introduced in Release 10 to enable indication of the larger buffer sizes. The usage of the new table is controlled by RRC. All UEs which support uplink carrier aggregation or uplink MIMO must support the new BSR table.

28.2.3.6 Logical Channel Prioritization

Different CCs may provide similar QoS, and therefore the UE is allowed complete freedom in how it maps uplink data to granted resources on different CCs.

When the UE is provided with uplink grants in multiple serving cells in one subframe, the order in which the grants are processed during logical channel prioritization, and whether joint or serial processing is applied, are left up to UE implementation. A variety of approaches can fulfil the long-term Prioritized Bit Rate (PBR) (see Section 4.4.2.6) for each logical channel.

28.3 Physical Layer Aspects

At the physical layer, each transport block is mapped to a single CC of a serving cell, as shown in Figure 28.5. Even if a UE is scheduled on multiple CCs simultaneously, HARQ, modulation, coding and resource allocation, together with the corresponding signalling, are performed independently on each CC.

28.3.1 Downlink Control Signalling

Each downlink CC carries a control signalling region for the PCFICH, PDCCH and PHICH⁸ at the start of each subframe, as in Release 8 (see Figure 9.5).

⁷The introduction of uplink MIMO in Release 10 (see Section 29.4) further increases the uplink data rates, with corresponding impact on the BSR.

⁸Physical Control Format Indicator Channel, Physical Downlink Control Channel and Physical Hybrid ARQ Indicator Channel.




Figure 28.5: MAC to physical layer mapping for carrier aggregation.

28.3.1.1 PDCCH

As in Release 8, it is possible for a PDCCH on each downlink CC to carry downlink resource assignments applicable to the same CC, and uplink resource grants applicable to the associated uplink CC (according to the linkage indicated in SIB2).

In addition, a key feature of carrier aggregation is *cross-carrier scheduling*. This enables a PDCCH on one CC to schedule data transmissions on another CC by means of a new 3-bit *Carrier Indicator Field (CIF)* inserted at the beginning of the PDCCH messages. The rest of the Release 8 PDCCH Control Channel Element (CCE) structure, coding and message contents (as described in Section 9.3.5) is unchanged for carrier aggregation. The presence or absence of the CIF on each CC is configured semi-statically (i.e. by RRC signalling) for each UE. When configured, the CIF is only present in PDCCH messages in the UE-specific search space (see Section 9.3.5.5), not the common search space.

For data transmissions on a given CC, a UE expects to receive scheduling messages on the PDCCH on just one CC – either the same CC, or a different CC via cross-carrier scheduling; this mapping from PDCCH to PDSCH is also configured semi-statically.⁹ Some example configurations are shown in Figure 28.6.

For the CC of a serving cell on which PDCCH is monitored, the UE searches for PDCCH messages at least for the same CC of the serving cell. In the example in Figure 28.6(b), the UE monitors the PDCCH on CC1 of serving cell 1 for assignments on CC1, and CC1 resources cannot be cross-scheduled from any other CC of the serving cells. The UE also searches for PDCCHs with CIF on CC1 for assignments for CC2 and CC3, without monitoring the PDCCH on CC2 or CC3 (of serving cells 2 and 3 respectively).

⁹Note that PDSCH transmissions on the PCell cannot use cross-carrier scheduling – their corresponding PDCCH messages must also be transmitted on the PCell.

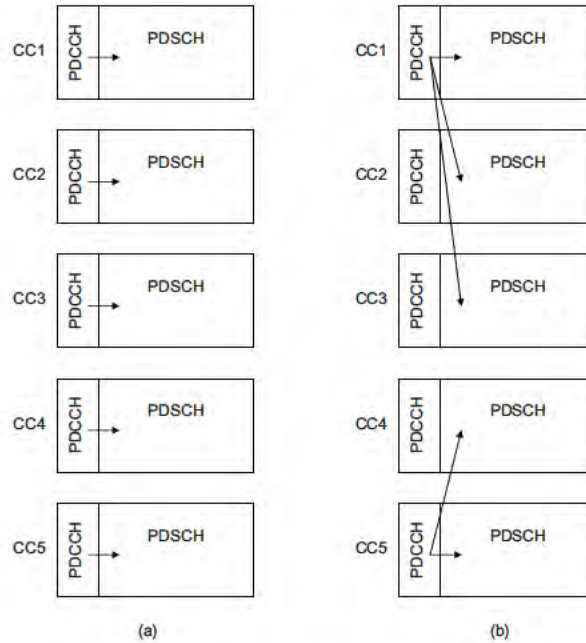


Figure 28.6: Examples of semi-statically configured mappings from PDCCH scheduling messages to CCs for data transmission:
 (a) without cross-carrier scheduling; (b) with cross-carrier scheduling.

A UE configured with the CIF for a serving cell uses the CIF value from the detected PDCCH to identify the serving cell on which the corresponding PUSCH or PDSCH transmission will take place. For simplicity, the CIF value is set to be the same as the cell index.

If CIF is not configured for a UE, an uplink grant or downlink assignment received on a given serving cell corresponds to PUSCH or PDSCH transmission on the same serving cell.

In general, it is obvious that the amount of processing that the UE must perform in case of carrier aggregation is significantly larger than in Release 8. This is also true for the PDCCH decoding, where, in the worst case, the number of blind decodes the UE must perform is expected to increase linearly with the number of configured CCs. As discussed in Section 9.3.5.5, the maximum number of blind decodes in any subframe for single-carrier operation is 44 (12 in the common search space and 32 in the UE-specific search space).¹⁰

When carrier aggregation is configured, the maximum total number of blind decodes a UE is required to perform is 44 for the PCell, plus 32 for each active downlink SCC.¹¹

For any downlink CC where the UE monitors PDCCH without the CIF being configured, the search space is the same as in Release 8.

¹⁰In addition, in Release 10 a further 16 blind decodes are necessary to support uplink MIMO, taking the total number of blind decodes per CC to 60.


¹¹A further 16 blind decodes are needed for each uplink CC that is configured for uplink MIMO operation.

In the case of cross-carrier scheduling, the total search space size (in terms of number of CCEs) is extended beyond the Release 8 size. For a given UE, the UE-specific search spaces located in the control region of a CC are individually defined per aggregation level for each PDSCH/PUSCH CC linked to that CC for control signalling. UE-specific search spaces corresponding to different CCs in a given control region are shared if the Downlink Control Information (DCI) format size is the same between the CCs. For any downlink CC with CIF where the UE monitors the PDCCH, a UE-specific search space $S_k^{(L)}$ in subframe k for the PDSCH/PUSCH CC c at aggregation level $L=\{1, 2, 4, 8\}$ is defined by a set of PDCCH candidates. The CCEs corresponding to PDCCH candidate m of the search space $S_k^{(L)}$ are given by:

$$S_k^{(L)} = L \cdot \left\{ Y_k + m + M^{(L)} \cdot c \right\} \bmod [N_{CCE,k}/L] + i. \quad (28.1)$$

where Y_k is the output of the UE-specific subframe-to-subframe search space hopping sequence (see Section 9.3.5.5 and [3, Section 9.1.1]), $i = 0, \dots, L - 1$, and $m = 0, \dots, M^{(L)} - 1$, $M^{(L)}$ is the number of PDCCH candidates to monitor in the given search space. $N_{CCE,k}$ is the total number of CCEs in the control region of subframe k .

This UE-specific search space design is shown in Figure 28.7.



$M^{(L)}$ is the number of PDCCH candidates to monitor in the given search space for aggregation level L . $||$ is the index of the i^{th} $||$.

Figure 28.7: UE-specific search space for multiple CCs.

For the common search space, the term $M^{(L)} \cdot c$ in Equation (28.1) is set to zero. The same applies if a UE is not configured with CIF.

Cross-Carrier Scheduling in Heterogeneous Networks

The main motivation for cross-carrier scheduling in LTE-Advanced is to provide support for Inter-Cell Interference Coordination (ICIC) for the PDCCH in heterogeneous network deployments, as mentioned in Section 28.1. Figure 28.8 shows a typical heterogeneous network scenario where macrocells and small cells share two downlink CCs, denoted CC1 and CC2. The small cells use both CCs at low transmit power, and the macrocells use CC1 at high power and CC2 at reduced power. The macrocells' transmissions on CC1 would cause high interference to the small cells, and therefore it is beneficial for the small cells to be able to use PDCCH messages on CC2 to perform cross-carrier scheduling for data transmissions

on CC1. To facilitate this, the macrocells can refrain from transmitting PDCCHs on CC2 (or transmit only with low power), instead using CC1 to schedule data transmissions on both CC1 and CC2, with cross-carrier scheduling for the latter.¹² This effectively provides ICIC for the PDCCH, while the Release 8 ICIC mechanisms may be utilized for PDSCH data.¹³




Figure 28.8: A typical heterogeneous network deployment with macrocells and small cells sharing two CCs.

28.3.1.2 PCFICH

Because of the potential for different loadings on different CCs, the number of OFDM¹⁴ symbols used for the downlink control region, and hence the starting point of the PDSCH, can be set independently on each CC. However, in cases when cross-carrier scheduling is employed due to unreliable PDCCH reception on an SCell CC that is used for data transmission, the same inter-cell interference affecting the PDCCH would typically also affect the PCFICH, as they are both in the same control region of the CC. To address this issue without incurring a high signalling overhead, Release 10 provides a mechanism whereby the index of the first OFDM symbol of any cross-scheduled PDSCH can be signalled semi-statically for each CC. This obviates the need for cross-scheduled UEs to decode the PCFICH on the target CC. It should be noted, however, that this does not prevent the eNodeB from

¹²These issues are not problematic for homogeneous networks of macrocells transmitting at similar power levels, where the Release 8 design techniques aimed at maximizing frequency diversity and randomizing the inter-cell interference are entirely appropriate.

¹³For the data, the Release 8 ICIC mechanisms (e.g. based on Relative Narrowband Transmit Power (RNTP) indicators for the downlink – see Section 12.5) enable Fractional Frequency Reuse (FFR) to be configured between cells with a granularity of one RB; this X2-interface-based ICIC mechanism for the data channels is equally applicable in heterogeneous network carrier aggregation scenarios to share frequency resources on the different CCs between macrocells and neighbouring small cells.

¹⁴Orthogonal Frequency Division Multiplexing.

varying the control region size dynamically (i.e. from subframe to subframe) on each CC (although a relatively static control region size is likely to be suitable for ICIC purposes in many heterogeneous network scenarios). If a shorter control region than the semi-statically signalled one is used, the additional OFDM symbol(s) can still be used for data transmission in RBs assigned to non-cross-scheduled UEs; if a longer control region is used, it will cause some degradation to the PDSCH decoding for cross-scheduled UEs.

28.3.1.3 PHICH

The design of the PHICH (used for transmission of Hybrid ARQ ACK/NACKs in response to uplink data transmissions) for carrier aggregation is based on that defined in Release 8: the physical transmission aspects (orthogonal code design, modulation, scrambling sequence and mapping to resource elements – see Section 9.3.4) are all identical. The PHICH is transmitted on the downlink CC that was used to transmit the corresponding uplink resource grant; this is particularly beneficial for heterogeneous network deployments where some CCs may experience high inter-cell interference in the control channel region as explained above. If cross-carrier scheduling is used, one downlink CC may have to carry PHICH transmissions for multiple uplink CCs, and therefore there is an increased probability of PHICH collisions occurring (since the PHICH index is determined from the lowest PRB of the corresponding PUSCH transmission (see Section 9.3.4), which may be the same on multiple uplink CCs). To mitigate this, the PHICH index can be shifted by configuring different cyclic shifts of the PUSCH demodulation Reference Signals (RSs) among the uplink CCs whose PHICHs are transmitted in the same downlink CC control region (see Figure 9.9). In addition, the eNodeB scheduler can aim to avoid collisions by selecting different starting PRBs for the uplink resource allocations on the different CCs.

28.3.2 Uplink Control Signalling

For carrier aggregation, the uplink control signalling (HARQ ACK/NACK signalling, scheduling requests and Channel State Information (CSI) feedback) has to support up to five downlink CCs.

A UE may send a HARQ ACK/NACK for every downlink transport block – i.e. up to ten per subframe in the case of downlink spatial multiplexing with five downlink CCs. Since the Release 8 Physical Uplink Control Channel (PUCCH) (see Section 16.3.2) was not designed to carry such large numbers of ACK/NACK bits, new mechanisms are defined for carrier aggregation in Release 10.

Similarly, CSI feedback may be needed for up to five downlink CCs, although not necessarily all in the same subframe.

All PUCCH control signalling (corresponding to all configured CCs) is transmitted on the uplink PCC of the PCell (the uplink CC that is SIB2-linked to the configured downlink PCC). Thus PUCCH is never transmitted on more than one uplink CC.

Uplink control signalling may also be mapped to the PUSCH, as in Release 8 (see Section 16.4).

In addition to the Release 8 multiplexing modes, Release 10 supports simultaneous transmission of PUCCH for control information and PUSCH for data. Some potential benefits of simultaneous PUCCH and PUSCH transmission include:

- The IoT operating point can be set independently for control and data, which can improve efficiency as HARQ enables the PUSCH to operate at higher IoT levels than is possible for reliable control information reception on the PUCCH;
- Interference fluctuations on RBs used for PUCCH may be reduced.

28.3.2.1 HARQ Feedback

In order to provide HARQ feedback (ACK/NACK) for PDSCH transmissions on multiple CCs, new multibit ACK/NACK PUCCH formats are defined in Release 10 in support of carrier aggregation:

- PUCCH format 3;
- PUCCH format 1b with PUCCH ‘channel selection’, whereby some of the ACK/NACK information to be conveyed is indicated by selecting one of a number of possible PUCCH resources, in a similar way to Release 8 TDD operation – see Section 23.4.3.

For UEs that support no more than four ACK/NACK bits and are configured with up to two CCs, PUCCH format 1b with channel selection is utilized. For UEs that support more than four ACK/NACK bits, both PUCCH format 1b with channel selection and format 3 are supported, where PUCCH format 1b with channel selection can be used for up to four ACK/NACK bits and two configured CCs and format 3 for the full range of ACK/NACK bits; RRC signalling configures which PUCCH format is used in this case.

PUCCH Format 3

PUCCH format 3 is designed to convey large ACK/NACK payloads. Unlike the Release 8 PUCCH formats (namely PUCCH formats 1, 1a, 1b and 2 – see Section 16.3.2), PUCCH format 3 is not based on Zadoff-Chu sequences and is more similar to PUSCH transmissions. It has the following characteristics, illustrated in Figure 28.9:

- DFT-S-OFDM¹⁵ waveform;
- The same demodulation RS structure as PUCCH format 2 (for both normal and extended cyclic prefix);
- Orthogonal cover sequence applied to the SC-FDMA symbols used for ACK/NACK data: these sequences are DFT sequences of length 5, allowing multiplexing of up to 5 format 3 transmissions in the same RB;
- No orthogonal cover sequence applied in the SC-FDMA symbols used for demodulation RSs; the RSs of multiple UEs are multiplexed by means of different cyclic shifts;
- A shortened format is defined in which the last SC-FDMA symbol is punctured and the orthogonal cover sequence is shortened to a Hadamard sequence of length 4.
- QPSK modulation.

The resulting PUCCH format 3 supports transmission of 48 coded bits. The actual number of bits of ACK/NACK feedback is determined from the number of configured CCs, the

¹⁵Discrete Fourier Transform Spread Orthogonal Frequency Division Multiplexing.

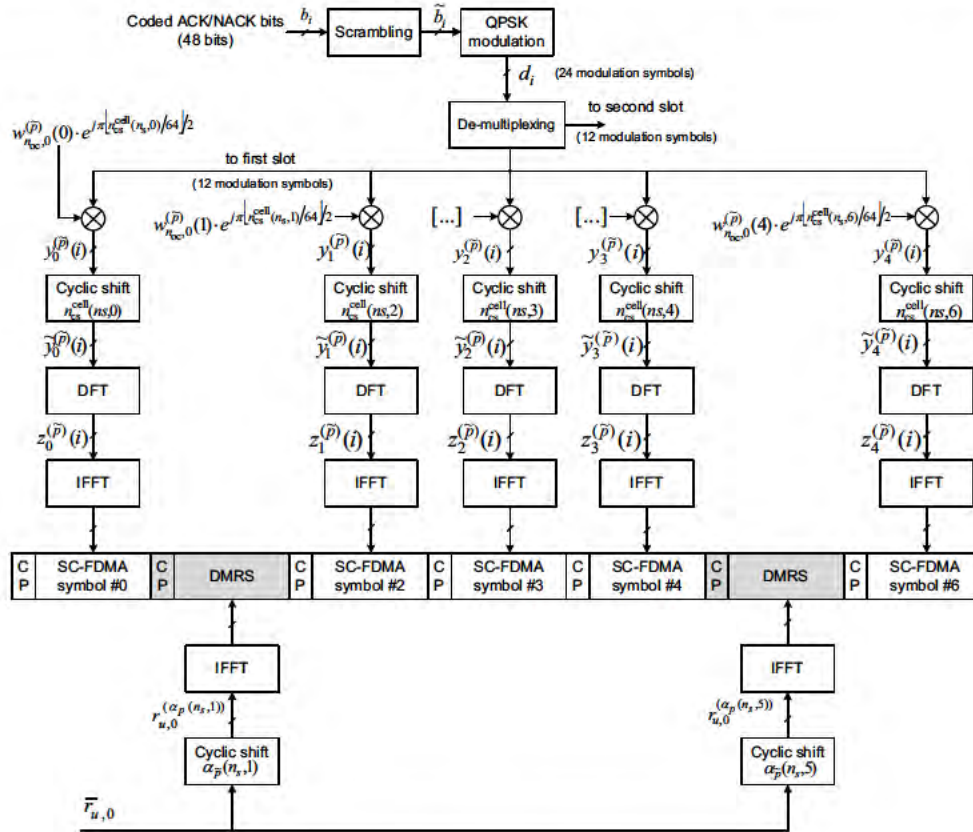


Figure 28.9: Structure of PUCCH format 3.

configured transmission modes on each of them, and, in TDD, the ACK/NACK bundling window size (the number of downlink subframes associated with a single uplink subframe – see Section 23.4.3). For FDD, a maximum payload of 10 ACK/NACK bits is supported, covering up to five CCs configured for MIMO transmission (i.e. two ACK/NACK bits per CC). For TDD, PUCCH format 3 supports an ACK/NACK payload size of up to 20 bits; if the number of ACK/NACK bits to be fed back for multiple downlink subframes associated with a single uplink subframe is greater than 20, ‘spatial bundling’ (i.e. a logical AND) of the ACK/NACK bits corresponding to the two codewords within a downlink subframe is performed for each of the serving cells. The maximum payload size carried by PUCCH format 3 in Release 10 is 21 bits (i.e. a code rate of 0.4375), corresponding to 20 bits of ACK/NACK information and one bit for a Scheduling Request (SR)¹⁶ appended at the end of the ACK/NACK bits.

The ACK/NACK bits are concatenated in ascending order of the downlink CC index. For payload sizes less than or equal to 11 bits, channel coding uses the Reed-Muller (RM)

¹⁶See Section 4.4.2.2.

code from Release 8, with circular buffer rate matching (as explained in Section 10.3.2.4). When the payload is larger than 11 bits, alternate ACK/NACK bits are input to two separate RM encoders. Finally, in order to mitigate inter-cell interference, cell-specific scrambling per SC-FDMA symbol is introduced. This structure is shown in Figure 28.10.

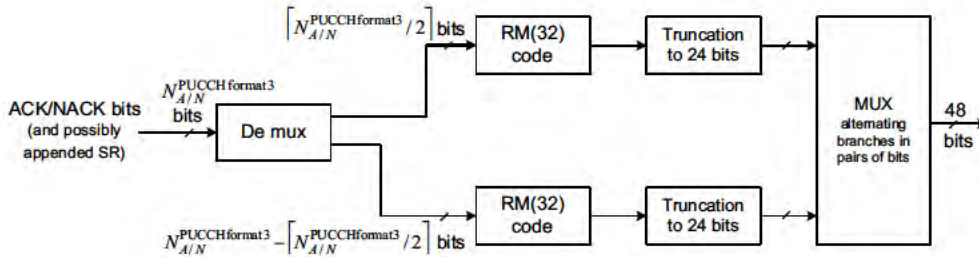


Figure 28.10: Coding and scrambling for PUCCH format 3.

The PUCCH resource to be used for format 3 is signalled explicitly to the UE. A set of four resources is configured by RRC signalling, of which one resource is then indicated dynamically for each ACK/NACK occasion using an indicator transmitted in the Transmitter Power Control (TPC) field of the PDCCH corresponding to PDSCH on the SCCs. All SCC PDCCH assignments in a given subframe indicate the same value.

If transmit diversity is used for PUCCH format 3 (see Section 29.4.2), the RRC signalling configures four pairs of PUCCH resources, and the PDCCH assigning resources for the SCC PDSCH indicates one of these pairs to be used by the two antenna ports.

If no PDCCH corresponding to PDSCH on an SCC is received in a given subframe and a single PDSCH is received on PCC, a UE that is configured for PUCCH format 3 would instead use the Release 8 format 1a or 1b.

PUCCH Format 1b with Channel Selection

PUCCH format 1b with channel selection involves configuring up to four PUCCH format 1b resources (‘channels’); the selection of one of these resources indicates some of the ACK/NACK information to be conveyed.

For FDD, the use of PUCCH format 1b with channel selection to convey the ACK/NACK information for two CCs is straightforward. For TDD, it is necessary to use spatial bundling of ACK/NACK bits across the two codewords within a downlink subframe for each of the serving cells if the number of ACK/NACK bits to be fed back is greater than four. If the number of ACK/NACK bits after performing spatial bundling is still larger than four, time-domain bundling is employed in addition.

Mapping tables are specified for the cases of two, three or four ACK/NACK bits to define the mapping of ACK/NACK combinations to the configured PUCCH resources. These tables are designed to support fully implicit resource indication, fallback to Release 8 operation in the case of a single configured CC, and equalization of the performance of individual ACK/NACK bits. Separate mapping tables are defined depending on whether or not time-domain bundling of the ACK/NACK feedback is performed. A Release 10 UE configured

with a single TDD serving cell can be configured to use either the mapping tables defined for TDD carrier aggregation without time-domain bundling or the tables defined for Release 8.

If a UE is configured for PUCCH format 1b with channel selection, implicit ACK/NACK resource allocation is used for dynamically scheduled PDSCH transmissions on the PCC and for cross-carrier scheduling of SCCs from the PCC, as well as for release of an SPS resource. In the case of PDSCH transmissions on SCCs, for non cross-carrier scheduling or for cross-carrier scheduling from an SCC, explicit PUCCH resource allocation is used: a set of resources is configured by RRC signalling, and the PDCCH(s) corresponding to the PDSCH on the SCC(s) indicate resources derived from this set using the TPC field.¹⁷

The rules for implicit PUCCH resource allocation are based on those defined in Release 8. For dynamic scheduling, a set of implicit ACK/NACK resources is derived from the indexes of the first CCEs used for transmission of the corresponding DCI assignments. When time domain bundling is used and the Downlink Assignment Index (DAI – see Section 23.4.3) is equal to either ‘1’ or ‘2’, the PUCCH resources are derived from the combination of the DAI value and the indexes of the first CCEs of the corresponding PDCCHs.

In the case of simultaneous ACK/NACK feedback and SR transmission, a UE that is configured to use PUCCH format 1b with channel selection transmits spatially bundled ACK/NACK feedback on a single SR resource. For TDD, the UE transmits two bits on the SR resource, representing the number of ACKs over both CCs within the bundling window.

ACK/NACK Repetition

ACK/NACK repetition (see Section 16.5) on PUCCH is not supported for carrier aggregation, because it can impact downlink performance and a UE configured for carrier aggregation is generally not assumed to suffer from a transmission power limitation for control signalling.

28.3.2.2 Channel State Information Feedback

Periodic Channel State Information (CSI) feedback (see Section 10.2.1.2) is independently configured for each downlink CC by RRC signalling. With carrier aggregation, periodic CSI is reported for only one downlink CC in any given subframe. Different offsets and periodicities should as far as possible be configured for each CC to aim to minimize collisions between CSI reports of different CCs in one subframe. If a collision of multiple CSI reports does occur, one report is selected for transmission according to defined prioritization rules (the others being dropped):

- First priority is given to the CSI reports that contain a Rank Indicator (RI) or a wideband ‘first’¹⁸ Precoding Matrix Indicator (PMI);
- Second priority is given to other wideband CQI and/or PMI reports;
- The lowest priority is assigned to the sub-band CQI/PMI reports.
- If there are multiple CCs with a report of the same priority, the CC of the serving cell with lowest cell index is prioritized.

¹⁷This indication on the PDCCH does not increase the DCI message size, since these TPC fields are not used for PUCCH power control (the PUCCH power being controlled by the PCC grant).

¹⁸‘First’ refers to the first part of the dual-stage PMI codebook structure introduced in Release 10 for the case of 8 downlink transmit antennas – see Section 29.3.3.

If simultaneous PUSCH and PUCCH transmission is supported by the UE and is enabled, a collision between periodic CSI feedback and ACK/NACK can be resolved by transmitting the CSI feedback on the PUSCH and ACK/NACK on PUCCH. Also, if there is a PUSCH transmission and the uplink control signalling consists only of periodic CSI feedback, the periodic CSI feedback is in this case transmitted on the PUCCH (not the PUSCH as in Release 8).

28.3.2.3 Uplink Control Information (UCI) on PUSCH

If a UE is configured with multiple serving cells and simultaneous PUCCH and PUSCH is not enabled, and there is at least one PUSCH transmission, all Uplink Control Information (UCI) is multiplexed onto a PUSCH. Any periodic CSI feedback is normally¹⁹ transmitted on the PCC PUSCH if it is transmitted. If the PCC PUSCH is not transmitted, the UCI is transmitted on an SCC PUSCH if one is transmitted; if more than one SCC PUSCH transmission occurs in the subframe, the periodic UCI is transmitted on the SCC PUSCH of the serving cell with the lowest cell index.

The ACK/NACK payload size for transmission on the PUSCH is determined from the number of configured downlink CCs, the configured transmission mode for each downlink CC, and, in the case of TDD, the bundling window size and the signalled DAI value in the UL grant.

Aperiodic CSI is reported on a PUSCH if it is triggered by a request in an uplink DCI message or a Random Access Response grant. When a UE is configured with multiple serving cells in carrier aggregation, a CSI request transmitted in the UE-specific PDCCH search space can trigger CSI reports for one or more downlink CCs, as shown in Table 28.1.

Table 28.1: Combinations of downlink CCs for which aperiodic CSI may be triggered.

Value of CSI request field	Meaning
00	No CSI report is triggered
01	CSI report is triggered for the cell on which the trigger is sent
10	CSI is triggered for a first set of serving cells configured by higher layers
11	CSI is triggered for a second set of serving cells configured by higher layers

If aperiodic CSI feedback is triggered using the common search space, the feedback is transmitted for an RRC-configured set of CCs.

In the case of aperiodic CSI triggering when multiple PUSCH transmissions are taking place on different uplink CCs in the same subframe, the UCI is always transmitted on the PUSCH on the CC indicated by the uplink grant containing the aperiodic CSI trigger. UEs are not expected to receive more than one positive aperiodic CSI trigger for a given subframe.

In the case of a collision between periodic CSI and aperiodic CSI for the same or different downlink CCs, the periodic CSI is dropped and only the aperiodic CSI feedback

¹⁹The exceptions relate to PUSCH transmissions that are part of the random access procedure.

is transmitted. This applies even if the periodic and aperiodic CSIs are for different downlink CCs.

28.3.3 Sounding Reference Signals

Sounding Reference Signals (SRSs) can be triggered on any serving cell either by higher-layer signalling or dynamically via the DCI messages for UL grants, as explained in Section 29.2.2. When carrier aggregation is configured, a UE may be configured with SRS parameters for both types of SRS triggering on each serving cell.

The following rules are defined for SRSs in relation to transmissions on multiple CCs:

- SRSs may be transmitted simultaneously on different CCs;
- If a PUCCH transmission on the PCC coincides with SRS on an SCC, the UE transmits the SRS on the SCC if the PUCCH uses a shortened format, and not otherwise;
- If a PUSCH transmission coincides with SRS on different CCs:
 - PUSCH transmitted in any cell-specifically configured SRS subframe is rate matched around the SRS resources on the same CC;
 - If PUSCH is transmitted in the same SC-FDMA symbol on a different CC from the same UE, the SRS is dropped.

28.3.4 Uplink Timing Advance

As mentioned in Section 28.1, the timing of the PCell and all SCells configured for a UE is expected to be synchronized. A single Timing Advance (TA) command (see Section 18.2) is therefore sufficient to control the UE's uplink transmission timing for all the uplink CCs together. This simplifies the UE implementation since a single time reference for the baseband processing can be used. In addition, for contiguous CCs, a single Inverse Fast Fourier Transform (IFFT) can be used for the generation of the signals for multiple uplink CCs.

It is, however, expected that Release 11 will introduce the possibility of independent TA per CC, which may be beneficial for some scenarios such as the use of frequency-selective repeaters on certain of the configured CCs.

28.3.5 Uplink Power Control

Uplink power control with carrier aggregation follows the same principles as for single carrier transmission in Release 8 (see Section 18.3). When multiple CCs are configured, uplink power control operates independently for each CC. This allows the different operating conditions of each CC (e.g. different frequency bands or different interference scenarios) to be taken into account. The parameters for open-loop power control (P_0 for both PUSCH and PUCCH, to set the operating point, and α , the fractional path-loss compensation factor) are therefore all CC-specific, as are the closed-loop TPC commands and any MCS-dependent offsets. TPC commands in uplink resource grants are applied to the PUSCH on the uplink CC for which the grant applies. The TPC commands in PCC downlink resource assignments

are applied to the uplink PCC on which the corresponding HARQ ACK/NACK signalling is transmitted.²⁰

Power control for groups of UEs using DCI Formats 3 and 3A is supported only for the same CC on which the TPC commands are transmitted. Cross-carrier scheduling of grouped power control commands is not supported, since SPS and periodic CSI reporting, which are some of the main uses for group power control, take place only on the PCC.

28.3.5.1 Path-Loss Estimation

Since the uplink transmission power is based on the path-loss estimated on a downlink CC, a reference downlink CC is defined for each uplink CC. The path-loss reference for an uplink CC can be either the SIB2-linked downlink CC or the downlink PCC, according to network configuration. The downlink CC used for path-loss estimation should always be in the same frequency band as the uplink CC.

Configurability of the path-loss reference allows appropriate operation for different deployment scenarios. For example, in heterogeneous network deployments it may happen that reliable uplink transmission on an SCC is possible, but the path-loss estimation on the SIB2-linked downlink SCC is not sufficiently reliable due to interference. The ability to configure the path-loss to be estimated on the downlink PCC for power control of the transmissions on such an uplink SCC can facilitate efficient resource usage and better load balancing among CCs.

A configured but deactivated CC can be used as the path-loss reference according to the configuration (SIB2 or PCC), in order to provide reasonably good path-loss estimation for power control upon SCell reactivation. The UE would measure a deactivated CC less frequently (similarly to the case of long DRX cycles) in order not to reduce too significantly the potential power savings from the CC deactivation.

28.3.5.2 Maximum Power Behaviour

As in Release 8, the UE's transmission power is limited by the power class to which it belongs (e.g. 23 dBm for class 3 – see Section 21.3.1.2). In addition, for carrier aggregation a maximum transmission power is set for each individual CC; this must be satisfied first, although this is also typically 23 dBm.

The possibility of simultaneous PUCCH and PUSCH transmission also affects the behaviour of the UE when it reaches its maximum allowed transmission power. Since the control information is essential for the correct reception of data, and, unlike PUSCH transmissions, the PUCCH cannot benefit from HARQ, any PUCCH transmission is always prioritized over concurrent PUSCH transmissions. Hence, the required power is first set for the PUCCH, and then any remaining power is used for PUSCH transmissions. Similarly, among the PUSCH transmissions, a PUSCH carrying uplink control information is prioritized over PUSCH transmissions without it. When a UE reaches its maximum power, it therefore first scales down the power of the PUSCHs without control information. The scaling factor is normally common for all serving cells, although for some serving cells it may be set to

²⁰Note that TPC commands in SCC downlink resource assignments may be used for PUCCH resource assignment – see Section 28.3.2.1.

zero (e.g. if the power after scaling falls below a useful level). Full details can be found in [3, Section 5.1.1.1].

Other factors affecting the UE's maximum transmission power are discussed in Section 21.3.1.2.

28.3.5.3 Power Headroom Reporting

For a UE configured with multiple uplink CCs, the Power Headroom Reports (PHRs) are independent for each CC. PHRs are used to provide the serving eNodeB with information about the difference between the nominal UE maximum transmission power and the estimated power for PUSCH transmission, and the difference between the nominal UE maximum power and the estimated power for PUSCH and PUCCH transmission on the PCell.

To support simultaneous PUCCH and PUSCH transmission, two types of PHR (known as extended PHRs (ePHR)) are supported in Release 10:

- Type 1, which only takes into account the PUSCH transmission power;
- Type 2, which indicates the power headroom when both PUCCH and PUSCH are present.

For a UE configured with simultaneous PUCCH and PUSCH, a Type 2 PHR for the PCC is always reported when Type 1 PHR is reported. For subframes where PUCCH is not actually transmitted, a hypothetical PUCCH format 1a transmission is assumed. Similarly, the UE assumes a hypothetical reference format for the PUSCH when reporting a PHR if no PUSCH transmission is scheduled on the PCC. PHRs based on these hypothetical reference formats are known as virtual PHRs.

In order to provide sufficient information on the total available UE transmission power, each Release 10 ePHR additionally includes the value of the current maximum power for the CC, $P_{\text{CMAX},c}$, after taking into account the maximum power reductions explained in Section 21.3.1.2. A single MAC control element contains both $P_{\text{CMAX},c}$ and the PHR for a given CC; a bitmap is used to indicate for which SCC the information is reported, and a 'virtual PHR indication' indicates whether the PUSCH power is based on an assumed reference format or not. $P_{\text{CMAX},c}$ is not reported for virtual PHRs.

The use of the new Release 10 ePHR is configured by RRC signalling. The support of ePHR is mandatory for UEs supporting uplink carrier aggregation and for UEs supporting simultaneous PUSCH and PUCCH transmission.

ePHR is calculated based on the power before any power scaling due to maximum power limitations and can therefore be positive or negative (similar to the single-carrier PHR of Release 8 – see Section 18.3.3). There is one prohibit timer per UE to control how often the PHR can be transmitted. The prohibit timer is started after a PHR transmission, and a new PHR cannot be transmitted until the timer expires.

In addition to the PHR triggers described in Section 18.3.3, the activation of an SCell with a configured uplink CC triggers a PHR. This is useful for the eNodeB scheduler, since the power state of a UE changes upon activation of a CC.

28.3.6 Uplink Multiple Access Scheme Enhancements

Although not directly related to carrier aggregation, some further enhancements to the uplink multiple access scheme are introduced in Release 10 for LTE-Advanced.

28.3.6.1 Analysis of Candidate Enhancements

As described in Chapter 14, the uplink multiple access scheme of LTE in Releases 8 and 9 is Single Carrier Frequency Division Multiple Access (SC-FDMA) (also known as DFT-Spread-OFDM (DFT-S-OFDM)). This has the desirable property of being ‘single carrier’ and therefore maintaining a low Cubic Metric (CM). The main motivations for modifying this scheme for LTE-Advanced in Release 10 were the possibilities of performance improvements, especially in conjunction with the introduction of uplink Single-User Multiple-Input Multiple-Output (SU-MIMO), and the opportunity to increase the flexibility of uplink resource allocation to maximize the utilization of the spectrum.

Figures 28.11 and 28.12 show block diagrams of two candidate schemes considered for LTE-Advanced, namely clustered DFT-S-OFDM and multiple SC-FDMA respectively.




Figure 28.11: Block diagram of clustered DFT-S-OFDM.




Figure 28.12: Block diagram of multiple SC-FDMA.

Clustered DFT-S-OFDM retains a single DFT operation but modifies the resource element mapping at the output of the DFT operation from a single cluster (as used for SC-FDMA) to multiple clusters which are multiplexed with $N - M$ zeros to form the input of the IFFT²¹ operation over N -virtual subcarriers. The resulting waveform is no longer single-carrier but still has a low CM.

Multiple SC-FDMA consists of a number of DFT operations, where the x^{th} DFT is of size M_x ; the output symbols are then multiplexed with $N - \sum_{x=1}^X M_x$ zeros to fit the N -point IFFT. This waveform is no longer single carrier and experiences a worse CM than that of clustered DFT-S-OFDM for the same number of clusters.

²¹Inverse Fast Fourier Transform.

Figure 28.13 shows the CM of the schemes described above, compared to the SC-FDMA of Release 8 and to OFDMA, for different numbers of clusters and different modulation schemes. It can be seen that the CM for SC-FDMA and QPSK is the lowest, while the CM for OFDMA is the largest and invariant with respect to the number of clusters. The CM for clustered DFT-S-OFDM and multiple SC-FDMA increases with the number of clusters but never gets worse than that of OFDMA. The CM of multiple SC-FDMA is always higher than that of clustered DFT-S-OFDM for the same number of clusters.




Figure 28.13: CM of SC-FDMA, clustered DFT-S-OFDM, multiple SC-FDMA and OFDMA.

Figure 28.14 shows the UE throughput Cumulative Distribution Function (CDF) for SC-FDMA, clustered DFT-S-FDMA and OFDMA. The simulation is for scenario D1 from [4] with 10 MHz system bandwidth, two RBs reserved for PUCCH transmission, a sub-band size of 6 RBs for scheduling and feedback reporting, Proportional Fair Scheduling (PFS), a maximum of 4 uplink grants per subframe, 10 UEs/cell each with a 1×2 antenna configuration, ideal channel estimation and Interference over Thermal (IoT) target of 7 dB. Different power backoffs are modelled in the simulation according to the difference in CM of the various schemes (see Figure 28.13). The figure shows that the performance of OFDMA and clustered DFT-S-OFDMA is practically identical. Both these schemes outperform SC-FDMA thanks to the additional flexibility in resource allocation which maximizes utilization of the spectrum.




Figure 28.14: System-level performance comparison between SC-FDMA, clustered DFT-S-OFDMA and OFDMA.

28.3.6.2 Enhancements Included in Release 10

As a result of these considerations, DFT-S-OFDM continues to be the basis of the uplink multiple access scheme for the PUSCH in LTE-Advanced, but the possibility of frequency-non-contiguous resource allocation is introduced for an individual Release 10 UE, using clustered DFT-S-OFDM with a maximum of two clusters.

The uplink grant on the PDCCH (i.e. DCI formats 0 or 4 – see Section 9.3.5.1) indicates whether the PUSCH resource allocation is multiclustered or not, by means of a *resource allocation type* bit. If this bit is set, in the case of DCI format 0 the ‘frequency hopping flag’ is used as an extra bit for the multiclustered resource allocation signalling. In general, PUSCH frequency hopping is not supported in conjunction with multiclustered PUSCH transmission.

In order to signal the frequency-domain locations of the RB Groups (RBGs) of the non-contiguous resource allocations, the indexing scheme used for signalling the CQI subbands in Release 8 is re-used (see Section 10.2.1.1). The RBG size depends on the system bandwidth in the same way as for downlink resource allocation type 0 (see Section 9.3.5.4 and Table 9.5).

If the resource allocation is non-contiguous, the DeModulation Reference Signals (DM-RSs) transmitted in the PUSCH (see Section 15.5) are adapted to match the resource allocation: a single DM-RS base sequence is generated according to the total number of allocated RBs, and then split into sections for transmission in the PUSCH RBs.

In the case of carrier aggregation, one DFT is used per CC, thus yielding a multiple SC-FDMA scheme. If a different Power Amplifier (PA) is used for each CC, each PA will amplify a single-carrier waveform and hence benefit from the low CM (unless simultaneous PUCCH and PUSCH transmission is used).

28.4 UE Transmitter and Receiver Aspects

In LTE-Advanced Release 10 the spacing between the centre frequencies of contiguously aggregated CCs is a multiple of 300 kHz. The rationale behind this choice is to preserve backward compatibility with the 100 kHz frequency raster used in LTE Release 8 as well as preserving the orthogonality of the subcarriers with the 15 kHz spacing. Depending on the aggregation scenario, the actual spacing (a multiple of 300 kHz) may be facilitated by insertion of a number of unused subcarriers between contiguous CCs.

28.4.1 UE Transmitter Aspects of Carrier Aggregation

The output power dynamics are impacted by the UE architecture, which may be based on single or multiple PAs. Figure 28.15 [2] illustrates various options for PA architectures at the UE which can be used to support carrier aggregation.

When considering the PA configuration, it is necessary to take into account any additional back-off requirements that may exist. The CM, introduced as a predictor of required back-off in Section 21.3.3, is only a good predictor of the additional power back-off required if the third-order InterModulation (IM3) distortion product lands in the Adjacent Channel Leakage Ratio (ACLR)²² band (as it does, for instance, for LTE Release 8 with full resource allocation, or for WCDMA-based system such as UMTS and HSPA – see Chapter 21).

The new multiple SC-FDMA and clustered DFT-S-OFDM waveforms supported in Release 10 (due to carrier aggregation and the concurrent transmission of PUSCH and PUCCH) impose more stringent linearity requirements on the PA than was the case for LTE Release 8.

The factors that determine the necessary UE PA back-off are compliance to the ACLR, Spectrum Emission Mask (SEM), spurious emissions and Error Vector Magnitude (EVM) requirements [2, 5].

Small resource assignments at the band edge behave as tones and hence produce highly concentrated InterModulation Distortion (IMD) products. Therefore, for the concurrent transmission of PUCCH and PUSCH, the SEM is expected to be the limiting requirement.

28.4.2 UE Receiver Aspects of Carrier Aggregation

For the baseband aspects of the UE receiver, the main impact of carrier aggregation is on the soft buffer allocation, where the total HARQ buffer has to be shared between the configured CCs.

For the RF aspects, two options were considered for the baseline UE receiver architecture as part of the carrier aggregation feasibility study:

- **Option A:** Single RF, and baseband processing with bandwidth ≥ 20 MHz;
- **Option B:** Multiple RF, and baseband processing with bandwidth ≤ 20 MHz.

Clearly, Option A is only applicable for intra-band aggregation of contiguous CCs, but it has the advantage of keeping the UE receiver complexity low. Option B is applicable for intra-band and inter-band aggregations for contiguous or non-contiguous scenarios, but this flexibility comes at the expense of increased complexity.

²²See Chapter 21.





	Transmitter architecture	Aggregation Scenarios
A	 <p>Single (baseband + IFFT + DAC + mixer + PA)</p>	Intra-band contiguous CCs
B	 <p>Multiple (baseband + IFFT + DAC), single (stage-1 IF mixer + IF combiner + stage-2 RF mixer + PA)</p>	Intra-band contiguous CCs Intra-band non-contiguous CCs
C	 <p>Multiple (baseband + IFFT + DAC + mixer), low-power RF combiner, and single PA</p>	Intra-band contiguous CCs Intra-band non-contiguous CCs
D	 <p>Multiple (baseband + IFFT + DAC + mixer + PA), high-power combiner to single antenna OR dual antenna</p>	Intra-band contiguous CCs Intra-band non-contiguous CCs Inter-band non-contiguous CCs (depending on the specific bands being aggregated)

Figure 28.15: Some examples of PA configuration options for carrier aggregation.
Reproduced by permission of © 3GPP.

28.4.3 Prioritized Carrier Aggregation Scenarios

Many carrier aggregation scenarios of relevance to different operators around the world were considered as part of the feasibility study for LTE-Advanced [2].

In order to focus the work to define RF requirements in 3GPP, some carrier aggregation combinations were prioritized for the timeframe of Release 10, based on the priorities of network operators.

For intra-band carrier aggregation, the first supported carrier bandwidths are 15 and 20 MHz in E-UTRA Band 1, and 10, 15 and 20 MHz in Band 40.²³ A maximum of two aggregated carriers are supported initially.

For inter-band carrier aggregation, the first defined scenarios are likely to include:

- E-UTRA Bands 1 and 5 for 10 MHz CCs, one CC per band;
- E-UTRA Bands 3 and 7 for 10, 15, and 20 MHz CCs, one CC per band;
- E-UTRA Bands 4 and 13 for 10 MHz CCs, one CC per band;
- E-UTRA Bands 4 and 17 for 10 MHz CCs, one CC per band.

Other combinations will be added in a release-independent manner, so that they can be implemented by UEs of any release from Release 10 onwards.

28.5 Summary

There are three main motivations for introducing carrier aggregation for LTE-Advanced in Release 10: support of high data rates, efficient utilization of fragmented spectrum, and support of heterogeneous network deployments by means of cross-carrier scheduling. In addition, enhancements to the uplink transmission scheme are included, allowing clustered DFT-S-OFDM transmission with non-contiguous resource allocation within one CC, and simultaneous PUCCH/PUSCH transmission.

References²⁴

- [1] 3GPP Technical Report 36.913, 'Requirements for further advancements for Evolved Universal Terrestrial Radio Access (E-UTRA) (LTE-Advanced)', www.3gpp.org.
- [2] 3GPP Technical Report 36.815, 'LTE-Advanced Feasibility Studies in RAN WG4 (Release 9)', www.3gpp.org.
- [3] 3GPP Technical Specification 36.213, 'Physical Layer Procedures (Release 10)', www.3gpp.org.
- [4] 3GPP Technical Report 25.814, 'Physical Layer Aspects for Evolved Universal Terrestrial Radio Access (UTRA) (Release 7)', www.3gpp.org.
- [5] 3GPP Technical Specification 36.101, 'Evolved Universal Terrestrial Radio Access (E-UTRA); User Equipment (UE) Radio Transmission and Reception (Release 10)', www.3gpp.org.

²³See Section 21.2 for the E-UTRA band definitions.

²⁴All web sites confirmed 1st March 2011.

Multiple Antenna Techniques for LTE-Advanced

Alex Gorokhov, Amir Farajidana, Kapil Bhattad, Xiliang Luo and Stefan Geirhofer

Multiple antenna techniques play a key role in LTE-Advanced. In the downlink, the goals are to support higher data rates than LTE Releases 8 and 9 through high-order Single-User Multiple-Input Multiple-Output (SU-MIMO), and higher spectral efficiency via enhanced Multi-User MIMO (MU-MIMO) techniques. To support these advances, new reference signals and enhanced UE feedback are introduced. In the uplink, SU-MIMO is introduced, and the control channel performance is enhanced using transmit diversity.

This chapter explains the techniques adopted for Release 10 and gives some insight into additional MIMO enhancement features that may be developed in future releases of LTE-Advanced.

29.1 Downlink Reference Signals

As discussed in detail in Section 8.2, LTE Release 8 provides cell-specific Reference Signals (RSs), also known as Common Reference Signals (CRSs) for up to 4 antenna ports. Cell-specific RSs are used by UEs both to perform channel estimation for demodulation of data and to derive feedback on the quality and spatial properties of the downlink radio channel. Together with signalling of the precoder used for data transmissions to the UE, the four cell-specific RS ports enable spatial multiplexing of up to four layers using codebook-based precoding, as explained in Section 11.2.2.2. LTE Release 9 additionally supports two-layer beamforming spatial multiplexing using precoded UE-specific RSs, which enable non-codebook-based precoding to be used, as explained in Section 11.2.2.3.

LTE – The UMTS Long Term Evolution: From Theory to Practice, Second Edition.

Stefania Sesia, Issam Toufik and Matthew Baker.

© 2011 John Wiley & Sons, Ltd. Published 2011 by John Wiley & Sons, Ltd.

For LTE-Advanced, downlink SU-MIMO transmission is extended to support up to eight spatial layers, and for this purpose the precoded UE-specific RS approach is further developed for the data demodulation. This has the advantage that transmission of the new Release 10 UE-specific RSs can be limited to only those Resource Blocks (RBs) where they are needed for demodulation, thus avoiding incurring a large overhead across the whole system bandwidth (as would have been the case with an extension to the cell-specific RSs, which would also have adversely impacted ‘legacy’ pre-Release 10 UEs that are unaware of the new RS structure).

In addition, in order to enable the UE to estimate and feed back the Channel State Information (CSI) corresponding to up to eight antenna ports across a wide bandwidth, new RSs referred to as CSI-RSs are provided. Since CSI-RSs are used only for feedback purposes they can be sparse and incur only a small overhead.

These two new types of RS are explained in detail in the following subsections.

29.1.1 Downlink Reference Signals for Demodulation

In the same way as the UE-specific RSs of earlier releases (see Sections 8.2.3 and 8.2.3), the extended UE-specific RSs of LTE-Advanced are embedded in the RBs used for the Physical Downlink Shared CHannel (PDSCH) for a specific UE. The UE-specific RSs for each layer undergo the same precoding as the data symbols, and therefore there is no need for explicit signalling of precoding information to the UE. A variety of multi-antenna beamforming techniques can therefore be supported efficiently and transparently to the UE.

The pattern of Resource Elements (REs) designed for the new UE-specific RSs had to satisfy certain criteria. Firstly, it had to avoid overlapping with the cell-specific RSs and the control channels in order to ensure backward compatibility. Secondly, the UE-specific RSs of different layers should be orthogonally multiplexed to avoid inter-layer RS interference degrading the channel estimation accuracy; this could in principle be achieved by assigning different sets of REs to the RSs of different layers (by Frequency-Division and/or Time-Division Multiplexing FDM/TDM), and/or by Code Division Multiplexing (CDM). CDM has the advantages of more flexible power balancing across different layers and potentially better interference estimation accuracy in MU-MIMO scenarios.

For Release 10, the UE-specific RS pattern for up to 2 layers (referred to as ‘rank-2’ transmission) is identical to that defined in Release 9, in order to ensure backward compatibility. For up to 4 layers, the pattern is obtained by extending the Release 9 rank-2 UE-specific RS pattern in a hybrid CDM/FDM fashion, as shown in Figure 29.1(a) (for the case of the normal Cyclic Prefix (CP) length)¹. The four precoded layers (antenna ports 7–10) are grouped into two groups of 2 REs, with the same length-2 Walsh–Hadamard Orthogonal Cover Codes (OCC) as in Release 9 (i.e. [1, 1] and [1, –1]) being used to multiplex the layers within each group. The UE-specific RSs in different groups are frequency-multiplexed on adjacent subcarriers. This pattern has been shown to provide reasonable performance for a variety of channel conditions and UE speeds [2, 3].

¹By following the same principles, similar patterns are defined in Release 10 for the case of the extended CP and for the Downlink Pilot TimeSlot (DwPTS) field of the special mixed downlink-uplink subframe in TDD operation – see [1, Section 6.10.3]. Use of the extended CP with more than two UE-specific antenna ports is not supported in Release 10.




Figure 29.1: Release 10 UE-specific RS patterns.

For 8-layer transmission (antenna ports 7–14), the UE-specific RS structure of Figure 29.1(a) is further extended using hybrid CDM/FDM with two CDM groups each using a Walsh–Hadamard OCC of length 4 as shown in Figure 29.1(b). Exactly the same set of REs are used as for the rank-4 case, with the UE-specific RSs for each layer now being spread across four REs, all having the same frequency location but different time locations within the subframe. This approach maintains the power-balancing property of the rank-4 design and is optimized for pedestrian mobility – the most likely scenario for 8-layer SU-MIMO.

The length-2 and length-4 OCCs have a nested structure (the OCC of length two is nested into the OCC of length four) which ensures backward compatibility with the Release 9 design.

The mapping of the OCCs to REs is shown in Figure 29.2 for the case of normal CP. The mapping for odd and even RBs (in the frequency domain) is different in order to provide orthogonality in both time and frequency. Such orthogonality improves the performance of channel estimation in higher Doppler scenarios and reduces the inter-subcarrier interference across the two CDM groups in the presence of frequency offsets. This OCC mapping also enables power balancing across OFDM² symbols by peak power randomization [4].

The RS sequence for up to 8 layers is the same as in Release 9 (as described in Section 8.2.3), except that for antenna ports 9–14 only one sequence initialization is possible.

²Orthogonal Frequency Division Multiplexing.

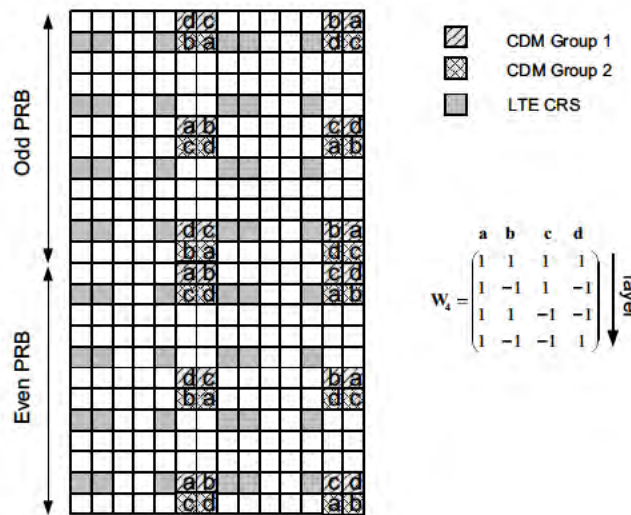


Figure 29.2: Length-4 OCC mapping to REs.

29.1.2 Downlink Reference Signals for Estimation of Channel State Information (CSI-RS)

The main goal of CSI-RSs is to obtain channel state feedback for up to eight transmit antenna ports to assist the eNodeB in its precoding operations. Release 10 supports transmission of CSI-RS for 1, 2, 4 and 8 transmit antenna ports. CSI-RSs also enable the UE to estimate the CSI for multiple cells rather than just its serving cell, to support future multicell cooperative transmission schemes (see Section 29.5.1).

The following general design principles can be identified for the CSI-RS:

- In the frequency domain, uniform spacing of CSI-RS locations is highly desirable, as explained in Chapter 8.
- In the time domain, it is desirable to minimize the number of subframes containing CSI-RS, so that a UE can estimate the CSI for different antenna ports and even different cells with a minimal wake-up duty cycle when the UE is in Discontinuous Reception (DRX) mode, to preserve battery life.
- The overall CSI-RS overhead involves a trade-off between accurate CSI estimation for efficient operation and minimizing the impact on legacy pre-Release 10 UEs which are unaware of the presence of CSI-RS and whose data are punctured by the CSI-RS transmissions. Figure 29.3 shows that a CSI-RS density of one RE per RB per antenna port is a good choice, as the throughput degradation compared to ideal CSI estimation is negligible.
- CSI-RSs of different antenna ports within a cell, and, as far as possible, from different cells, should be orthogonally multiplexed to enable accurate CSI estimation.

- To ensure backward compatibility, CSI-RSs should avoid REs used for cell-specific RSs and control channels, as well as avoiding REs used for the Release 10 UE-specific RSs.




Figure 29.3: Throughput performance with a CSI-RS density of 1 RE per RB per antenna port: uncorrelated 8x2 SU-MIMO, 5 MHz, Typical Urban channel model, 3 km/h.

Taking these considerations into account, the CSI-RS patterns selected for Release 10 are shown in Figure 29.4. CDM codes of length 2 are used, so that CSI-RSs on two antenna ports share two REs on a given subcarrier.

The pattern shown in Figure 29.4(a) can be used in both frame structure 1 (FDD)³ and frame structure 2 (TDD).⁴ In Figure 29.4, the REs used for CSI-RSs are labelled using two letters, the first indicating the cell index and the second referring to the antenna ports of the CSI-RS transmitted on that RE. These patterns follow a ‘nested’ structure, meaning that the REs used in the case of two CSI-RS antenna ports are a subset of those used for four and eight antenna ports; this helps to simplify the implementation. The total number of supported antenna ports is 40, which can be used to give a frequency-reuse factor of 5 between cells with 8 antenna ports per cell, or a factor of 20 in the case of 2 antenna ports. It can be seen that collisions may occur with REs used for the UE-specific RS antenna ports defined in Release 8 for PDSCH transmission mode 7 (see Section 8.2.2); it may therefore be desirable to avoid scheduling UEs in transmission mode 7 in subframes containing CSI-RS.

The pattern shown in Figure 29.4(b) can only be used for frame structure 2 for TDD operation. This pattern is designed to avoid collisions with the Release 8 UE-specific RS antenna port, as this port is more suited to TDD operation where channel reciprocity can be more effectively exploited to support the beamforming. However, this pattern only offers frequency reuse factors of between 3 and 12 depending on the number of antenna ports per cell.

³Frequency Division Duplex.

⁴Time Division Duplex.

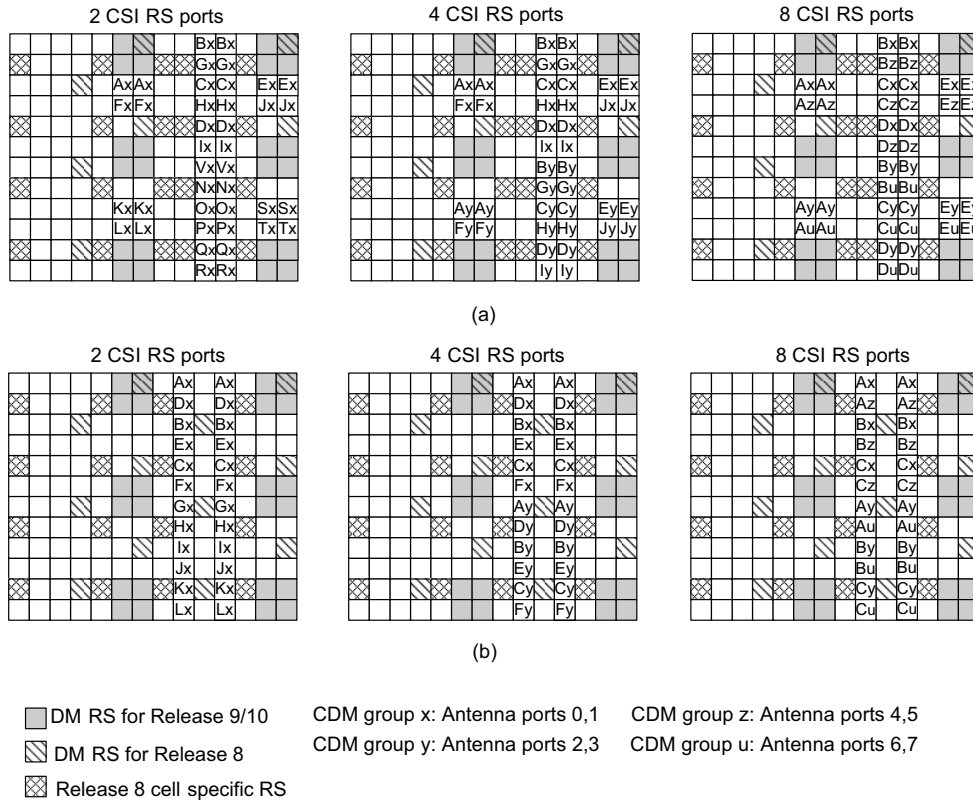


Figure 29.4: CSI-RS patterns for LTE-Advanced.

Similar CSI-RS patterns are provided for the case of the extended CP and the DwPTS field of the special mixed downlink-uplink subframe in frame structure 2 (see [1, Section 6.10.5]).

The CSI-RS configuration is UE-specific. When configured, CSI-RSs are present only in some specific subframes following a given duty cycle and subframe offset. The duty cycle and offset of the subframes containing CSI-RSs and the CSI-RS pattern used in those subframes are provided to a Release 10 UE through RRC signalling. The duty cycle and subframe offset are jointly coded, while the CSI-RS pattern is configured independently of these two parameters.

It should be noted that, in subframes containing CSI-RSs, rate-matching (see Section 10.3.2.4) for PDSCH transmissions to Release 10 UEs assumes that the CSI-RS REs are not available for PDSCH data, and the coded PDSCH data is only mapped to the surrounding REs. However, PDSCH transmissions to Release 8 and 9 UEs are punctured by CSI-RS transmissions. The density and periodicity of the CSI-RSs should therefore be chosen such that the impact of puncturing on the performance of these users is acceptable.

In the context of cooperative MIMO, it may be possible to improve the performance of channel estimation, and especially interference estimation, by coordinating CSI-RS

transmissions across multiple cells. In Release 10 it is therefore possible to ‘mute’ a specific set of REs in data transmissions from a cell. The locations of these REs, known as the ‘muting pattern’, can be chosen to avoid colliding with CSI-RS transmissions from other cells and hence improve the inter-cell measurement quality. The muting pattern is indicated to the UEs by a 16-bit bitmap, where each bit corresponds to a 4-port CSI-RS configuration for frame structure 1 or 2 as shown in Figure 29.4(a). All the REs in an indicated 4-port CSI-RS configuration are muted, and the UE can assume zero transmission power on those REs unless they are used to transmit CSI-RSs.

In subframes in which muting is configured, data transmissions for Release 10 UEs are rate-matched around the muted REs in the same way as for CSI-RSs. Data transmitted to Release 8 and 9 UEs is punctured in the muted REs.⁵

29.2 Uplink Reference Signals

In Chapter 15, the LTE Release 8 uplink RS design is explained, including DeModulation RSs (DM-RSs) and Sounding RSs (SRSs). In Release 10, the uplink DM-RSs are extended to support uplink SU-MIMO transmission with up to four spatial layers, and the SRSs are enhanced to provide improved support for channel sounding.

29.2.1 Uplink DeModulation Reference Signals (DM-RS)

In order to support uplink SU-MIMO transmission with up to 4 spatial layers, DM-RSs for all spatial layers need to be multiplexed together to enable channel estimation of each layer at the eNodeB. Like the downlink UE-specific RSs, the uplink DM-RSs are precoded using the same precoding as the Physical Uplink Shared CHannel (PUSCH) data transmissions in the same RBs when multiple antenna transmission is used.

Release 10 uses different Cyclic Shifts (CSs) of the DM-RS base sequence to multiplex the DM-RSs for the different layers. For example, in a rank-4 SU-MIMO transmission, CS offsets 0, 6, 3 and 9 are used relative to the base sequence, for layers 1, 2, 3 and 4 respectively. In addition to separation of the antenna ports by selection of different CSs, two length-2 OCCs can be applied to the two DM-RS symbols in the two slots of one subframe to further separate the multiplexed DM-RSs. The codes used are $[+1, +1]$ and $[+1, -1]$; they enable better orthogonality to be achieved between DM-RSs that use close CSs. Figure 29.5 shows an example of the link-level throughput performance improvement that can be achieved thanks to the better orthogonality provided by the use of the OCCs in addition to the CS separation.

Different OCCs can also be assigned to different UEs for uplink MU-MIMO transmission. In this way, the DM-RSs from different UEs can be kept orthogonal to each other even if the PUSCH transmission bandwidths differ between the UEs (provided that there is no sequence group hopping (see Section 15.3.1) between the two slots within one subframe).

In order to reduce the control signalling overhead and make the benefits of the OCCs available to both SU-MIMO and MU-MIMO UEs, the identities of the OCCs applied to each spatial layer are implicitly derived from the 3-bit cyclic shift index signalled in the corresponding PDCCH uplink grants (see Section 9.3.5.1). The cyclic shift indices are divided into two groups: one group of 4 cyclic shift indices is targeted for SU-MIMO, and

⁵See [5, 6] for more background.




Figure 29.5: The benefit of the uplink DM-RS OCCs in terms of link throughput.

the layers with close CSs are assigned different OCCs; the other group is optimized for MU-MIMO, and the same OCC is applied to the different layers.

29.2.2 Sounding Reference Signals (SRSs)

The SRSs are important for uplink channel sounding to support dynamic uplink resource allocation, as well as for reciprocity-aided beamforming in the downlink, especially in TDD operation.

As explained in Section 15.6, the Release 8 SRSs allow for uplink sounding from a single transmit antenna only. With the introduction of uplink multiple-antenna transmission in Release 10, SRS transmission from all the uplink antenna ports is needed in order to enable all the spatial dimensions of the channel to be sounded. The same mechanisms as are available in Release 8 are used to separate the SRS transmissions from different antenna ports – i.e. different CSs and different transmission comb indices of the IFDMA⁶ structure.

Nevertheless, the need for SRSs from up to four antenna ports increases the demand on SRS resources if all transmit antennas are to be sounded at a reasonable rate, and therefore semi-static configuration of the resources via higher layer (RRC) signalling is less appropriate than it was for Release 8. Therefore, Release 10 introduces the possibility of dynamically triggering aperiodic SRS transmissions via the Physical Downlink Control Channel (PDCCH); these dynamic aperiodic SRS transmissions are known as ‘type 1’ SRSs, while the Release 8 RRC-configured SRSs described in Section 15.6 are known as ‘type 0’ in Release 10.

⁶Interleaved Frequency Division Multiple Access.

An indicator in an uplink resource grant on the PDCCH can be used to trigger a single type 1 SRS transmission. This facilitates rapid channel sounding to respond to changes in traffic or channel conditions, without tying up SRS resources for a long period. In DCI Format 0, one new bit can indicate activation of a type 1 SRS⁷ according to a set of parameters that is configured beforehand by RRC signalling. In DCI Format 4, which is used for scheduling uplink SU-MIMO transmissions, two new bits allow one of three sets of RRC-configured type 1 SRS transmission parameters to be triggered (the remaining state indicates no type 1 SRS activation).

In the case of uplink multiple-antenna transmission (see Section 29.4), SRSs may be triggered on multiple antenna ports simultaneously.

Other enhancements to the SRSs, including coordination of SRS transmission resource configurations and frequency-hopping patterns, may also be considered for cooperative MIMO schemes and inter-cell interference coordination in the future (see, for example, [7] and references therein).

29.3 Downlink MIMO Enhancements

The main enhancements to the downlink multiple antenna transmission schemes in Release 10 are the extension of SU-MIMO to support 8-layer transmission, and improved support for MU-MIMO. These are introduced in a new PDSCH transmission mode, known as Transmission Mode 9.

29.3.1 Downlink 8-Antenna Transmission

The new RSs to support transmissions from 8 antenna ports in Release 10 are explained in Section 29.1.

For the data transmissions from 8 antenna ports, the Release 8/9 transmit diversity and spatial multiplexing modes are extended. For transmit diversity, no new schemes are specified in Release 10. Instead, if an eNodeB has more than four antennas and wishes to use transmit diversity, the Release 8 SFBC/FSTD⁸ schemes are re-used via ‘antenna virtualization’; this means that multiple antennas are treated as a single port by the application of suitable precoding, with a single Release 8 cell-specific RS being transmitted from all the antennas comprising the antenna port. The antenna virtualization is thus transparent to the UE.

For 8-layer spatial multiplexing, two aspects of the codeword-to-layer mapping needed to be considered, namely the number of codewords and the mapping of the codewords to layers. An increase in the number of codewords from the two supported in Release 8 could in theory have provided a throughput gain for some receivers such as those based on Minimum Mean Squared Error Serial Interference Cancellation (MMSE-SIC), but on the other hand a larger signalling overhead would have been incurred in both downlink and uplink, and the complexity of link adaptation would have been increased. Figure 29.6 shows the throughput performance of an 8×8 SU-MIMO transmission scheme with different numbers of codewords for a SIC receiver and a Linear MMSE (LMMSE) receiver. It can be seen that the performance gain from increasing the number of codewords beyond two is small even for the MMSE-SIC receiver.

⁷Only applicable if the DCI message is sent in the UE-specific search space – see Section 9.3.5.5.

⁸Space-Frequency Block Code / Frequency Switched Transmit Diversity (SFBC/FSTD).




Figure 29.6: Throughput performance of LMMSE and MMSE-SIC receivers with different numbers of codewords.

Therefore Release 10 continues to use a maximum of two codewords even in the case of eight antenna ports. The codeword-to-layer mapping for 5 to 8 layers is illustrated in Figure 29.7. Note that, for up to 4 antenna ports, the mapping is the same as in Release 8.




Figure 29.7: Codeword-to-layer mapping for up to 8 antenna ports.

As noted in Section 29.1, the Release 10 high-rank PDSCH transmissions use UE-specific RSs. This means that they are not constrained to being transmitted in subframes where the Release 8 cell-specific RSs are transmitted, and PDSCH transmission mode 9 in Release 10 therefore allows the PDSCH to be transmitted in the data region of ‘MBSFN’⁹ subframes, provided that those subframes are not actually used for MBMS¹⁰ transmissions.

The use of non-codebook-based precoding for the PDSCH means that the eNodeB has a substantial degree of freedom in determining the actual precoding to apply for each UE

⁹Multimedia Broadcast Single Frequency Network.

¹⁰Multimedia Broadcast/Multicast Service.

and for each RB. However, in some cases the eNodeB may use the same precoding matrix across several RBs. For example, when a UE is configured to feed back sub-band Precoding Matrix Indicators (PMIs), the eNodeB is quite likely to use the same precoding matrix across the RBs in a sub-band, in the absence of other relevant information. In this case, the UE can improve its channel estimation quality if it has information about the RBs on which the same precoding matrix is applied. Therefore, in Release 10 it is specified that when the UE is configured to feed back PMI, it can assume that the precoder is the same across all the RBs within a *Precoding Resource block Group* (PRG). In resource allocation type 0 (see Section 9.3.5.4), the data allocation to the UE is performed in groups of contiguous RBs called Resource Block Groups (RBG); in this case, the PRG size is chosen such that it is a factor of the RBG size, in order to be consistent with the resource allocations to individual UEs. The sizes of the PMI sub-bands, RBGs and PRGs for different bandwidths are shown in Table 29.1.

Table 29.1: PRG sizes.

System bandwidth	Sub-band size	RBG size	PRG size
<10	4	1	1
11–26	4	2	2
27–63	6	3	3
64–110	8	4	2

29.3.2 Enhanced Downlink Multi-User MIMO

MU-MIMO was introduced in LTE Releases 8 and 9 as explained in Section 11.1.4. Further enhancements to MU-MIMO operation are included in PDSCH transmission mode 9 in Release 10.

Like PDSCH transmission mode 8 in Release 9, transmission mode 9 in Release 10 supports dynamic switching between SU-MIMO and MU-MIMO. The provision of UE-specific RS patterns with rank greater than one enables the scheduling of more than one spatial stream to UEs participating in MU-MIMO. Cell-specific scrambling of the UE-specific RS sequences, as well as orthogonal multiplexing of UE-specific RS ports assigned to co-scheduled UEs, enables accurate channel estimation and improved interference nulling at the UE.

Although Release 10 supports orthogonal multiplexing of up to 8 layers, the maximum intended multi-user multiplexing order is 4 in Release 10. The gains, if any, of supporting more than 4 layers with MU-MIMO are expected to be small since the channel estimation quality is expected to drastically decrease, and feedback inaccuracies make efficient selection of UEs for co-scheduling difficult.

Another salient aspect of MU-MIMO that was introduced in LTE Release 9 and is carried over into LTE-Advanced is the transparency of MU-MIMO operation from the point of view of the UE. Specifically, no signalling is provided to inform UEs of the presence of

co-scheduled transmissions, nor of any parameters of such transmissions (such as rank, antenna port, RB allocation or power offset); in particular, the use of UE-specific RSs whose amplitude can be varied along with the data symbols means that the Release 8 signalling of a power offset to indicate that the eNodeB's transmit power in an RB is divided between multiple UEs is no longer necessary. Control signalling is limited to indicating the set of UE-specific RS ports allocated to the UE itself. As in Release 9, the 4-layer multiplexing is achieved by using two scrambling sequences and the rank-2 UE-specific RS pattern.

A new DCI format (Format 2C) is introduced in PDSCH transmission mode 9 to support dynamic switching between SU-MIMO transmission (up to rank 8) and MU-MIMO transmission (up to rank 4). DCI Format 2C is based on Format 2B with the addition of 2 new bits; also, the SCrambling IDentity (SCID) bit of Format 2B is reused to encode jointly the number of layers, UE-specific RS SCID and antenna port mapping. The encoding is shown in Table 29.2.¹¹

Table 29.2: Encoding of antenna ports, layers and SCID in DCI Format 2C.

One codeword: Codeword 0 enabled, Codeword 1 disabled		Two codewords: Both codewords enabled	
State	Message	State	Message
0	1 layer, port 7, SCID=0	0	2 layers, ports 7–8, SCID=0
1	1 layer, port 7, SCID=1	1	2 layers, ports 7–8, SCID=1
2	1 layer, port 8, SCID=0	2	3 layers, ports 7–9
3	1 layer, port 8, SCID=1	3	4 layers, ports 7–10
4	2 layers, ports 7–8	4	5 layers, ports 7–11
5	3 layers, ports 7–9	5	6 layers, ports 7–12
6	4 layers, ports 7–10	6	7 layers, ports 7–13
7	Reserved	7	8 layers, ports 7–14

It is also well known that the system capacity of MU-MIMO is highly dependent on the accuracy of the CSI feedback from the UEs. Therefore, improved feedback is also a key factor for improved MU-MIMO performance; this is discussed in the next section.

29.3.3 Enhanced CSI Feedback

In Release 10, the same CQI/PMI/RI¹² feedback types are used as in Release 8. As explained in Section 11.2.3, this is often described as *implicit* feedback, as it provides an implicit representation of the channel consisting of an indication of the data rate that could be achieved if the eNodeB used a certain precoder.

¹¹Note that transmissions with more than one layer but a single codeword only occur in the case of retransmissions of SU-MIMO transmissions. Hence, for transmission with one codeword and 2 layers, only one scrambling sequence is used.

¹²Channel Quality Indicator / Precoding Matrix Indicator / Rank Indicator.

The concept of implicit (i.e. recommended precoder) feedback can be compared to what is sometimes called *explicit feedback* (not supported in LTE or LTE-Advanced), whereby a UE would instead explicitly report a quantized representation of the physical CSI without making assumptions about the nature of the eNodeB precoder. Since the capacity achieved by MU-MIMO is very dependent on the accuracy of the CSI at the transmitter, explicit feedback mechanisms are often favoured in theoretical studies. However, the use of implicit feedback for MU-MIMO in LTE does have some practical advantages, including minimizing the implementation effort required to support both SU- and MU-MIMO. Furthermore, the accuracy of a UE's implicit feedback can very easily be tested, by transmitting data to the UE using the UE's recommended precoder and the Modulation and Coding Scheme (MCS) corresponding to the UE's CQI report.

Moreover, it can be argued that, under certain assumptions, implicit and explicit feedback schemes have similar results. As an example, in the case of rank-1 feedback, an explicit short-term CSI report would signal the direction of the highest channel gain, i.e. the strongest eigenvector of the channel matrix, after some form of quantization to reduce the overhead to a reasonable level. On the other hand, with implicit feedback, a UE would select the vector from the precoding codebook that yields the largest Signal-to-Interference-plus-Noise Ratio (SINR), which is necessarily an approximation of the strongest eigenvector, since the latter identifies the direction of maximum beamforming gain of the channel.

29.3.3.1 PMI Feedback Codebook Enhancements

Enhanced spatial feedback in PDSCH transmission mode 9 is available for the case when 8 CSI-RS ports are configured.

The principle of this enhanced feedback is based on an observation that large antenna arrays at the eNodeB typically comprise two cross-polarized subarrays, in which each subarray is a Uniform Linear Array (ULA). A close spacing of the antenna elements (of the order of half a wavelength) gives rise to fairly strong long-term correlation within a subarray, which is seen in particular in many practical scenarios in macrocellular networks where the angular spread at the eNodeB is relatively small. A high long-term correlation within each subarray can be exploited to reduce the rate of feedback required for each subarray, in time and/or frequency.

In general, the relatively persistent nature of the spatial signatures of the subarrays facilitates robust MU-MIMO scheduling with predictable multi-user interference. Conversely, the low correlation between the cross-polarized subarrays yields improved diversity for SU-MIMO scheduling. Following these principles, the PMI feedback codebooks for 8 CSI-RS antenna ports have been designed such that precoders within each 4-element ULA subarray are represented by a grid of beams, with the number of available beams from which to choose being 32 when 1- or 2-layer transmission is preferred, 16 for 3 or 4 layers, 4 for 5 to 7 layers and only one for 8 layers. Cross-polarized dimensions (corresponding to low correlation) are combined with *co-phasing factors* represented by quaternary symbols, following the same codebook design principle as in Release 8 for uncorrelated antennas.

An example of the resulting Release 10 codebook design for the case of 2-layer transmission being preferred is shown in Table 29.3.3.1. Full details of the PMI feedback codebooks for 8 CSI-RS antenna ports when other transmission ranks are preferred can be found in [1, Section 6.3.4.2.3].

Table 29.3: PMI feedback codebook design for the case of 2-layer feedback with 8 CSI-RS antenna ports.

i_2	0	1	2	3
Codebook	$W_{2i_1,2i_1,0}$	$W_{2i_1,2i_1,1}$	$W_{2i_1+1,2i_1+1,0}$	$W_{2i_1+1,2i_1+1,1}$
i_2	4	5	6	7
Codebook	$W_{2i_1+2,2i_1+2,0}$	$W_{2i_1+2,2i_1+2,1}$	$W_{2i_1+3,2i_1+3,0}$	$W_{2i_1+3,2i_1+3,1}$
i_2	8	9	10	11
Codebook	$W_{2i_1,2i_1+1,0}$	$W_{2i_1,2i_1+1,1}$	$W_{2i_1+1,2i_1+2,0}$	$W_{2i_1+1,2i_1+2,1}$
i_2	12	13	14	15
Codebook	$W_{2i_1,2i_1+3,0}$	$W_{2i_1,2i_1+3,1}$	$W_{2i_1+1,2i_1+3,0}$	$W_{2i_1+1,2i_1+3,1}$

$$W_{m,k,n} = \frac{1}{4} \begin{bmatrix} \mathbf{v}_m & \mathbf{v}_k \\ e^{j\pi n/2} \mathbf{v}_m & -e^{j\pi n/2} \mathbf{v}_k \end{bmatrix}, \quad \mathbf{v}_m = \begin{bmatrix} 1 & e^{j2\pi m/32} & e^{j4\pi m/32} & e^{j6\pi m/32} \end{bmatrix}^T.$$

The index $0 \leq i_1 \leq 15$ selects a cluster of beams from a set of partially overlapping clusters in a co-polarized subarray. The LSB of the index $0 \leq i_2 \leq 15$ selects one of the two co-phasing factors while the 3 MSBs of i_2 select one of 8 pairs of beams within a cluster.

This codebook structure can be viewed as a two-stage codebook, in which the precoding matrix \mathbf{W} corresponding to the UE's CSI feedback can be decomposed into a product of two matrices, such that $\mathbf{W} = \mathbf{W}_1 \cdot \mathbf{W}_2$, where \mathbf{W}_1 corresponds to the ULA grid-of-beams component of the PMI and \mathbf{W}_2 corresponds to the frequency-selective co-phasing component of the PMI. The codebooks for PMI feedback when less than 8 CSI-RS antenna ports are configured, and indeed the Release 8/9 codebooks, can be considered to have the same structure with \mathbf{W}_1 being set to the identity matrix.

29.3.3.2 Periodic and Aperiodic CSI Reporting Enhancements

Aperiodic CSI feedback on the Physical Uplink Shared Channel (PUSCH) as described in Section 10.2.1.1 follows the same design in Release 10 as in Release 8, except for the use of the new PMI codebooks described above if 8 CSI-RS ports are configured. However, the periodic feedback on the PUCCH (Physical Uplink Control CHannel) described in Section 10.2.1.2 is extended in Release 10 to take advantage of the new two-stage PMI feedback codebook structure explained above.

As spatial signatures within a co-polarized subarray are not expected to change much in the frequency domain, sub-band PMI feedback for this component (referred to as \mathbf{W}_1 above) was not considered to be necessary. On the other hand, sub-band reporting can be useful for the spatial component that captures co-phasing across polarizations (referred to as \mathbf{W}_2).

These, together with the aim of keeping the uplink overhead as low as possible, led to the following PUCCH feedback modes being supported for the case of 8 CSI-RS ports:

Wideband periodic feedback. The RI/PMI/CQI information is carried in two separate types of PUCCH feedback report: the first contains the RI, while the second contains associated PMI and CQI information. Two alternative modes of operation exist in this case, known as submode 1 and submode 2, which are selected by means of an RRC signalling parameter ‘PUCCH_format1-1_CSI_reporting_mode’. In submode 1, the wideband grid-of-beams component of the PMI is included in the first feedback report together with the RI. In submode 2, all of the PMI information (i.e. including both the grid-of-beams component and the sub-array co-phasing component) is contained in the second PUCCH report; in order to achieve this, the PMI feedback codebook has to be subsampled (i.e. only certain specified entries are available from which the UE may select) so that the second report can fit within the available number of bits on the PUCCH. Submode 2 minimizes the code rate for the RI, whereas submode 1 avoids the need for subsampling of the PMI codebook.

Sub-band periodic feedback. This mode allows the fact that the grid-of-beams component of the feedback may be relatively time-invariant [8] to be exploited, by reducing the feedback rate for that component and thereby making some feedback bits available for sub-band reporting of the co-phasing component. The RI/PMI/CQI information is carried in three separate types of PUCCH feedback report. A one-bit Precoder Type Indication (PTI) report is added to the RI reports so that the UE can indicate the contents of the second and third reports. If the PTI is set to ‘0’, the second report contains the grid-of-beams PMI component and the third report contains the wideband co-phasing component of the PMI together with the associated CQI; on the other hand, if the UE estimates that the grid-of-beams component has not changed since the previous time it was reported, the PTI is set to ‘1’, the second report contains wideband co-phasing PMI feedback together with associated CQI, and the third report contains sub-band co-phasing PMI and associated CQI; multiple such third reports are transmitted, one for each sub-band. An example of this sub-band periodic feedback using the PTI is illustrated in Figure 29.8, showing how the sub-band resolution of the co-phasing feedback is increased if the grid-of-beams state is static.




Figure 29.8: Example of a possible sequence of PUCCH CSI feedback reports when 8 CSI-RS ports and sub-band PMI feedback are configured.

29.4 Uplink Multiple Antenna Transmission

In Release 10, SU-MIMO transmission with up to four spatial layers is introduced to increase the data rate for the PUSCH, and the reliability of the control signalling on the PUCCH is increased using transmit diversity.

29.4.1 Uplink SU-MIMO for PUSCH

In order to support uplink SU-MIMO, the concept of transmission modes is introduced for the PUSCH in Release 10. The transmission modes are as follows:

PUSCH Transmission Mode 1: Transmission from a single antenna port;¹³

PUSCH Transmission Mode 2: Transmission from multiple antenna ports; within this mode, the UE can be configured to transmit from either 2 or 4 antenna ports.

If a UE is configured in PUSCH transmission mode 2, it can transmit up to two transport blocks per subframe.¹⁴ The transport blocks are mapped onto one or more layers according to the same rules as are used in the downlink in Release 8, as shown in Table 11.2.

Each transport block is independently acknowledged using the Physical HARQ Indicator Channel (PHICH). The PHICH index for the first codeword is the same as is used in Release 8 for single-codeword PUSCH transmission (i.e. associated with the lowest PRB index of the PUSCH resource allocation); the PHICH index for the second codeword is related to the first by an offset.

The PDCCH resource grant for uplink MIMO transmissions includes two independent New Data Indicators (NDIs) to indicate whether a retransmission is expected.

Closed-loop codebook-based precoding is used for the PUSCH in a very similar way to the Release 8 PUSCH in transmission mode 4 (see Section 11.2.2.2). The UE is instructed by the eNodeB as to which rank and precoding matrix to use, by means of a dynamic precoding matrix indicator transmitted in the uplink grant on the PDCCH.

In designing the codebooks for uplink SU-MIMO precoding, the primary considerations were preserving the single-carrier waveform at each of the transmit antennas and achieving as much precoding gain as possible. In order to comply with the former, there can be at most one non-zero entry in each row of the precoding matrix. For the rank-1 codebook, some vectors are introduced to turn off certain transmit antennas to save power.

The number of codewords in the codebook is a trade-off between performance and signalling overhead. As a result, for 2 transmit antennas, 6 precoding vectors are defined for rank-1, and a single identity precoding matrix is defined for rank-2; this can be signalled using 3 bits. For 4 transmit antennas, there are 24 precoding vectors for rank-1, 16 for rank-2, 12 for rank-3, and a single identity matrix for rank-4, requiring 6-bit signalling. Full details of the PUSCH precoding codebooks can be found in [1, Section 5.3.3A].

In view of the performance of closed-loop precoding for the PUSCH, transmit diversity is not considered useful. In Figure 29.9, the throughputs using Space-Time Block Coding (STBC) and long-term closed-loop precoding are compared at 3 km/h and 120 km/h under different transmit antenna correlations. Even with uncorrelated transmit antennas, STBC only

¹³Two configurations exist within this mode: one is the Release 8 PUSCH transmission scheme, while the other supports both contiguous and non-contiguous resource allocation (see Section 28.3.6.2).

¹⁴Per component carrier if carrier aggregation is employed – see Chapter 28.

shows better link throughput at high SINRs, corresponding to a Block Error Rate (BLER) below 10% (which is an unlikely operating point). Furthermore, as the transmit antennas become more correlated, long-term precoding exhibits more beamforming gain.




Figure 29.9: Link throughput performance of STBC vs long-term precoding [9].

If uplink control signalling is carried on the PUSCH¹⁵ in a subframe when two codewords are transmitted, any Acknowledgement/Negative Acknowledgement (ACK/NACK) and RI are replicated across all layers of both codewords before channel coding, with the coded modulation symbols for ACK/NACK and RI being time-domain multiplexed with the data and time-aligned across all layers. CQI and PMI signalling is mapped only to the codeword with the highest MCS as indicated by the initial uplink grant. When the two codewords have the same MCS, codeword 0 is always selected. The same multiplexing and channel interleaving mechanisms are used as in Release 8, with the control information being mapped to the same REs on the two spatial layers if the codeword is transmitted on two layers. In Figure 29.10, the placement of the Uplink Control Information (UCI) is illustrated for the case when the uplink PUSCH transmission is of rank 2 and codeword 0 has the higher MCS.

¹⁵See Section 16.4.




Figure 29.10: Multiplexing of control signalling with PUSCH data in the case of rank-2 SU-MIMO.

29.4.2 Uplink Transmit Diversity for PUCCH

The transmit diversity scheme introduced for the PUCCH in Release 10 is designed to ensure backward compatibility with the Release 8 PUCCH design (see Chapter 16). For PUCCH formats 1/1a/1b (see Section 16.3.2), Space Orthogonal-Resource Transmit Diversity (SORTD) [10] is used, whereby the UE transmits the same control information from different transmit antennas with different orthogonal resources, including cyclic shifts and OCCs. For PUCCH formats 1a/1b, the two orthogonal resources for SORTD transmission are derived from n_{CCE} and $n_{\text{CCE}} + 1$, where n_{CCE} is the number of the first Control Channel Element (CCE) used for the transmission of the corresponding uplink grant on the PDCCH. SORTD is also defined for PUCCH formats 2/2a/2b and 3. Transmit diversity is not supported in Release 10 for PUCCH format 1b with channel selection (see Section 28.3.2.1).

If the transmit antennas are uncorrelated, SORTD provides the best diversity performance, at the expense of consuming multiple orthogonal resources which potentially reduces the PUCCH multiplexing capability.

29.5 Coordinated MultiPoint (CoMP) Transmission and Reception

Coordinated MultiPoint (CoMP) transmission/reception, also known as Cooperative MIMO, has received significant attention in academic literature and is being studied as a technique to increase performance, especially at the cell edge, within the evolution of LTE-Advanced for Release 11 or beyond.

While the scope of CoMP includes both downlink and uplink cooperation, the downlink has received significantly more attention in the literature, primarily due to the more challenging nature of the transmission coordination problem. From the point of view of the air interface design, uplink CoMP basically consists of coordination of eNodeB scheduling and/or receiver processing, and hence the main standardization effort would lie in the definition of appropriate information exchange protocols between eNodeBs if multivendor operation is required. Indeed, uplink CoMP schemes can already be realized in Release 8 by proprietary mechanisms; even in the downlink, basic CoMP schemes can be realized in Release 8 between the cells controlled by a given eNodeB.

In addition to backhaul protocol support, downlink cooperation would require enhancements to the CSI feedback design. In the remainder of this section, the downlink is the primary focus.

29.5.1 Cooperative MIMO Schemes and Scenarios

Various forms of downlink cooperation are illustrated in Figure 29.11. Case (A) corresponds to the traditional uncoordinated single-cell transmission. Case (B) highlights a simple form of cooperation known as *coordinated scheduling* where Cell_B reduces the transmission power on a set of time-frequency resources for the benefit of UE_A and UE_C, served by Cell_A and Cell_C respectively. Note that Inter-Cell Interference Coordination (ICIC) via partial frequency reuse is a form of coordinated scheduling where the coordination is achieved by exchanging transmission power levels used on different time-frequency resources between eNodeBs via the X2 interface (see Section 12.5). Dynamic scheduling coordination is valuable in Home eNodeB (HeNB) Closed Subscriber Group (CSG) deployments where severe interference conditions can arise whenever UE moves into the coverage area of a different CSG, as explained in Chapter 24.

Case (C) illustrates an example of coordinated beamforming in the context of CSG deployments which adds a spatial dimension to ICIC and is applicable in deployments with multi-antenna eNodeBs. Unlike coordinated scheduling, coordinated beamforming allows an interfering HeNB₁ to transmit to UE₁ as long as transmit beams are steered away from UE₂ which is being served by HeNB₂ but located in the coverage area of HeNB₁. Note that opportunistic beamforming [11, 12] may be viewed as a simple form of coordinated beamforming whereby each cell performs beam-swapping according to predefined patterns while cells schedule UEs based on CSI feedback computed according to these patterns. This type of operation is well suited to the presence of a large number of UEs per cell and full buffer traffic. It is worth mentioning, however, that recent studies [13] have shown very limited spectral efficiency improvement due to coordinated beamforming forms of CoMP compared to single cell MU-MIMO in homogeneous macro-cellular deployments, even with full buffer traffic.




Figure 29.11: Downlink transmission schemes:

- (A) non-cooperative transmission; (B) coordinated scheduling;
(C) coordinated beamforming; (D) coherent joint transmission.

Coherent Joint Transmission (JT) (also known as multicell MU-MIMO or ‘network MIMO’) involves simultaneous transmission of data packets to one or more UEs from multiple cells with cophasing. This form of coordinated transmission requires fast connectivity (low latency) between the transmission points to enable the exchange of control information, as well as high bandwidth backhaul to the eNodeBs involved in data sharing. While general applicability of JT in macrocellular scenarios needs further evolution of the network interfaces, practical scenarios of interest include eNodeBs equipped with multiple Remote Radio Heads (RRHs), as illustrated in case (D) of Figure 29.11, or heterogeneous networks of macro-eNodeBs and picocells. Multiple RRHs or picocells can be placed according to coverage and/or capacity requirements and connected to one or more host eNodeBs via fast broadband links. In such configurations, the processing functionality and resource management may be centralized in a host eNodeB. Deployment of RRHs or picocells can be an attractive way to enhance throughput and coverage, especially in hot-zone and indoor scenarios where JT can play a significant role. From the air interface point of view, RRHs can be treated as cells of the same eNodeB. A key enabler of JT in heterogeneous network scenarios would be multicell channel state feedback from the UE.

While coherent JT across multiple eNodeBs is the ultimate approach towards maximizing spectral efficiency for low-mobility UEs, the practically achievable gains remain to be understood. Most of the studies available to date have been focused on intra-eNodeB JT as well as intra-cluster JT with network-based fixed clustering of a limited number of cells [14]. This kind of limited technique, however, shows limited spectral efficiency gains [15]. On

the other hand, theoretical analysis of JT transmission across a large number of eNodeBs suggests potential for substantial gains [16]. Besides the basic requirement of low latency and high throughput backhaul, major practical challenges yet to be addressed include dynamic and distributed multicell multi-user scheduling and transmit processing at the network side, and multicell channel quality prediction, and means for accurate measurement and delivery of the multicell channel state feedback from the UEs.

The greatest benefits in practice are likely to be seen in heterogeneous networks, where the proportion of UEs in interference-limited scenarios is relatively high.

29.6 Summary

As explained in this chapter, Release 10 incorporates some important enhancements to multi-antenna transmission in both uplink and downlink, aimed at substantially increasing peak data rates and system spectral efficiency. This is achieved by the introduction of 4×4 SU-MIMO for uplink data and transmit diversity for uplink control signalling, as well as in the downlink higher-order (8×8) SU-MIMO and enhancements for MU-MIMO. The downlink developments are supported by new reference signals for data demodulation and CSI estimation, and enhanced feedback for improved CSIT accuracy with low overhead.

Further evolution of the multi-antenna transmission schemes supported by LTE-Advanced is likely by means of CoMP, for which studies of potential performance improvements with practical signalling schemes are ongoing.

References¹⁶

- [1] 3GPP Technical Specification 36.211, 'Evolved Universal Terrestrial Radio Access (E-UTRA); Physical Channels and Modulation (Release 10)', www.3gpp.org.
- [2] Ericsson and ST-Ericsson, 'R1-093485: Downlink demodulation RS design for Rel-9 and beyond', www.3gpp.org, 3GPP TSG RAN WG1, meeting 58, Shenzhen, China, August 2009.
- [3] Qualcomm Europe, 'R1-094211: UE-RS Patterns for Rank 3-4,' www.3gpp.org, 3GPP TSG RAN WG1, meeting 58bis, Miyazaki, Japan, October 2009.
- [4] Qualcomm Europe, 'R1-104795: Details of OCC mapping for DL MIMO operation,' www.3gpp.org, 3GPP TSG RAN WG1, meeting 62bis, Xi'an, China, October 2010.
- [5] Huawei, LG Electronics, Samsung, Panasonic, Intel, HiSilicon, New Postcom, CATR, Potevio and CMCC, 'R1-105132: Proposal for specification of PDSCH Muting', www.3gpp.org, 3GPP TSG RAN WG1, meeting 62bis, Xi'an, China, October 2010.
- [6] Nokia Siemens Networks and Nokia, 'R1-094648: Inter-cell CSI-RS design and performance', www.3gpp.org, 3GPP TSG RAN WG1, meeting 59, Jeju, Korea, November 2009.
- [7] Alcatel-Lucent and Alcatel-Lucent Shanghai Bell, 'R1-094609: UL SRS enhancements to support CoMP and Transmit Diversity,' www.3gpp.org, 3GPP TSG RAN WG1, meeting 59, Jeju, Korea, November 2009.
- [8] Alcatel-Lucent Shanghai Bell and Alcatel-Lucent, 'R1-104088: Discussion of two-stage feedback proposals,' www.3gpp.org, 3GPP TSG RAN WG1, meeting 61bis, Dresden, Germany, June 2010.
- [9] Qualcomm Europe, 'R1-092716: Rank-1 Precoding vs OL Tx Diversity for PUSCH', www.3gpp.org, 3GPP TSG RAN WG1, meeting 57bis, Los Angeles, USA, July 2009.

¹⁶All web sites confirmed 1st March 2011.

- [10] LG Electronics, ‘R1-090786: PUCCH TxD Schemes for LTE-A’, www.3gpp.org, 3GPP TSG RAN WG1, meeting 56, Athens, Greece, February 2009.
- [11] P. Vishwanath, D. Tse, and R. Laroia, ‘Opportunistic Beamforming Using Dumb Antennas’, *IEEE Transaction on Information Theory*, Vol. 48, pp. 1277–1294, June 2002.
- [12] Huawei, ‘R1-093037: Performance of DL coordinated beam switching with bursty traffic’, www.3gpp.org, 3GPP TSG RAN WG1, meeting 58, Shenzhen, China, August 2009.
- [13] NTT DOCOMO, ‘R1-094953: TP for TR36.814 on Self-evaluation results’, www.3gpp.org, 3GPP TSG RAN WG1, meeting 59, Jeju, Korea, November 2009.
- [14] P. Marsch and G. Fettweis, ‘A Decentralized Optimization Approach to Backhaul-Constrained Distributed Antenna Systems’ in *Proceedings of the 16th IST Mobile & Wireless Communications Summit*, July 2007.
- [15] NTT DOCOMO, ‘R1-090314: Investigation on Coordinated Multipoint Transmission Schemes in LTE-Advanced Downlink’, www.3gpp.org, 3GPP TSG RAN WG1, meeting 55bis, Ljubljana, Slovenia, January 2009.
- [16] P. Marsch and G. Fettweis, ‘A Framework for Determining Realistic Capacity Bounds for Distributed Antenna Systems’, *Proceedings of the IEEE Information Theory Workshop*, October 2006.

Relaying

**Eric Hardouin, J. Nicholas Laneman,
Alexander Golitschek, Hidetoshi Suzuki, Osvaldo Gonsa**

30.1 Introduction

30.1.1 What is Relaying?

Relays are a key new feature of LTE-Advanced, introduced in Release 10 of the LTE specifications. These additional network nodes are designed to complement a macro-cellular network of regular eNodeBs with reduced cost, by expanding coverage or increasing capacity.

Early relays, in the form of repeaters, are already present in legacy radio interface technologies such as UMTS and Release 8 of LTE [1, 2]. Repeaters simply amplify the radio-frequency signal received from a macro base station. Compared to a base station, repeaters have a lower cost since they involve no baseband processing, backhaul network installation or subscription fee for access to the fixed public network. They are typically used to improve coverage in zones where the traffic is too light to justify the deployment of a base station, or where there is no easy backhaul network access, such as road segments in rural areas. Repeaters are also convenient for the provision of indoor coverage (e.g. in shopping malls). In the majority of cases, repeaters are deployed by network operators, but user-deployed low-power repeaters are also now available (e.g. to cover a shop or flat).

Although useful, repeaters exhibit two significant drawbacks:

- Since a repeater only amplifies the received signal, any received interference is also amplified. The signal quality in the zone covered by the repeater is thus degraded compared to a regular base station.
- Repeaters are operated independently of the radio access network. Separate Operation and Maintenance (O&M) functionality is therefore required for the repeater to signal

LTE – The UMTS Long Term Evolution: From Theory to Practice, Second Edition.

Stefania Sesia, Issam Toufik and Matthew Baker.

© 2011 John Wiley & Sons, Ltd. Published 2011 by John Wiley & Sons, Ltd.

its status and any potential malfunction to the network, which results in additional Operational EXpenditure (OPEX) for the network operator.

Relays can be seen as an evolution of repeaters to solve the above drawbacks. A Relay Node (RN) is a network node connected wirelessly to a source eNodeB, called the *donor eNodeB*. An important characteristic of RNs is that they are under the full control of the radio access network, which permits similar monitoring and remote control capabilities as for an eNodeB. In contrast to a repeater, an RN processes the received signal before forwarding it; this may involve Layer 1, Layer 2 or Layer 3 operations, such that an RN can, in principle, range from an enhanced repeater to a fully fledged eNodeB with a wireless backhaul connection, as discussed further in Section 30.3.

In contrast to the minimal delay introduced by the simple amplification process of a repeater, the processing performed by an RN requires at least two transmission occasions to deliver the signal from the donor eNodeB to a UE, as illustrated in Figure 30.1.




Figure 30.1: Reception and transmission phases of a repeater and a relay node.

In 3GPP, the following terminology is used, some of which is illustrated in Figure 30.2:

- **Donor eNodeB/cell.** The source eNodeB/cell from which the RN receives its signal.
- **Relay cell.** The coverage area of the RN.
- **Backhaul link.** The link between the donor eNodeB and the RN.
- **Access link.** The link between the RN and a UE.
- **Direct link.** The link between the donor eNodeB and a UE.
- **Inband/outband.** An inband RN uses the same carrier frequency for the backhaul link as for the access link; otherwise, the RN is said to be outband.
- **Half/full duplex.** A half-duplex RN cannot receive on the backhaul link at the same time as transmitting on the access link, and vice versa, whereas a full-duplex RN has sufficient antenna isolation to be able to operate without this restriction. This distinction applies to inband RNs only, since outband RNs are always full-duplex.
- **Donor and coverage antennas.** At the RN, the donor antenna(s) are used for the backhaul link, while the coverage antenna(s) are used for the access link. In some cases, the physical donor and coverage antennas may be the same.




Figure 30.2: Terminology for relaying.

30.1.2 Characteristics of Relay Nodes

The characteristics of RNs have a significant impact on the network operation, giving rise to a trade-off between the cost and efficiency of RNs.

Inband RNs are expected to be the most common, in which case the backhaul link consumes radio resources from the donor cell, thus reducing the capacity of the latter. If the RN is also half-duplex, there are further complications for the system design, as discussed in Section 30.3. Full-duplex RNs provide better relay cell capacity and allow simpler system design than half-duplex RNs. On the other hand, full-duplex RNs need separate antennas for the access and backhaul links, with sufficient isolation between them, and are therefore expected to be more costly than half-duplex RNs. Outband relaying increases the relay cell capacity and simplifies the system design by allowing the backhaul link to operate on a separate carrier frequency, but this is expensive for the network operator, if indeed a second carrier frequency is available.

The transmit power and the type and number of antennas directly impact the complexity, cost, size and weight of an RN, which in turn affect the site costs (engineering and site rental).

The RN transmission power depends on the deployment scenario (see Section 30.1.4) and can range from 30 dBm or less to levels similar to a macro eNodeB (e.g. 46 dBm).

If the RN uses a single physical antenna for both the backhaul and the access links, an omni-directional antenna is typical. However, the backhaul link quality can be enhanced by the use of a directional antenna towards the donor eNodeB, which also reduces interference to other cells when the RN transmits in the uplink. In this case, a separate coverage antenna is needed, which may also be directional, away from the donor eNodeB. Multiple physical antennas at the RN are needed to support Multiple-Input Multiple-Output (MIMO) techniques, which can be especially beneficial for enhancing the backhaul link capacity – expected to be the bottleneck of relaying.

The required processing capabilities of an RN are also affected by the number of UEs the RN is designed to support, which may depend on whether it is deployed indoors or outdoors.

The main expected benefit of installing an RN in an LTE network compared to an additional eNodeB is a reduction of the infrastructure and operational costs, as well as allowing easier and quicker installation.

Unlike eNodeBs, RNs do not require any wired or microwave backhaul connection to the network. This avoids a significant part of the network CAPital EXpenditure (CAPEX) and OPEX, since wired backhaul is typically costly for both the initial connection to the site and the monthly subscription to a network provider, while microwave backhaul suffers from high initial equipment costs as well as licence fees for the microwave carrier frequency. For an RN, the backhaul-related costs are much reduced, since the same radio interface technology is used for the backhaul and access links, and, if the RN operates inband, no additional spectrum licence is needed. An RN also allows more flexibility in site selection than a wired node, by alleviating the constraint on fixed network availability. The main requirements for an RN site are a power supply and good radio link quality towards the donor eNodeB. This can also enable RNs to be deployed very quickly, so that RNs can provide a convenient solution for temporary or emergency deployments (see Section 30.1.4).

Nevertheless, RNs also have some drawbacks compared to eNodeBs. In addition to the increased delay imposed on the transmissions compared to normal eNodeB operation, the capacity offered by an RN will be reduced compared to an additional eNodeB, due to the radio resources consumed by the wireless backhaul. In terms of coverage extension, however, simulations show that the performance difference between outdoor inband RNs and pico eNodeBs (or outband RNs) is small (e.g. [3]).

30.1.3 Protocol Functionality of Relay Nodes

RNs can be categorized by the protocol layer functionality provided [4, 5]. Simplified diagrams of the protocol stacks for Layer 1, 2 and 3 RNs are shown in Figure 30.3.

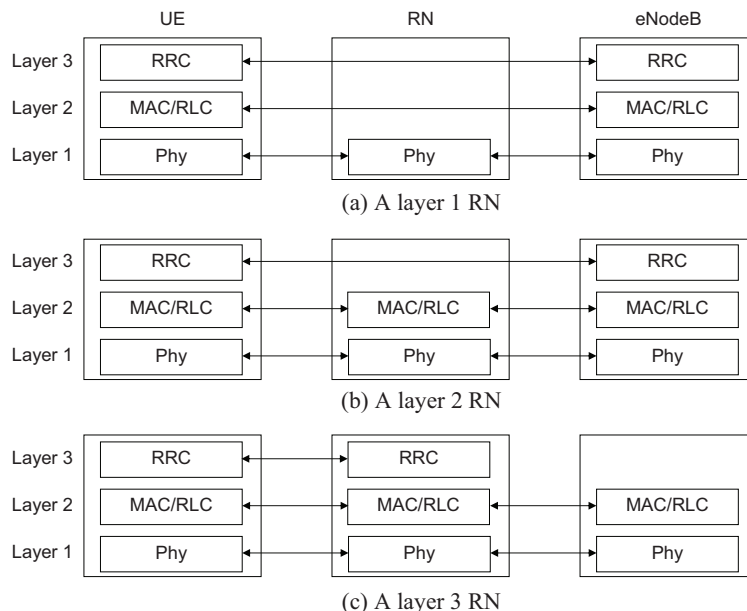


Figure 30.3: Protocol stack alternatives for an RN.

30.1.3.1 Layer 1 RN

In a Layer 1 RN, no Layer 2 processing is carried out. If the RN only performs Radio Frequency (RF) processing, it is a repeater (sometimes known as a ‘Layer 0’ RN) [2], as described in Section 30.1.1. Layer 1 RNs that are not simple repeaters can include baseband processing such as Forward Error Correction (FEC), but do not include scheduling functions.

30.1.3.2 Layer 2 RN

In a Layer 2 RN, Medium Access Control (MAC) functions such as scheduling (see Section 4.4.2.1) are carried out within the RN. In order to be identified by a UE as a network transmission/reception point, a Physical Cell-ID (PCI) signalled by the Primary and Secondary Synchronization Signals (PSS/SSS, see Section 7.2) is required. The presence of a PCI requires Layer 3 identification by the UE. On the other hand, some Layer 3 functions, such as Radio Resource Control (RRC, see Section 3.2), can be located within the eNodeB rather than the RN. Thus although the cell itself is realized within the RN, only Layer 2 functions are implemented there. Alternatively, a Layer 2 RN could be realized without the UE identifying it as a network transmission/reception point (no PCI is required in this case); in such a case, the relay might work with the eNodeB in a cooperative way, for example only for retransmission of packets. A Layer 2 RN may also support Radio Link Control (RLC) Automatic Retransmission reQuest (ARQ) and segmentation/concatenation functions (see Section 4.3.1).

30.1.3.3 Layer 3 RN

A Layer 3 RN has its own PCI signalled by the PSS/SSS. Mobility between RNs and eNodeBs is based on RRC in the Layer 3 control plane (see Chapter 3). In the user plane, a Layer 3 RN supports protocols at least up to PDCP (i.e. IP packet handling) (see Chapter 4). All Layer 1 and Layer 2 functions are supported within such an RN.

30.1.4 Relevant Deployment Scenarios

RNs can be used either to extend the cell coverage or to provide additional traffic capacity (hot spots). Coverage is defined by the geographical area of reception of the control channels – i.e. the *cell footprint*. Increased traffic capacity is provided by higher Signal-to-Interference-plus-Noise Ratios (SINRs), and possibly by cell splitting within the existing cell footprint.

The main practical deployment scenarios are illustrated in Figure 30.4 and discussed below.

30.1.4.1 Cell Coverage Extension

In rural scenarios, RNs can be used to provide cost-effective coverage of a large zone to extend the cell footprint provided by a macro eNodeB, as in Figure 30.4(a). Such an RN is typically mounted on a high mast, with a transmit power similar to that of a macro eNodeB, i.e. 46 dBm. RNs can also be useful for deployments in a higher frequency band than a 3G network, enabling the same coverage to be retained despite the worse radio propagation conditions, without increasing the number of base station sites.




Figure 30.4: Typical use cases for relays: (a) cell coverage extension; (b) capacity boost; (c) indoor coverage enhancement; (d) dead spot mitigation.

30.1.4.2 Urban or Suburban Outdoor Capacity Boost

In this scenario, illustrated in Figure 30.4(b), the RNs are used to provide capacity boost in particular areas of the cell (hot spots), or at the cell edge to homogenize the throughput offered throughout the cell. Lamp posts often provide convenient outdoor sites for RNs in such environments. Transmit powers of 30 dBm or less are typically appropriate for urban environments due to the small cell size and inter-cell interference issues. Suburban deployments might use up to 37 dBm.

30.1.4.3 Indoor Coverage Enhancement

RNs can be placed inside a building to provide enhanced indoor coverage (Figure 30.4(c)); in such a case, they can be seen as femto eNodeBs (or Home eNodeBs – see Chapter 24) with a wireless backhaul. This can be useful in cases where the user has no ADSL link, or in countries where there is no wired network infrastructure to provide the backhaul access. Indoor RNs must have low transmission power to avoid interference problems.

30.1.4.4 Dead Spot Mitigation

RNs may be used for the mitigation of dead spots – i.e. filling coverage holes in the macro network caused by large obstacles such as high buildings, as in Figure 30.4(d). In such a scenario, the RN is typically placed in line of sight of the donor eNodeB and radiates toward the dead spot. If the extra capacity required by the dead spot compared to the donor cell is limited, an RN represents a more cost-effective solution than a conventional eNodeB.

30.1.4.5 Temporary Deployments

RNs provide a convenient solution for temporary network deployments such as special events (e.g. sports games, street fairs) or emergency deployments (e.g. if network equipment has

been damaged by a natural disaster). Suitable RNs would be low power in order to be able to operate without a special eNodeB-type power supply.

30.1.4.6 Group Mobility

An RN can provide coverage to passengers in a moving vehicle such as a bus, train or ship. The RN is placed on the vehicle and connects to the most appropriate donor eNodeB as the vehicle moves. Unlike the other scenarios, such an RN is required to support mobility procedures (handover). In-vehicle RNs offer better coverage than macro eNodeBs and save the UEs' batteries since the UEs' uplink transmission power is reduced. A vehicle-mounted RN can have a higher transmit power and/or more efficient antennas than a UE and can therefore reduce the incidence of call drops during handovers.

30.2 Theoretical Analysis of Relaying

The capacity of RN-assisted cellular network deployments can be considered in the framework of the point-to-point *relay channel* [6], the *multi-access relay channel* [7], and the *broadcast relay channel* [8]. Fundamental performance limits for the channel capacity are generally unknown for these cases, but a number of communication strategies and upper/lower bounds on the fundamental limits have been obtained. For LTE-Advanced, the most relevant model for the cell coverage extension scenario of Figure 30.4(a) is a collection of relay channels, while the most relevant model for the capacity boost scenario of Figure 30.4(b) consists of multiple-access and broadcast relay channels for the uplink and downlink respectively.

This section summarizes some of the relevant theoretical developments and fundamental limits in the context of LTE-Advanced and illustrates how practical system constraints reduce the complexity of the models and of the architectural designs. A more detailed introduction to relaying and cooperative communications can be found in [9].

30.2.1 Relaying Strategies and Benefits

We begin with a summary of the general Gaussian relay channel [6, 10], on which much of the following analysis is derived.

A general Gaussian relay channel relevant to both the coverage extension and capacity boost cases of Figure 30.4 can be modelled as follows (also illustrated for the case of the downlink in Figure 30.5):

$$Y_2(t) = H_{2,1}X_1(t) + Z_2(t)$$

$$Y_3(t) = H_{3,1}X_1(t) + H_{3,2}X_2(t) + Z_3(t)$$

where $X_1(t)$ is the source transmission, $Y_3(t)$ is the destination received signal, and $X_2(t)$ and $Y_2(t)$ are the RN transmitted and received signal, respectively. The additive noises $Z_j(t)$ are modelled as white Gaussian, with one-sided power spectral densities N_j , $j = 2, 3$. The channel coefficients $H_{j,i}$, $i = 1, 2$, $j = 2, 3$ capture the effects of path-loss and shadowing; for simplicity, we assume here that these coefficients vary slowly relative to the signalling interval, are non-selective with frequency and are known to the source, the RN and the

destination. The model can readily be extended to incorporate other channel effects, such as multiple antennas, resolvable multipath and limited knowledge of the channel coefficients at the transmitters and receivers.

For the cell coverage extension scenario of Figure 30.4(a), the destination is outside the coverage range of the control channels, and therefore $H_{3,1}$ is set to zero.




Figure 30.5: The general Gaussian relay channel model in the context of downlink transmission in an RN-assisted cellular network. The same model can be applied to the uplink case, with the UE and the donor eNodeB reversed.

The source and RN are assumed to have average power constraints P_1 and P_2 respectively. The waveforms are essentially band-limited to the same frequency band of width W .

The received Signal-to-Noise Ratio (SNR) between transmitter $i = 1, 2$ and receiver $j = 2, 3$ can then be written as:

$$\rho_{ji} = \frac{|H_{ji}|^2 P_i}{N_j W}. \quad (30.1)$$

We consider here distance-dependent path-loss models such that the average channel gain $E[|H_{ji}|^2] \propto d_{ji}^{-\nu}$, where d_{ji} is the distance between transmitter i and receiver j , and $2 \leq \nu \leq 5$ is the path-loss exponent [11, 12].

For the capacity boost scenario of Figure 30.4(b), we first analyse the potential capacity increase compared to the case when the RN is not present. If $X_2(t) = 0$, then the channel model reduces to a point-to-point Gaussian channel with capacity given by [10, 13]

$$C_{\text{direct}} = W \log_2 (1 + \rho_{3,1}). \quad (30.2)$$

By processing the received signal $Y_2(t)$, the RN can create its transmission $X_2(t)$ so as to enhance performance beyond Equation (30.2) which therefore represents a lower bound on the capacity of an RN-aided channel.

An upper bound on the capacity of the general Gaussian relay channel can be obtained by considering the total information flows across a set of ‘cuts’ in the RN-assisted network, as shown in Figure 30.6. It is therefore known as the *cut-set bound* [6, 10].

The cut-set upper bound is given by the minimum of the rate of the broadcast cut from the source to the RN and destination, and the rate of the multiple-access cut from the source and RN to the destination, maximized over the parameter α ($0 \leq \alpha \leq 1$) which controls the amount of correlation between the signals $X_1(t)$ and $X_2(t)$: $C_{\text{relay}} \leq C_{\text{cutset}}$, where

$$C_{\text{cutset}} = W \max_{0 \leq \alpha \leq 1} \min \left\{ \log_2 (1 + (1 - \alpha)(\rho_{2,1} + \rho_{3,1})), \right. \\ \left. \log_2 (1 + \rho_{3,1} + \rho_{3,2} + \sqrt{\alpha \rho_{3,1} \rho_{3,2}}) \right\} \quad (30.3)$$




Figure 30.6: Broadcast and multiple-access cuts through a relay network.

As introduced in Section 30.1.1, a decode-and-forward RN decodes the source transmission $X_1(t)$ and re-encodes it to produce $X_2(t)$. This approach in theory allows for the design of a joint signalling scheme among $X_1(t)$ and $X_2(t)$, roughly approximating a system with two transmit antennas and one receive antenna. The resulting lower bound on channel capacity for this approach is $C_{\text{relay}} \geq C_{\text{decode}}$ [6, 10], where

$$C_{\text{decode}} = W \max_{0 \leq \alpha \leq 1} \min \left\{ \log_2 (1 + (1 - \alpha)\rho_{2,1}), \log_2 (1 + \rho_{3,1} + \rho_{3,2} + \sqrt{\alpha\rho_{3,1}\rho_{3,2}}) \right\}. \tag{30.4}$$

The first term in the minimum in Equation (30.4) corresponds to the capacity of the link between the source and RN assuming that the fraction $(1 - \alpha)$ of the source power P_1 is used to convey information to the RN. The second term in the minimum in Equation (30.4) corresponds to the multiple-access cut from the source and RN to the destination, and is identical to the second term in the minimum in Equation (30.3).

Examination of Equations (30.2)–(30.4) provides some theoretical insights about the potential benefits of relaying. To see these benefits, assume that the RN lies between the source and destination, so that we can expect both $\rho_{2,1} \geq \rho_{3,1}$ and $\rho_{3,2} \geq \rho_{3,1}$. If $\alpha = 0$ (i.e. no correlation between $X_1(t)$ and $X_2(t)$), then both terms in the minima of Equations (30.3) and (30.4) are larger than Equation (30.2). These observations, illustrated in Figure 30.7, suggest that relaying increases the capacity for a given amount of source power and bandwidth relative to direct transmission, extends coverage for the same target data rate, or some combination of these two benefits.

Additional observations result from carefully examining the structure of the gains that relaying can provide. These gains can be classified into the following three basic types:

- **Multihop Gain** results from the fact that, if the RN lies physically between the source and destination, then transmissions from the source to the RN and from the RN to the destination occur over shorter distances than direct transmission from the source to the destination. If we assume $\rho_{2,1} \gg \rho_{3,1}$ and $\rho_{3,2} \gg \rho_{3,1}$ in Equations (30.3) and (30.4), as is the case in the coverage extension scenario, then in the best case of a full-duplex relay

$$C_{\text{relay}} \approx W \min \{ \log_2 (1 + \rho_{2,1}), \log_2 (1 + \rho_{3,2}) \} \gg C_{\text{direct}} .$$

It needs to be borne in mind that this gain comes at the expense of power and bandwidth usage in the cell of the donor eNodeB. From a system perspective, there may be cases




Figure 30.7: Example theoretical performance of various relaying schemes with the RN located halfway between the source and destination, path-loss exponent $\nu = 3$, and no correlation between the backhaul and access links:

(a) spectral efficiency C/W versus SNR between the source and RN,

(b) fractional bandwidth β allocated to RN reception for a half-duplex RN strategy.

(As explained in Section 30.2.2, the parameter β is the fraction of the channel dimensions during which the RN can receive when it is operating in half-duplex mode; the RN can transmit in the remaining $1 - \beta$ opportunities.)

in which total capacity is higher if these resources are instead devoted to UEs closer to the donor eNodeB. However, in general, the use of RNs will have the effect of increasing fairness, as the capacity to distant terminals is increased at the expense of capacity to terminals close to the donor eNodeB.

- **Power / Diversity Gain** results from the summing and averaging of SNRs in the various expressions. For example, the second term in the minimum in Equation (30.4) has effective SNR at least

$$\rho_{3,1} + \rho_{3,2} = \left(\frac{|H_{3,1}|^2 P_1 + |H_{3,2}|^2 P_2}{P_1 + P_2} \right) \cdot \left(\frac{P_1 + P_2}{N_3 W} \right). \quad (30.5)$$

The first term in Equation (30.5) corresponds to a weighted average of the channel gains, $|H_{3,1}|^2$ and $|H_{3,2}|^2$, that is reminiscent of transmit diversity [14, 15]. Relaying protocols that provide diversity benefits are analysed in [16, 17] and references therein. The second term in Equation (30.5) highlights the additional power that the relay can provide to the system.

- **Coherence Gain** results from the ability to create correlated signals $X_1(t)$ and $X_2(t)$ with the parameter α , leading to the additional SNR of $\sqrt{\alpha \rho_{3,1} \rho_{3,2}}$ in the second term in the minima in Equations (30.3) and (30.4). Note that coherence gain is only possible with accurate carrier and symbol timing synchronization, which is implicitly assumed here. Such tight synchronization may in theory be possible on the downlink but may

not be realistic on the uplink. It is possible to show that, with phase variations known to the receivers but not the transmitters, the optimal choice is $\alpha = 0$, i.e. no correlation between $X_1(t)$ and $X_2(t)$ [8].

30.2.2 Duplex Constraints and Resource Allocation

Due to circuit isolation and radio cost issues, it is often necessary to preclude the relay from transmitting and receiving simultaneously in the same frequency band. This *half-duplex* constraint is often modelled by allowing the relay to receive in a fraction $0 \leq \beta \leq 1$ of the channel dimensions (time and frequency) and to transmit in the remaining fraction $\bar{\beta} = 1 - \beta$; the source utilizes a fraction $0 \leq \gamma \leq 1$ of its power while the relay receives and the remaining fraction $\bar{\gamma} = 1 - \gamma$ while the relay transmits. Time-sharing can then be performed between two general Gaussian relay channels of the form discussed in Section 30.2.1, one in which the relay only receives ($X_2(t) = 0$), and the other in which the relay only transmits ($Y_2(t) = 0$). As an example, under this half-duplex constraint, Equation (30.4) becomes:

$$C_{\text{decode,hd}} = W \max_{0 \leq \beta, \gamma \leq 1} \min \left\{ \beta \log_2 (1 + \gamma \rho_{2,1} / \beta), \right. \\ \left. \beta \log_2 (1 + \gamma \rho_{3,1} / \beta) \right. \\ \left. + \bar{\beta} \log_2 \left(1 + (\bar{\gamma} \rho_{3,1} + \rho_{3,2} + \sqrt{\bar{\gamma} \rho_{3,1} \rho_{3,2}}) / \bar{\beta} \right) \right\}. \quad (30.6)$$

The rate expression in Equation (30.6) simplifies considerably for the special case where $\beta = 1/2$ and $\gamma = 1$. In this case, the half-duplex constraint affects the multihop gain in the sense that, if $\rho_{2,1} \gg \rho_{3,1}$ and $\rho_{3,2} \gg \rho_{3,1}$, then

$$C_{\text{decode,hd}} \approx W/2 \min \{ \log_2(1 + 2\rho_{2,1}), \log_2(1 + 2\rho_{3,2}) \}.$$

Here the SNR gain of multihop transmission can be eliminated by the bandwidth cost of half-duplex, so that $C_{\text{decode,hd}} < C_{\text{direct}}$, especially for high rates [18].

Because modern cellular systems such as LTE use basically orthogonal transmission schemes, the key challenges within a cell relate to resource allocation in time and frequency. This view is oversimplified relative to general information-theoretic treatments of multiple access and broadcast, but is generally adopted in practice to avoid excessive complexity. The result of transmission resources being orthogonal is that the general multiple access channel and broadcast channel can be converted into a collection of point-to-point channels with limited interference over which power and transmission resources can be allocated. Similarly, orthogonal transmissions convert the general multiple-access relay channel and broadcast relay channel into collections of simpler channels. For example, orthogonal transmissions on the uplink in the capacity-boost scenario of Figure 30.4(b) can result in a subset of point-to-point channels from the UEs to the eNodeB and subsets of channels between a set of UEs and each relay whose rates are aggregated into a backhaul channel between the relay and eNodeB. For the decode strategy, a relay must decode *all* signals of the UEs associated with it, and convey the sum rate over the backhaul channel. Careful resource allocation must be performed by the cell scheduler in order to ensure that the half-duplex constraint and the backhaul throughput do not become system bottlenecks.

Finally, it should be noted that, although not currently considered in the context of LTE-Advanced, network coding of data streams at the relays can lead to increased system capacity and more efficient use of resources [19–21].

30.3 Relay Nodes in LTE-Advanced

As outlined in Section 30.1, the term ‘Relay Node’ can encompass a wide variety of different implementations and operation modes. The RN functionality defined in Release 10 is the result of careful consideration of these different possibilities.

30.3.1 Types of RN

An important aspect of RN design is the extent to which the RN can be seen by a UE. RNs in Release 10 are designed to be backward compatible to legacy UEs (i.e. UEs conforming to Releases 8 and 9 of the LTE specifications). During the study for LTE-Advanced, 3GPP identified two major types of RNs – Type 1 and Type 2 – and some subtypes of Type 1 RNs.

30.3.1.1 Type 1, 1a and 1b RNs

These types of RN are Layer 3 RNs. A UE sees them as separate cells (i.e. normal eNodeBs) with their own PCIs. The RN transmits all control and data channels including the PSS/SSS. UEs receive scheduling information such as the Physical Downlink Control Channel (PDCCH) from the RN, and send feedback information such as Channel State Information (CSI) and ACKnowledgement/Negative ACKnowledgements (ACK/NACKs) to the RN. The handover and cell reselection procedures are identical to Release 8 (see Chapter 3), and the scheduler is located in the RN in order to respond quickly to a UE’s feedback.

Type 1 RNs are inband half-duplex RNs, while Type 1a RNs operate outband.

Type 1b RNs have sufficient isolation between the received and transmitted signals to enable full-duplex operation – i.e. the backhaul and access links can be active simultaneously without the need for time-division multiplexing. In one example of a Type 1b RN, the backhaul antenna may be located outside the building while the coverage antenna is located inside the building for the improved support of UEs inside the building. Antenna isolation may also be aided by mechanical or adaptive beam-forming.

LTE-Advanced supports Type 1 and Type 1a RNs in Release 10 of the specifications. No specific support is provided for Type 1b RNs, although such RNs may be able to be deployed by implementation-dependent means.

30.3.1.2 Type 2 RNs

Type 2 RNs, which were also studied for LTE-Advanced but not adopted, are Layer 2 RNs. They cannot be identified by a UE since these RNs would not transmit control channels or have their own physical cell-IDs. These RNs would only transmit the Physical Downlink Shared Channel (PDSCH), and the scheduler is located in the eNodeB.

Type 2 RNs would be operated in a non-cooperative way with the donor eNodeB [22] or in a cooperative way. In the latter case, the eNodeB and RN would jointly transmit and receive

the signal to and from a UE, as illustrated in Figure 30.8. The initial transmission from the eNodeB would be received at both the RN and the UE at the same time. As the backhaul link should usually experience good radio conditions, the RN would be expected to receive the data correctly more often than the UE. For the retransmissions, both the RN and the eNodeB would transmit to the UE, and the UE would combine the two signals. The coverage area of a Type 2 RN would always overlap the coverage area of the donor eNodeB, as the UE would rely on control channel reception from the donor eNodeB.




Figure 30.8: Cooperative operation of Type 2 RNs.

30.3.2 Backhaul and Access Resource Sharing

It can be seen from Figure 30.2 that an RN competes for radio resources with UEs in the donor eNodeB's cell coverage areas. This can, in principle, be solved either by dynamic resource allocation, where the RN is scheduled like any other UE in the donor eNodeB's cells, or by reserving a specific band or part of a band for the backhaul. On the backhaul link, an RN has to receive and transmit signals from and to the donor eNodeB, while on the access link, an RN has to transmit and receive signals to and from the UEs. Some coordination or separation is necessary in order to avoid the transmitted signal of the RN causing interference to its own receiver. This separation can be in frequency, in time or, in the case of Type 1b relays as discussed above, in space. Alternatively a transmission signal canceller can be implemented in the receiver [23] to limit the self-interference, although this increases the RN complexity.

30.3.2.1 Separation in Frequency

Outband RNs (i.e. Type 1a RNs) use different frequencies for the backhaul and access links. Nevertheless, if the frequencies for the backhaul and access links are not sufficiently separated, interference can still occur due to out-of-band and spurious emissions.

Alternatively, frequency separation in the same band could in theory be achieved if the eNodeB scheduler were to allocate dynamically the Resource Blocks (RBs) for the backhaul and access links, provided that the RN's receiver could be sufficiently isolated against sidelobes from the transmitted signal. This would allow dynamic characteristics of the data traffic to be taken into account. However, the scheduler design for such an arrangement may be complex, and it would require additional signalling from the eNodeB to the RN to indicate the access link resource assignments. Therefore, the simpler realization is to assign the resources for the backhaul and access links relatively statically.

30.3.2.2 Separation in Time

Separation in time is realized by time-division multiplexing between the backhaul and access links within the same spectrum. The schedulers for the backhaul and access links are each responsible for not scheduling transmissions when the respective link is not available.

If arbitrary time instances were used for backhaul downlink reception at the RN, with no access link transmissions from the RN at these times, a UE could lose its connection to the RN due of the lack of Reference Signals (RSs), synchronization signals and control channels such as the Physical Control Format Indicator CHannel (PCFICH) and Physical Broadcast CHannel (PBCH). Therefore, to meet the requirement for backward compatibility, the time instances at which no transmission will occur from the RN have to be designed in such a way that even Release 8 UEs are not disturbed by them.

The Release 8 LTE specifications already provided a mechanism to inform UEs by RRC signalling [24] that certain subframes are designated as ‘Multimedia Broadcast Single Frequency Network’ (MBSFN) subframes. In such subframes, the UE expects control signals and RSs only in the first 1–2 OFDM symbols and ignores the rest of the subframe.¹ Originally this mechanism was introduced purely for MBSFN transmissions (see Section 13.4), but the signalling can be reused in Release 10 for other purposes, such as relaying, while avoiding disturbance to Release 8 and 9 UEs. Therefore these subframes in which RN backhaul downlink reception occurs are still referred to as ‘MBSFN’ subframes, even though no actual MBSFN transmission takes place there.

MBSFN subframes are therefore used for time-division multiplexing between the backhaul and access links for Type 1 RNs for LTE-Advanced.

As shown in Figure 30.9, the number of available OFDM symbols for an RN to receive the backhaul signal in an MBSFN subframe is less than the full subframe length because the RN has to transmit control signalling on the access link in the first 1–2 OFDM symbols for backward compatibility; additionally, some time is required for the RN to switch from transmission to reception and vice versa.




Figure 30.9: Transmission from an RN to a UE in a normal subframe, and from the donor eNodeB to the RN in an MBSFN subframe.

The starting symbol of the RN’s backhaul downlink reception window is configured by RRC signalling as the start of the second, third or fourth OFDM symbol, and can therefore be changed during the operation of an RN by RRC reconfiguration. The last symbol of the reception window is a function of the eNodeB–RN synchronization, the backhaul signal

¹ UEs of Release 9 and later may be instructed to receive the Physical Multicast Channel (PMCH) for MBMS in MBSFN subframes (see Section 13.4.1).

propagation delay and the RN's transmission–reception switching time, all of which can be considered to be matters of RN implementation and deployment. Therefore the last symbol of the RN's backhaul downlink reception window does not need to be configurable by RRC, and the reception window continues up to the last or second-to-last OFDM symbol. Consequently, the resulting number of OFDM symbols within an RN's backhaul downlink reception window ranges from 10 to 13 (in the case of the normal cyclic prefix being configured). These reception restrictions need to be taken into account by the donor eNodeB when transmitting on the backhaul link, by suitable RRC configuration, resource assignment and link adaptation.

30.3.3 Relay Architecture

RNs in LTE-Advanced support the full eNodeB functionality, including termination of the radio protocols and the S1 and X2 interfaces (see Sections 2.5 and 2.6 respectively). The network interfaces relevant to RNs are shown in Figure 30.10. The backhaul interface is known as the Un interface, over which the RN is connected to the donor eNodeB via S1 and X2 interfaces.




Figure 30.10: Network architecture and interfaces for RNs.
Reproduced by permission of © 3GPP.

The donor eNodeB provides S1 and X2 proxy functionality between the RN and other eNodeBs, Mobility Management Entities (MMEs) and Serving Gateways (S-GWs) (see Chapter 2); this includes passing UE-specific S1 and X2 signalling messages, as well as GPRS² Tunneling Protocol (GTP) data packets, between the S1/X2 interfaces associated

²General Packet Radio Service.

with the RN and the S1/X2 interfaces associated with other network nodes. Due to the proxy-like functionality, the donor eNodeB appears to the RN as an MME (for S1), as an eNodeB (for X2) and as an S-GW.

30.3.3.1 Relay Node User Plane

The donor eNodeB provides the S-GW and Packet Data Network (PDN) Gateway (P-GW) functionality for the RN (see Figure 30.11). This includes creating a session for the RN and managing Evolved Packet System (EPS) bearers (see Section 2.4) for the RN. The RN (in the uplink) and donor eNodeB (in the downlink) also perform mapping of signalling and data packets onto EPS bearers that are set up for the RN. The mapping is based on existing Quality of Service (QoS) mechanisms defined for the UE and the P-GW such as Quality Class Identifier (QCI) values (see Section 2.4). On the backhaul link, there can be up to eight Data Radio Bearers (DRBs) per RN. Data on these RN bearers is mapped from the EPS bearers of the UEs connected to the RN according to the QCI values of the EPS bearers. The mapping is configured by O&M and supports many-to-one mapping. The set up, timing and modification of Un bearers is left to the donor eNodeB implementation.




Figure 30.11: RN user plane protocol stack. Reproduced by permission of © 3GPP.

30.3.3.2 Relay Node Control Plane

There is one S1 interface relation between an RN and the donor eNodeB, and one S1 interface between the donor eNodeB and each MME in the ‘MME pool’ (see Section 2.2.2). The donor eNodeB processes and forwards all S1 messages between the RN and the MMEs for all UE-dedicated procedures (see Figure 30.12). The processing of S1 Application Protocol (S1-AP) messages includes modifying S1-AP UE identities and GTP Tunnelling End IDs (TEIDs – see Section 2.5) but leaves other parts of the messages unchanged. All non-UE-dedicated procedures are handled locally, between the RN and donor eNodeB, and between the donor eNodeB and MMEs.

There is one X2 interface relation between an RN and the donor eNodeB, and one X2 interface relation between the donor eNodeB and each other eNodeB with which the donor eNodeB has an X2 interface relationship. X2 messages for mobility are processed in a similar way to the S1 messages.




Figure 30.12: RN control plane protocol stack. Reproduced by permission of © 3GPP.

30.3.4 RN Initialization and Configuration

RN initialization is similar to the process of a UE performing ‘initial attach’ (see Sections 2.2.1.1 and 3.2.3.2). The only differences are that the S-GW/P-GW functionality is performed by the donor eNodeB and that during the RRC connection establishment, an RN signals an RN Indicator to the donor eNodeB. Based on this indicator, the donor eNodeB selects an MME that supports relay functionality.

After initial attachment, performed as if it was a UE, an RN has to enable and configure its RN functionality; this is done in two phases:

Phase I. The RN retrieves initial configuration parameters, including a list of donor eNodeBs to which it is allowed to attach, from an RN O&M server. The RN then detaches from the network as a UE and triggers Phase II.

Phase II. The RN connects to a donor eNodeB selected from the list acquired during Phase I. In order for the donor eNodeB to be authorized to pass S1-AP traffic to the RN, the donor eNB needs to know that the RN has an RN subscription rather than a UE subscription. The preferred approach for the donor eNodeB to obtain this information is via an S1-AP message from the MME, which can in turn obtain the relevant information from the HSS by means of an indication from the subscription profile in the HSS related to the Universal Subscriber Identity Module(s) (USIM(s)) in the RN.³ During Phase II, the donor eNodeB also embeds and provides the S-GW/P-GW-like functions needed to support the RN operation. This includes creating a session for the RN and managing EPS bearers for the RN, as well as terminating the S11 interface towards the MME serving the RN.

After the donor eNodeB has set up the bearers for the S1/X2 interfaces, the RN initiates the setup of S1 and X2 associations with the donor eNodeB and begins relay operations.

The RRC layer of the Un interface running on the backhaul link can send updated System Information (SI) in a dedicated message to RNs. When an RN receives such signalling, it applies it immediately.

In addition, RN-specific RRC functionality over the backhaul link is provided through a new Un reconfiguration procedure that applies after an RN has attached and commenced

³This indication may, for example, be based on subscription type, or on particular Access Point Names (APNs) which are used only by RNs.

relay behaviour; this procedure is independent of the Uu RRC reconfiguration procedure. The RN reconfiguration procedure can configure the backhaul subframe allocations and physical channels between the donor eNodeB and the RN. The RN may request such a configuration from the donor eNodeB during RRC connection establishment, or the donor eNodeB may initiate the RRC signalling for such configuration. The RN applies a new configuration immediately upon reception, so the subframe configurations on the Un interface and on the Uu interface in the RN cell may become temporarily misaligned until a corresponding reconfiguration of the latter can be completed.

30.3.5 Random Access on the Backhaul Link

RNs support both contention-based and non-contention-based random access (see Section 17.3 over the backhaul link in exactly the same way as a UE. When an RN performs the random access procedure, it suspends any current backhaul subframe configuration, meaning that it temporarily disregards any Un-interface-specific configuration. The backhaul subframe configuration is resumed after the random access procedure is successfully completed.

30.3.6 Radio Link Failure on the Backhaul Link

RNs detect radio link failure on the backhaul link using the same parameters as UEs (see Section 22.7). If an attempt to re-establish an RRC connection fails, then the RN goes to RRC_IDLE state and tries to recover. The details of RN behaviour from RRC_IDLE are left to RN implementation; it is assumed that an RN would always try to recover the RRC connection as quickly as possible and not remain in RRC_IDLE for longer than necessary.

If an RN experiences a Scheduling Request (SR)⁴ failure (e.g. a preconfigured number of SR attempts have been made), it follows the same procedure as a UE.

30.3.7 RN Security

RNs have different security considerations from an eNodeB or a UE. This arises from the fact that the behaviour of RNs comprises aspects of both eNodeB and UE behaviour, and because the connection of an RN to the network is wireless. The latter issue imposes additional security considerations and is especially important when the RN has to send RRC, S1 or X2 signalling messages over the backhaul radio bearers. Some such security threats are as follows (further details can be found in [25, Section 5.3]):

- Impersonation of an RN to attack the user(s) attached to the RN;
- Attacks on the backhaul link between the RN and donor eNodeB;
- Inserting a ‘Man in the Middle’ (MitM);
- Attacking the user data traffic;
- Impersonation of an RN to attack the network;
- Attacks on the RN itself;
- Denial of Service (DoS) attacks.

⁴See Section 4.4.2.2.

RNs support the Transport Network Layer (TNL) of the S1-MME and S1-U interfaces (see Figure 2.1), and hence a function to ensure secure transport over the backhaul link is needed. The solution for RN security (see [26, Annex D]) is to realize a one-to-one binding of an RN and a USIM called USIM-RN, which is not realized by eNodeBs. This one-to-one binding uses symmetric pre-shared keys or certificates. The donor eNodeBs store the IP security credentials in the secure part of the eNodeB platform, thus providing a secure anchor for Network Domain Security (NDS)/IP.

Over the Un interface, security is provided by Un PDCP between the RN and the donor eNodeB, and by NDS/IP between the donor eNodeB and the MME. The native Security Gateway (SEG) can be reused for NDS/IP traffic between the donor eNodeB and the MME. Un security is modified compared to Uu security such that integrity protection is provided in the Un user plane at least for PDCP PDUs carrying S1 signalling.

The security architecture for RNs is shown in Figure 30.13.

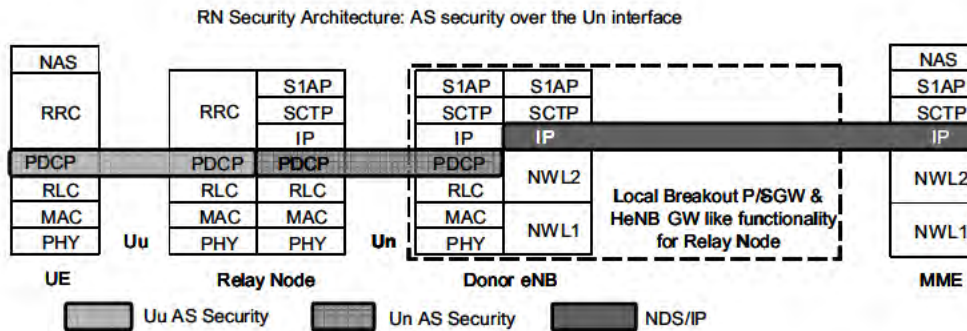


Figure 30.13: RN security architecture. Reproduced by permission of © 3GPP.

More information about the security considerations for the backhaul link and RNs in general can be found in [25] and in [26].

Security over the access link is provided by the RN in the same way as a normal eNodeB (see Sections 3.2.3.1 and 4.2.3).

30.3.8 Backhaul Physical Channels

The nature of the relay backhaul link results in a number of differences between the physical channels or signals defined for the backhaul link and those defined for the access link. These are reviewed below; full details can be found in [27].

30.3.8.1 Backhaul Reference Signals

Most of the Reference Signals (RSs) defined for the access link are also usable on the backhaul link, including the cell-specific RSs (see Section 8.2.1) and the new UE-specific RSs defined in Release 10 (see Section 29.1.1).

One exception may occur if the last OFDM symbol of a subframe is not within the RN's backhaul downlink reception window, in which case the second slot of subframes addressed

to such an RN does not carry the RSs of the UE-specific DeModulation RS (DM-RS) antenna ports 7 to 10. In this case, only the DM-RSs from the first slot are available for channel estimation and demodulation.

A further restriction is that the backhaul link in Release 10 is limited to a maximum of four spatial layers, and DM-RS antenna ports 11 to 14 are never transmitted in Physical Resource Blocks (PRBs)⁵ addressed to an RN.

30.3.8.2 Backhaul Control Channels

As explained in Section 30.3.2, the time-division separation of transmission and reception for the Type 1 RN in LTE-Advanced utilizes MBSFN subframes.

Due to the timing synchronization between the access and backhaul links, and the need to transmit one or two OFDM symbols for the PDCCH on the access link at the start of each MBSFN subframe, the RN is not able to receive the PDCCH from the donor eNodeB. To cater for this, a new backhaul control channel known as the Relay PDCCH (R-PDCCH) is provided. The R-PDCCH has to be multiplexed into the backhaul subframes within the RN's reception window, and it therefore occupies Resource Elements (REs) that otherwise would be used for the PDSCH. In the first slot of a subframe, the R-PDCCH starts in the fourth OFDM symbol (even though the RN's reception window may start earlier), and only R-PDCCHs containing downlink assignments are transmitted, enabling the RN to try to decode the downlink assignments immediately after the first slot. In the second slot, R-PDCCHs containing uplink grants are transmitted (the last OFDM symbol being determined by the length of the reception window (see Section 30.3.2.2)). This is illustrated in Figure 30.14.




Figure 30.14: Placement of the R-PDCCH.

Since the backhaul link conditions are likely to be relatively static and predictable, the R-PDCCH is mapped to certain PRBs instead of being distributed in frequency over the system bandwidth like the PDCCH. The PRBs used for the R-PDCCH are configurable for each RN. Within these PRBs, two alternative R-PDCCH formats are supported: a 'cross-interleaved' R-PDCCH and a 'non-cross-interleaved' R-PDCCH. These are illustrated in Figure 30.15 and described below.

Cross-interleaved R-PDCCH

The cross-interleaved R-PDCCH is only supported if the RN is configured to use the donor eNodeB's cell-specific RSs for demodulation of the R-PDCCH.

⁵See Section 6.2.




Figure 30.15: R-PDCCH formats: (a) cross-interleaved; (b) non-cross-interleaved.

The cross-interleaved R-PDCCH is designed to re-use as far as possible the design of the PDCCH. It is constructed in the same way as a regular Uu PDCCH (see Section 9.3.5) in respect of multiplexing, scrambling, modulation and mapping to resource elements and antenna ports, with the following exceptions:

- A Resource Element Group (REG) is composed of 4 consecutively available REs in one OFDM symbol in a PRB, counted in ascending order of subcarriers, where an RE is assumed not to be available for the R-PDCCH in the following cases:
 - if the RE is used for cell-specific RSs;⁶
 - if any CSI-RSs (see Section 29.1.2) are present in a subframe (even if they are muted), all REs for transmission of CSI-RSs on antenna ports 15 to 22 are unavailable for the R-PDCCH in that subframe.
- UE-specific RSs are not mapped onto PRB pairs used for the transmission of cross-interleaved R-PDCCH.

These exceptions are designed to simplify the RN receiver design and to reduce the need for excessive configuration parameters, even though it means that a few REs are not able to be used.

Non-cross-interleaved R-PDCCH

The non-cross-interleaved R-PDCCH can be supported regardless of whether the RN is configured to use the donor eNodeB's cell-specific RSs or the UE-specific DM-RSs for demodulation of the R-PDCCH.

⁶If the cell-specific RSs are configured to be transmitted only on antenna port 0, the REs that would be used for transmission of cell-specific RSs on antenna port 1 are also deemed to be unavailable for the R-PDCCH.

With this format, a single R-PDCCH occupies either 1, 2, 4, or 8 contiguous virtual resource blocks, depending on the aggregation level used for the R-PDCCH. The virtual resource blocks are configured for each RN individually by RRC signalling, using any of the PDSCH resource allocation types 0, 1 or 2 (see Section 9.3.5.4).

Scrambling, modulation and mapping to antenna ports follow the same methods as for the PDCCH. However, in contrast to the PDCCH and the cross-interleaved R-PDCCH, the mapping to REs fills REs across the aggregated VRBs before using the next OFDM symbol (not including REs that carry RS). This helps the RN to detect correctly the aggregation level of the R-PDCCH, which can be important for correctly determining whether R-PDCCH and PDSCH are multiplexed into the same PRB pair.

If the non-cross-interleaved R-PDCCH is transmitted to an RN that is configured to receive the PDSCH on the backhaul link in transmission mode 9 (see Section 9.2.2.1), the control information on the R-PDCCH indicates which and how many UE-specific RS antenna ports are transmitted. However, this information affects the R-PDCCH decoding due to the fact that REs used for UE-specific RSs are unavailable for the R-PDCCH itself. Therefore, in order to avoid excessive receiver complexity, REs should not be used for R-PDCCH transmission in the first slot if they could potentially be used for UE-specific RSs according to the maximum possible rank for PDSCH transmission according to the RRC configuration (the parameter ‘codebookSubsetRestriction-r10’).

R-PDCCH Search Space

For both R-PDCCH formats, similar blind decoding and search space functionality is defined as for the Uu link (see Section 9.3.5.5). However, for the R-PDCCH, the search space is defined separately in the first and second slots. The number of blind decoding trials per slot for the RN is 6 for aggregation levels 1 and 2, and 2 for aggregation levels 4 and 8.

Control Format Indicator

A backhaul version of the PCFICH is not needed, because the size of the control signalling region of the backhaul subframes can be configured semi-statically by RRC signalling. REs that are semi-statically assigned for RN control signalling but are not actually needed for that purpose can still be used by the donor eNodeB for normal data transmissions to UEs in its cells; in such cases, the RN’s blind decoding attempt for those REs would simply fail.

ACK/NACK feedback for backhaul uplink transmissions

On the Uu interface, the Physical HARQ Indicator CHannel (PHICH – see Section 9.3.4) carries the HARQ ACK/NACK information in response to uplink data transmissions on the Physical Uplink Shared CHannel (PUSCH). In a synchronous non-adaptive HARQ protocol (see Section 4.4.1.1) the PHICH can be seen as a resource-efficient retransmission request, since it is tied to a given PUSCH transmission. However, for the backhaul uplink, a synchronous non-adaptive HARQ operation is not always straightforward, and therefore an RN-specific PHICH for the backhaul was not considered to be sufficiently beneficial to be worth specifying such a channel. Consequently, all backhaul uplink retransmissions have to be triggered by an R-PDCCH, and are therefore adaptive. If an RN does not detect

an R-PDCCH to trigger a PUSCH transmission, the RN simply does not make a PUSCH transmission.

Backhaul uplink control signalling

The uplink control signalling is transmitted by the RN in the same way as specified for the Uu link (see Section 16.3). However, the transmission resources for the backhaul PUCCH are determined by higher layer signalling and can therefore be considered as reserved. Consequently, in contrast to the PDCCH, there is no relationship between the resources used for R-PDCCH transmission and the resources used for PUCCH transmission.

30.3.8.3 Backhaul Data Transmission

Data transmission on the backhaul link uses the same physical channels (PDSCH and PUSCH) as are defined for the access link. For the multiplexing of the PDSCH and R-PDCCH, three possibilities exist in terms of occupation by R-PDCCHs within a pair of backhaul PRBs (illustrated in Figure 30.16):

- a downlink assignment in the first slot and an uplink grant in the second slot;
- only a downlink assignment in the first slot;
- only an uplink assignment in the second slot.

In the second case, it is desirable to use the second slot of such a PRB pair for PDSCH data transmission to the RN to which the R-PDCCH in the first slot is addressed, instead of wasting the PRB by not using it. Therefore, a downlink assignment may include the PRB pair containing the downlink assignment itself. In this case, the RN is aware that the first slot was used to transmit the R-PDCCH, and can therefore correctly assume that only the second slot of that pair of PRBs carries PDSCH data. Although this is in principle applicable to both the cross-interleaved and the non-cross-interleaved R-PDCCH formats, the non-cross-interleaved format benefits more because a second slot is either completely used or completely unused for an R-PDCCH transmission; by contrast, for the cross-interleaved R-PDCCH, a partial occupation of the second slot by an R-PDCCH is likely.

On the other hand, if a donor eNodeB uses at least part of the second slot of a pair of PRBs for any kind of R-PDCCH transmission, that pair of PRBs should not be assigned to any RN;

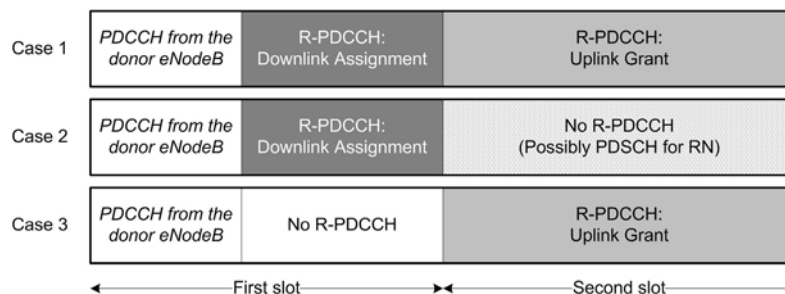


Figure 30.16: Supported R-PDCCH/PDSCH multiplexing possibilities.

otherwise the RN would have to perform additional blind decoding to establish whether the second slot contained data or an uplink grant.

In the third case, it could, in theory, be possible to use the resources of the first slot of such a pair of PRBs for a PDSCH transmission to the same or another RN. However, the practical benefit is smaller compared to the second case, because the available number of REs that could be used for a PDSCH in the first slot is substantially smaller than in the second slot due to the constraint of the RN's backhaul downlink reception window. Moreover, it is simpler to transmit a downlink assignment in the first PRB of the pair (as in the first case) than to use it for a PDSCH. Therefore, the transmission of PDSCH in the first PRB in the third case is not supported.

30.3.9 Backhaul Scheduling

In the case of an outband RN, there are no special restrictions for the backhaul scheduler located at the donor eNodeB, because no adverse effects from or for the access link need to be taken into account. The donor eNodeB can schedule an outband RN for backhaul uplink and downlink just like any UE in the donor eNodeB's cells.

On the other hand, in the case of an inband Type 1 RN, the donor eNodeB needs to consider at least the following constraints:

- which subframes are available for backhaul data transmission or reception, as dictated typically by the MBSFN subframe configuration advertised by the RN;
- the availability of transmission resources for feedback to or from the RN with suitable timing for the corresponding data transmissions.

Regardless of such constraints, the scheduler in the donor eNodeB should also try to take into account the added delay arising from the additional network hop from the donor eNodeB to the RN. Depending on the overall delay requirements, the donor eNodeB should limit the required number of backhaul retransmissions, as each backhaul retransmission adds at least 8 ms to the total delay experienced (see Section 10.3.2.5). A donor eNodeB may, for example, use more conservative link adaptation for an RN than it would for a regular UE.

For FDD operation (i.e. using Frame Structure Type 1), the subframes configured for backhaul transmission follow a periodicity of 8 ms, with the exception of subframes that cannot be used for the downlink backhaul because they cannot be configured as MBSFN subframes on the access link. An 8-bit bitmap is therefore sufficient to configure the downlink backhaul subframes. For every configured downlink backhaul subframe, a corresponding uplink backhaul subframe 4 ms later is configured. An uplink grant transmitted in a downlink backhaul subframe is valid for the corresponding uplink backhaul subframe. Similarly, a PDSCH transmission in a downlink backhaul subframe is associated with an ACK/NACK transmission in the corresponding uplink backhaul subframe.

For TDD operation (i.e. using Frame Structure Type 2), the backhaul and access subframe configurations must be mutually compatible. While Frame Structure Type 2 supports a total of seven uplink-downlink configurations (0 to 6 – see Section 6.2), an RN only needs to support five of these configurations for the access link. For most of these access link

configurations, several backhaul configurations exist, allowing a variety of numbers and/or positions of the backhaul downlink and uplink subframes, as shown in Figure 30.17.⁷

Uplink-downlink configuration	Backhaul DL:UL ratio	Backhaul subframe configuration	Subframe number									
			0	1	2	3	4	5	6	7	8	9
1	1:1	0	↘	⊗	↗	↗	▼	↘	⊗	↗	▲	↘
		1	↘	⊗	↗	▲	↘	↘	⊗	↗	↗	▼
	2:1	2	↘	⊗	↗	↗	▼	↘	⊗	↗	▲	▼
		3	↘	⊗	↗	▲	▼	↘	⊗	↗	↗	▼
		4	↘	⊗	↗	▲	▼	↘	⊗	↗	▲	▼
2	1:1	5	↘	⊗	▲	↘	↘	↘	⊗	↗	▼	↘
		6	↘	⊗	↗	▼	↘	↘	⊗	▲	↘	↘
	2:1	7	↘	⊗	▲	↘	▼	↘	⊗	↗	↘	▼
		8	↘	⊗	↗	▼	↘	↘	⊗	▲	↘	▼
		9	↘	⊗	▲	▼	↘	↘	⊗	↗	▲	▼
		10	↘	⊗	↗	▼	↘	↘	⊗	▲	↘	▼
3	2:1	11	↘	⊗	↗	▲	↗	↘	↘	▼	▼	▼
	3:1	12	↘	⊗	↗	▲	↗	↘	↘	▼	▼	▼
4	1:1	13	↘	⊗	↗	▲	↘	↘	↘	▼	▼	▼
	2:1	14	↘	⊗	↗	▲	↘	↘	↘	▼	▼	▼
		15	↘	⊗	↗	▲	↘	↘	↘	▼	▼	▼
	3:1	16	↘	⊗	↗	▲	↘	↘	↘	▼	▼	▼
	4:1	17	↘	⊗	↗	▲	▼	↘	↘	▼	▼	▼
6	1:1	18	↘	⊗	↗	↗	▲	↘	⊗	↗	▼	

Legend	
▼ Backhaul link, downlink subframe	↘ Access link, downlink subframe
▲ Backhaul link, uplink subframe	⊗ Access link, special subframe
	↗ Access link, uplink subframe

Figure 30.17: TDD subframe configurations.

For a backhaul uplink grant transmitted in backhaul downlink subframe n by the donor eNodeB, the corresponding backhaul PUSCH transmission by the RN occurs in backhaul uplink subframe $n + k_1$, where k_1 is determined in the same way as for the Uu interface.

For a backhaul downlink PDSCH transmission by the donor eNodeB in subframe n , the corresponding HARQ ACK/NACK feedback is transmitted by the RN in backhaul

⁷The RN is not required to use the same uplink-downlink configuration for the access link as for the backhaul link; however, it is likely that this can be assumed in most, if not all, practical deployments. Figure 30.17 reflects this assumption.

uplink subframe $n + k_2$, where k_2 is the smallest integer greater than or equal to 4. If multiple ACK/NACKs need to be transmitted in a single subframe, the same ‘bundling’ and ‘multiplexing’ methods are available as for the Uu interface (see Section 23.4.3).

30.3.10 Backhaul HARQ

30.3.10.1 Frame Structure Type 1

In the downlink, the access link HARQ is operated in an asynchronous manner. Each downlink transmission resource assignment on the PDCCH includes an explicit indication of the HARQ process number of the data transmission, so the same design can be applied to the downlink backhaul link.

For the uplink, the FDD access link HARQ is designed as a synchronous protocol with a periodicity of 8 ms. However, the first, fifth, sixth and tenth subframes in each 10 ms radio frame cannot be designated as MBSFN subframes on the access link due to the presence of synchronization and broadcast signalling (see Sections 7.2 and 9.2.1), and therefore a regular 8 ms pattern cannot be established for the backhaul link. As shown in Figure 30.18, the consequence is that two transmission opportunities for uplink data on the backhaul link (in the subframes marked 13 and 29) cannot be used since no corresponding resource grant can be received on the R-PDCCH 4 ms earlier (in the subframes marked 9 and 25).




Figure 30.18: Uplink backhaul HARQ timing with 8 ms periodicity (FDD).

The backhaul uplink therefore uses a modification of the access link protocol. The synchronous aspect of the identification of HARQ processes is retained in such a way that the HARQ processes are mapped onto backhaul uplink subframes sequentially; however, the time aspect of the synchronous operation is modified. This is similar to the way in which the uplink HARQ works in the TDD access link, where the availability of subframes for uplink transmission is configured by an RRC parameter. In order to simplify the backhaul uplink subframe configuration and to disrupt as few HARQ processes as possible on the access uplink, only configurations with a basic periodicity of 8 ms are supported. Since an uplink transmission needs to be assigned by an uplink grant 4 ms earlier by means of an R-PDCCH, this implies that if a downlink subframe cannot be used for backhaul transmission due to an MBSFN subframe restriction on the access link, the corresponding uplink subframe cannot be configured as a backhaul uplink subframe, even if it were necessary to do so to follow the 8 ms periodicity.

A minimum Round Trip Time (RTT) of 8 ms applies for the backhaul HARQ (just as for the access link HARQ). In order to minimize the delay, the number of backhaul uplink HARQ processes is minimized as a direct function of the backhaul uplink subframe configuration. This results in the number of HARQ processes on the backhaul uplink being between 1 and 6.

30.3.10.2 Frame Structure Type 2

In the downlink, the access link HARQ is operated in an asynchronous manner. Each downlink transmission resource assignment on the PDCCH includes an explicit indication of the HARQ process number of the PDSCH transmission, so the same design can be applied to the backhaul downlink. In the uplink, a synchronous protocol is employed with a fixed RTT of 10 ms, regardless of the uplink-downlink configuration. For backhaul subframe configuration 4 (see Figure 30.17), two uplink HARQ processes are employed, while for all other backhaul subframe configurations, only a single uplink HARQ process is used.

30.4 Summary

The introduction of relaying represents a significant new step in the LTE radio access network as it is evolved for LTE-Advanced. The incorporation of Relay Nodes (RNs) in an LTE network can offer attractive cost advantages for network operators compared to a homogeneous deployment of macro eNodeBs. Key deployment scenarios for RNs include coverage extension, to extend the range of macro-cells, and capacity enhancement, to boost the supported data rates within the coverage area of a macro-cell.

In LTE-Advanced, RNs operate as a type of eNodeB, with a wireless backhaul link that is specially designed to reuse as much of the LTE access link technology as possible to connect to a donor eNodeB. This wireless backhaul link may operate on the same or a different frequency as the access link from the RN to the RN-assisted UEs (inband and outband operation respectively), depending on the available spectrum of the network operator.

References⁸

- [1] 3GPP Technical Specification 25.106, 'UTRA repeater radio transmission and reception', www.3gpp.org.
- [2] 3GPP Technical Specification 36.106, 'Evolved Universal Terrestrial Radio Access (E-UTRA); FDD Repeater radio transmission and reception', www.3gpp.org.
- [3] Nokia Siemens Networks and Nokia, 'R1-100353: Comparing In-band vs. Out-band Relays in coverage limited scenario', www.3gpp.org, 3GPP TSG RAN WG1, meeting 59bis, Valencia, Spain, January 2010.
- [4] Panasonic, 'R1-082397: Discussion on the various types of relays', www.3gpp.org, 3GPP TSG RAN WG1, meeting 54, Warsaw, Poland, June 2008.
- [5] Ericsson, 'R1-082024: A discussion on some technology components for LTE-Advanced', www.3gpp.org, 3GPP TSG RAN WG1, meeting 53, Kansas City, USA, May 2008.
- [6] T. M. Cover and A. A. El Gamal, 'Capacity Theorems for the Relay Channel', *IEEE Trans. Inf. Theory*, Vol. 25, pp. 572–584, September 1979.

⁸All web sites confirmed 1st March 2011.

- [7] G. Kramer and A. J. van Wijngaarden, ‘On the White Gaussian Multiple-Access Relay Channel’ in *Proc. IEEE Int. Symp. Information Theory (ISIT)*, (Sorrento, Italy), p. 40, June 2000.
- [8] G. Kramer, M. Gastpar, and P. Gupta, ‘Cooperative Strategies and Capacity Theorems for Relay Networks’, *IEEE Trans. Inf. Theory*, Vol. 51, pp. 3037–3063, September 2005.
- [9] G. Kramer, I. Marić, and R. D. Yates, ‘Cooperative Communications’, Vol. 1 of *Foundations and Trends in Networking*. Hanover, MA: now Publishers, Inc., 2007.
- [10] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York: John Wiley & Sons, Inc., 1991.
- [11] J. G. Proakis, *Digital Communications*, Fourth Edition. New York: McGraw-Hill, Inc., fourth ed., 2001.
- [12] T. S. Rappaport, *Wireless Communications: Principles and Practice*, Second Edition. Upper Saddle River, New Jersey: Prentice-Hall, Inc., 2002.
- [13] C. E. Shannon, ‘A Mathematical Theory of Communication’, *Bell Syst. Tech. J.*, Vol. 27, pp. 379–423, 623–656, 1948.
- [14] I. E. Telatar, ‘Capacity of Multi-Antenna Gaussian Channels’, *European Trans. on Telecomm.*, Vol. 10, pp. 585–596, November-December 1999.
- [15] G. J. Foschini and M. J. Gans, ‘On Limits of Wireless Communications in a Fading Environment when Using Multiple Antennas’, *Wireless Personal Communications*, Vol. 6, pp. 311–335, March 1998.
- [16] J. N. Laneman, D. N. C. Tse, and G. W. Wornell, ‘Cooperative Diversity in Wireless Networks: Efficient Protocols and Outage Behavior’, *IEEE Trans. Inf. Theory*, Vol. 50, pp. 3062–3080, December 2004.
- [17] K. Azarian, H. El Gamal, and P. Schniter, ‘On the Achievable Diversity-Multiplexing Tradeoff in Half-Duplex Cooperative Channels’, *IEEE Trans. Inf. Theory*, Vol. 51, pp. 4152–4172, July 2005.
- [18] Sikora M., J. N. Laneman, M. Haenggi, D. J. Costello, and T. E. Fuja, ‘Bandwidth- and Power-Efficient Routing in Linear Wireless Networks’, *IEEE Trans. Inf. Theory*, Vol. 52, pp. 2624–2633, June 2006.
- [19] R. Koetter and M. Medard, ‘An Algebraic Approach to Network Coding’, *IEEE/ACM Trans. Netw.*, Vol. 11, pp. 782–795, October 2003.
- [20] S. Katti, S. Gollakota, and D. Katabi, ‘Embracing wireless interference: analog network coding’ in *Proc. Conf. Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM)*, pp. 397–408, 2007.
- [21] S. Katti, H. Rahul, W. Hu, D. Katabi, M. Médard, and J. Crowcroft, ‘XORs in the air: practical wireless network coding’, *IEEE/ACM Trans. Netw.*, Vol. 16, pp. 497–510, may 2008.
- [22] Alcatel-Lucent and CHTTL, ‘R1-092321: System Design Frameworks to Support Type II Relay Operation in LTE-A’, www.3gpp.org, 3GPP TSG RAN WG1, meeting 57bis, Los Angeles, USA, June 2009.
- [23] Samsung, ‘R1-100139: Full duplex configuration of Un and Uu subframes for Type I relay’, www.3gpp.org, 3GPP TSG RAN WG1 meeting 59bis, Valencia, Spain, January 2010.
- [24] 3GPP Technical Specification 36.331, ‘Evolved Universal Terrestrial Radio Access (E-UTRA) Radio Resource Control (RRC); Protocol specification’, www.3gpp.org.
- [25] 3GPP Technical Report 33.816, ‘Feasibility Study on LTE Relay Node Security (Release 10)’, www.3gpp.org.
- [26] 3GPP Technical Specification 33.401, ‘3GPP System Architecture Evolution (SAE); Security architecture’, www.3gpp.org.
- [27] 3GPP Technical Specification 36.216, ‘Evolved Universal Terrestrial Radio Access (E-UTRA); Physical Layer for Relaying Operation (Release 10)’, www.3gpp.org.

Additional Features of LTE Release 10

Teck Hu, Philippe Godin and Sudeep Palat

31.1 Introduction

In addition to the features described in Chapters 28 to 30, a number of additional enhancement features are included in Release 10 of the LTE specifications. These are explained in the following sections.

31.2 Enhanced Inter-Cell Interference Coordination

An important feature of LTE-Advanced is the support of heterogeneous network deployments consisting of small cells overlaid within the coverage area of a macrocellular network. The small cells may comprise picocells or femtocells (see Chapter 24) or Relay Nodes (RNs) (see Chapter 30). Overlaying small cells in this way enables higher spectral reuse due to cell-splitting, which delivers an increase in capacity and can support localized high traffic-densities ('hot-spots'). Overlaying small cells is particularly attractive when the availability of suitable sites for macrocells is limited. An alternative would be to use separate spectrum (see Section 28.3.1.1), but this depends on the availability of such additional spectrum.

However, the improved system capacity from heterogeneous deployments sharing the same spectrum present interference scenarios that are different from a homogeneous macrocellular networks. Sections 31.2.1 to 31.2.6 explore these scenarios in more detail and explain the enhanced Inter-Cell Interference Coordination (eICIC) techniques introduced in Release 10 to address them.

LTE – The UMTS Long Term Evolution: From Theory to Practice, Second Edition.

Stefania Sesia, Issam Toufik and Matthew Baker.

© 2011 John Wiley & Sons, Ltd. Published 2011 by John Wiley & Sons, Ltd.

The enhanced ICIC techniques complement those already available since Release 8 (see Sections 12.5.1 and 12.5.2) and are designed for the case of single-carrier operation (i.e. without carrier aggregation).

Co-channel heterogeneous network deployments can in general be categorized into two scenarios: macro-pico and macro-femto. In the macro-pico case, the small cells are Open Subscriber Group (OSG) cells, open to all users of the macrocellular network. In the macro-femto case, the small cells are Closed Subscriber Group (CSG) cells accessible only to a limited group of users. Although both scenarios are considered in Release 10, the main focus of the eICIC techniques, and hence of this section, is the OSG macro-pico scenario.¹ The differences in the applicable eICIC solutions between the two scenarios are mentioned where relevant. Figure 31.1 illustrates the interference scenarios for the macro-pico scenario in both downlink and uplink.

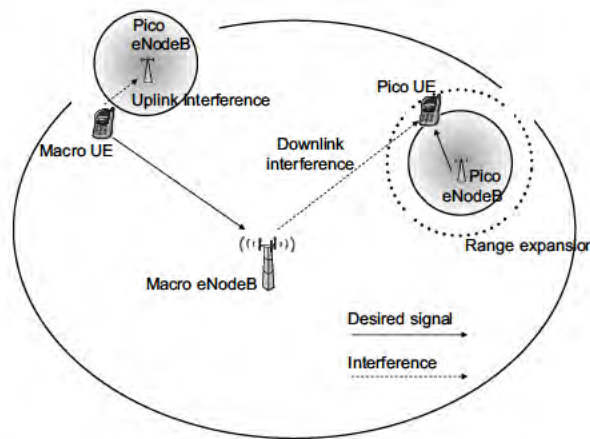


Figure 31.1: Heterogeneous network interference scenarios in downlink and uplink.

In a homogeneous macrocellular network, a UE is served by the strongest cell, and the number of UEs served by the picocells would be limited due to the lower power of the picocells compared to the macrocells. To derive the full gain of cell splitting, the serving eNodeB can intentionally ‘bias’ the handover offset values of some UEs in RRC_CONNECTED mode to transfer them to the picocells. This has the effect of expanding the coverage of the picocells and is therefore sometimes known as *Cell Range Expansion* (CRE). If a UE is transferred in this way, the signal from the pico eNodeB would be low due to its lower transmission power, while the interference received from the macro eNodeB would be significantly higher, as highlighted in Figure 31.1.

To illustrate the level of interference, Figure 31.2 plots the Cumulative Distribution Function (CDF) of the received Signal-to-Interference-plus-Noise Ratio (SINR) with different bias values used for the CRE. The locations of the picocells and the UE distributions are according to configuration 4b in [1, Section A.2.1.1.2], with four picocells per macrocell.² The figure clearly shows how the SINR degrades as the bias is increased.

¹The corresponding interference scenarios for the macro-femto case are discussed in Section 24.3.1.

²Additional simulation results can be found in [2].




Figure 31.2: SINR for macro-pico scenario with various degrees of range expansion.

31.2.1 LTE Interference Management

The interference mitigation and ICIC techniques available in Release 8 can be summarized as frequency-domain scheduling (including frequency-selective scheduling and frequency hopping), power setting, increasing robustness (for example by beamforming or interference cancellation for the data channels, or by increasing the number of Control Channel Elements (CCEs) used for each message on the PDCCH – see Section 9.3.5) and X2-interface-based data channel interference coordination (see Sections 12.5.1 and 12.5.2).

Interference avoidance based on frequency-domain partitioning between different cells are of limited benefit for the synchronization signals, Physical Broadcast Channel (PBCH), cell-specific Reference Signals (RSs) or control channels (PDCCH, PCFICH, and PHICH). These are needed for initial access to the network and/or thereafter for maintaining the radio link. Therefore, their time-frequency locations are fixed (with the exception of the cell-specific RSs, which can use a frequency reuse factor of 3 or 6 depending on the number of antenna ports configured) and frequency partitioning of these channels and signals would not be backward-compatible with Release 8/9 UEs. However, the interference experienced by the picocell UEs in a co-channel macro-pico deployment affects these channels also, and, if very large range expansion is employed, the control channel reception at the picocell UEs may fail, resulting in outage.

The need for interference mitigation of the control channels was the motivation for the specification of time-domain-based ICIC in Release 10. The overall objective of this eICIC is to mute certain subframes of one layer of cells in order to reduce the interference to the other layer. These muted subframes are called *Almost Blank Subframes* (ABSs) [3].

31.2.2 Almost Blank Subframes

ABSs are defined as subframes with reduced downlink transmission power and/or activity. Ideally, ABSs would be totally blank when configured by a macrocell in a macro-pico scenario, in order to remove as much as possible of the interference towards the UEs

served by the picocells. In practice, the transmission and exact content of ABSs would be an implementation choice, taking into account the gains from interference reduction versus the loss of transmission resources from being unable to transmit PDSCH data in the ABSs. As one example, uplink PUSCH data traffic could continue to be scheduled by means of uplink grants on the PDCCH, and corresponding HARQ³ ACK/NACK feedback could be transmitted, while downlink assignments might be stopped.

The important consideration of backward-compatibility means that cells must remain accessible and measurable for Release 8/9 UEs. As a result, the cell-specific RSs (and CSI-RS if configured), PSS, SSS, Paging Channel (PCH) and PBCH must also be transmitted in ABSs. Nevertheless, even with these transmissions, the ABSs can contain much less energy than normal subframes and thus reduce the interference towards both the data and control signalling of the corresponding subframes of the co-channel victim cell.

A subframe designated as an ABS can be either a normal subframe or an MBSFN subframe,⁴ as illustrated in Figure 31.3.




Figure 31.3: Normal and MBSFN Almost Blank Subframes.

ABSs which are also MBSFN subframes can reduce interference to the victim cell even more than ABSs which are not MBSFN subframes, since the former do not have cell-specific RSs in the PDSCH region and cannot contain PBCH or PCH transmissions. However, only subframes 1, 2, 3, 6, 7 and 8⁵ can be configured as MBSFN subframes, and, as a result, MBSFN subframes alone cannot fully support the 8 ms periodicity of the synchronous uplink HARQ processes. Therefore, if uplink scheduling is required during downlink ABSs, some of the downlink ABSs would need to occur in subframes which cannot be configured as MBSFN subframes.

³Hybrid Automatic Repeat reQuest.

⁴Multimedia Broadcast Single Frequency Network subframe – see Section 13.4.1.

⁵For FDD; indexing starting from zero.

31.2.3 X2 Interface Enhancements for Time-Domain ICIC

The characteristics of ABSs allow a high degree of flexibility in their implementation. To ensure their effectiveness when configured by one eNodeB, coordination is required with the neighbouring eNodeBs they are intended to benefit. For this purpose, time-domain inter-cell coordination signalling via the X2 interface is defined in Release 10, whereby the pattern of ABSs configured by one eNodeB can be indicated to its neighbours.

This signalling comprises ‘ABS bitmaps’, of which two can be signalled over the X2 interface from one cell (typically a macrocell) to another (typically a picocell).⁶

The ABS patterns are updated ‘semi-statically’, i.e. not faster than the Release 8/9 frequency-domain RNTP⁷ indications.

The first ABS bitmap, known as the ABS Pattern, indicates the complete set of ABS subframes to assist the receiving eNodeB with its scheduling operations. Like the RNTP indications, each bit of an ABS bitmap informs the receiving eNodeB of the sending cell’s intention regarding its transmit power – but on a per-subframe basis rather than a per-RB basis. This information can be used by the receiving eNodeB to arrange its scheduling operations to avoid transmitting in non-ABS subframes to UEs believed to be in high-interference areas such as the cell edge. Instead, the receiving eNodeB can schedule such UEs in the ABS subframes where they should experience less interference. As with frequency-domain ICIC, the exact operation of ABS-based time-selective scheduling is left up to the eNodeB implementation.

The second ABS bitmap, known as the ‘Measurement Subset’, is a subset of the first and is expected to be less variable; it is intended to recommend to the receiving eNodeB a suitable set of subframes that can be used for measurements by the UEs that would potentially be suffering interference in the subframes that are not configured as ABSs.

In order to align with the uplink HARQ round trip times, the ABS bitmaps have the following periodicities:

- 40 ms for FDD;
- 20 ms for TDD configurations 1–5;
- 70 ms for TDD configuration 0
- 60 ms for TDD configuration 6.

The ABS bitmap signalling can be transferred as part of the ‘Load Indication’ and ‘Resource Status Reporting Initiation’ procedures over the X2 interface. In addition, a message known as an ‘Invoke Indication’ allows one eNodeB to request ABS configuration at another. Such a request would be based on the requesting eNodeB’s assessment of the interference situation in its cell(s). The request initiates the Load Indication procedure, and the resulting ‘Load Information’ message from the eNodeB receiving the request contains the ABS bitmaps.

An eNodeB which has been informed about configured ABSs in another cell can return an ‘ABS Status’ message to the eNodeB that configured the ABSs. This message assists the latter eNodeB in determining whether the number of configured ABSs might need to be increased

⁶ABS bitmaps make no distinction as to the cell type; ABSs can be configured by macro, pico or femto cells, but, in Release 10, ABS bitmap signalling over the X2 interface is not supported from a macro eNodeB to a CSG Home eNodeB. ABS configuration in CSG cases is performed through Operation and Maintenance (O&M).

⁷Relative Narrowband Transmit Power – see Section 12.5.1.

or decreased, by indicating the percentage resource blocks in the currently configured ABSs that are allocated for UEs served by the victim eNodeB.

In scenarios where the UEs served by a pico eNodeB suffer interference from more than one strong interfering macro eNodeB, each with potentially different ABS patterns, the set of usable ABSs may be different from what any one individual macro eNodeB has indicated to the pico eNB. For the macro eNodeB to interpret correctly the ABS Status information from the pico eNodeB, the pico eNodeB should indicate the actual usable ABS pattern (in the form of a bitmap in the information element ‘Usable ABS Pattern Info’) to the macro eNodeB together with the ABS status information.

Figure 31.4 shows a typical message exchange over the X2 interface for ABS coordination between a macro eNodeB and a pico eNodeB. Box A in Figure 31.4 shows the invoke function being used to request the the macro eNodeB to configure ABSs. The macro eNodeB may take this request into consideration when configuring ABSs; the exact relationship between a request and any subsequent decision to configure ABSs is not standardized.




Figure 31.4: A typical exchange of ABS information over the X2 interface.

31.2.4 UE Measurements in Time-Domain ICIC Scenarios

For UEs which are suffering inter-cell interference, the measurement results will depend substantially on whether the measurements are restricted to the ABSs (which are protected from interference) or are made in all subframes.

Figure 31.5 illustrates the various UE measurement scenarios which are affected by the existence of protected subframes. The most common scenario for each of these measurements is when UEs served by a picocell suffer interference from a macrocell. UEs served by a macrocell do not need restrictions for their serving-cell measurements; however, they should use the configured ABSs to measure any picocells which might be handover targets, since during the ABS subframes the interference experienced at the picocell is lower than in other subframes; this allows interfered picocells to be detected by the macro UEs for possible handover.




Figure 31.5: UE measurements of serving and neighbour cells.

The affected measurements can be classified as Radio Link Monitoring (RLM) measurements, Radio Resource Management (RRM) measurements and Channel State Information (CSI) measurements.

31.2.4.1 RLM Measurements

The RLM procedure is used for a UE to monitor the downlink radio link quality of the serving cell for the purpose of determining whether synchronization has been lost and Radio Link Failure (RLF) has occurred (see Section 22.7). In order to avoid unnecessary loss of synchronization being reported by a UE, when ABSs are configured the RLM measurements should be made in those subframes only.

31.2.4.2 RRM Measurements

For RRM-related measurements such as intra-frequency cell reselection, the UE needs to perform RSRQ and RSRP measurements.⁸

RSRQ. Since the measured RSSI⁹ value used in deriving RSRQ will depend on whether the measurement is made in ABSs only or in all subframes, the RSRQ measurements need to be restricted to a limited set of subframes, which are indicated to the UE by RRC signalling. If the network restricts a UE's RSRQ measurements in this way, the RSSI definition is also modified so that it is measured over all OFDM symbols in the indicated subframes (rather than only the OFDM symbols containing cell-specific RSs). This modification reduces the impact of interference from the cell-specific RSs in the ABSs, especially in the case of ABSs which are MBSFN subframes; in MBSFN ABSs, cell-specific RSs are present only in the first

⁸Reference Signal Received Quality and Reference Signal Received Power – see Sections 22.3.1.2 and 22.3.1.1 respectively.

⁹Received Signal Strength Indicator.

symbol, and taking measurements only on this symbol would underestimate the RSRQ due to the strong neighbour-cell cell-specific RS interference. Figure 31.6 shows the effect of this modification on the RSRQ accuracy by comparing the CDFs of the RSRQ measurements with the Release 8 and Release 10 RSSI definitions.¹⁰ The case when the neighbour-cell cell-specific RS interference is perfectly cancelled by the UE is also shown for comparison.




Figure 31.6: RSSI measurement modification.

RSRP. Since RSRP is a measurement of cell-specific RSs, RSRP measurement values are not affected as strongly by whether they are taken in ABSs or non-ABSs, and therefore the RSRP definition is unchanged in Release 10. However, the RSRP measurement accuracy is likely to be strongly impacted when there is collision between the cell-specific RSs of an interfering cell and those of the cell being measured.

31.2.4.3 CSI Measurements

CSI measurements, especially the CQI reports, provide the serving eNodeB with information related to the SINR, which is useful for scheduling and selecting an appropriate Modulation and Coding Scheme (MCS) to use for each UE. When ABSs are configured in an interfering cell, the SINR experienced by the UEs in a victim cell will fluctuate dramatically from subframe to subframe. In order to exploit the potential performance gain of ABSs for scheduling and link adaptation, the CSI measurements made by these UEs should be performed on restricted subsets of the subframes.

31.2.5 RRC Signalling for Restricted Measurements

When an eNodeB receives the first and second bitmaps of the ABS information over the X2 interface from another eNodeB, it has to derive suitable subframe patterns to signal to the UEs it is serving, to restrict their measurements in the time-domain. The signalled patterns should

¹⁰The simulation assumptions used here can be found in [4].

be a subset of the ABS bitmaps received over the X2 interface, although the exact relationship between the patterns signalled over the X2 interface and the sets of subframes signalled to the UEs for restricted measurements is not standardized; it is entirely an implementation matter and can take into account other factors related to the scheduling algorithm.

Dedicated (i.e. UE-specific) RRC signalling is used to indicate to the UEs in connected mode¹¹ the sets of subframes to which their measurements should be restricted. The use of dedicated signalling (as opposed to broadcast) enables different UEs to be informed of different measurement subframe patterns, depending, for example, on their proximity to interfering neighbour cells.

The RRC signalling can configure three independent measurement-subframe patterns for a UE that supports time-domain measurement restrictions:

- **Pattern 1** applies to RLM and RRM measurements on the serving cell.
- **Pattern 2** applies to RRM measurements on certain intra-frequency neighbour cells with specifically indicated Physical Cell Identities (PCIs).¹² For neighbour cells whose PCIs are not listed, no RRM measurement restrictions apply; this is the default for neighbour cells which have not configured ABSs.
- **Pattern 3** applies to CSI measurements (including Channel Quality Indicator (CQI), Precoding Matrix Indicator (PMI) and Rank Indicator (RI)). This pattern enables the eNodeB to receive accurate CSI for the subframes in which the UE is expected to be scheduled. Pattern 3 comprises two independent sets of subframes per UE. This enables the serving eNodeB to compare a UE's CQI reports for different sets of subframes, for example for ABSs and non-ABSs, or for ABSs configured by two different neighbour (e.g. macro) eNodeBs which have both signalled their ABSs to the serving (e.g. pico) eNodeB. In the latter example, the two signalled sets of subframes could comprise one set that is common to both the interfering eNodeBs and another that contains only subframes configured as ABSs by one of the interfering eNodeBs; it should be noted that the two sets are not permitted to overlap.

For periodic CSI reporting (see Section 10.2.1.2), two sets of reporting periodicities and offsets are configured for the UE; each set is associated with one of the sets of subframes. For aperiodic CSI reporting (see Section 10.2.1.1), the UE reports CSI for whichever set of subframes contains the 'CQI reference resource' corresponding to the aperiodic CSI trigger (see [5, Section 7.2.3]).

Table 31.1 summarizes typical measurement restrictions that could be applied.

31.2.6 ABS Deployment Considerations

The decision as to whether and when to configure ABSs in a network involves a trade-off between capacity loss due to inability to schedule UEs in the ABSs of the interfering cells and capacity gain due to range expansion of the small cells. A number of factors should therefore be taken into account when reaching this decision for a practical deployment:

- Caution should be exercised if attempting to use CQI reports to identify the need for ABSs, since CQI reports are designed to indicate a suitable MCS for the PDSCH, for

¹¹Restriction of measurements to specific sets of subframes is not supported for UEs in idle mode in Release 10.

¹²Restricted sets of measurement subframes for inter-frequency measurements are not supported in Release 10.

which frequency-domain ICIC techniques have been shown to be very effective. CQI reports may not adequately predict the likelihood of PDCCH decoding failure.

- At low system load, the probability of users experiencing co-channel interference is expected to be low, as is the need for ABSs. At higher loadings, careful planning of the small cell coverage areas may result in sufficient UEs being served by the small cells without needing significant range expansion and ABSs to mitigate the outage probability for cell edge users.
- If ABSs are configured, the number of UEs handed over into the small cells due to range expansion will affect the number of ABSs required.
- The pattern of ABSs configured by an eNodeB needs to ensure that any active guaranteed bit-rate services can be supported in downlink and uplink. The uplink HARQ timeline must also be taken into account.

Some example ABS patterns are shown in Figure 31.7 for both normal ABSs and MBSFN ABSs. Normal ABS densities of 12.5% and 25% are shown, corresponding to 1 ABS and 2 ABSs respectively per eight subframes. An MBSFN ABS density of 15% is also shown. In all the ABS patterns shown, the 40 ms ABS pattern periodicity for FDD is maintained.

Table 31.1: Typical UE measurement restrictions.

	Time-domain measurement restrictions		
	RLM and serving-cell RRM measurements (Pattern 1)	Neighbour-cell measurements (Pattern 2)	CSI measurements (Pattern 3)
UE with severe interference	ABSs of interfering cell	ABSs of interfering cell (if in PCI list)	ABSs of interfering cell
UE without severe interference	No restriction	ABSs of interfering cell (if in PCI list)	No restriction

31.3 Minimization of Drive Tests

Drive tests, whereby UEs are driven through a deployed access network, are widely used to detect possible coverage holes or inefficiencies in the radio network planning. The Minimization of Drive Tests (MDT) feature aims to enable some of these drive test campaigns to be replaced by the automatic collection of UE measurements triggered by the eNodeBs, in order to reduce the cost of network optimization and carbon dioxide emissions.

Two modes of MDT are specified in Release 10 [6, 7]:

- **Logged MDT**, where the UE logs measurements in idle mode, together with time-stamps and, optionally, accurate location information;
- **Immediate MDT**, where the UE makes measurements in connected mode.

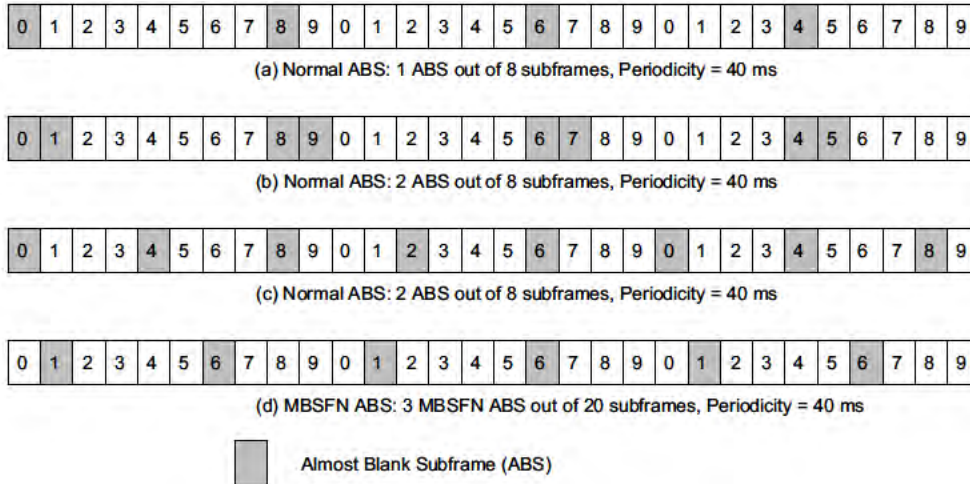


Figure 31.7: Examples of ABS patterns.

The two modes are configured by Operation and Maintenance (O&M) to give the operator more flexibility, even though the configuration and the measurements collected for the two modes may be different.

Both MDT modes can be requested as either a ‘Signalling-Based Trace’ or a ‘Management-Based Trace’ (see Section 2.5.8). Signalling-based MDT would typically be triggered when one specific UE is targeted, whereas with management-based MDT the eNodeB can pick any arbitrary UE that satisfies an MDT requirement.

The two MDT modes are specified for the control plane, but they could also potentially be used as diagnostic monitoring functions [8] in the user plane.

31.3.1 Logged MDT

To perform logged MDT, a UE is configured with the area in which it should log measurements (defined in terms of cells, tracking areas or a whole PLMN¹³), a logging interval, the logging duration and an absolute time reference which it has to feed back in the measurement log. Whenever the UE enters the configured area, it will start logging radio measurements (RSRP and RSRQ) together with the relative time compared to the configured time reference; it may additionally store the location (in terms of cell identity, GNSS¹⁴ coordinates or RF ‘fingerprints’ of neighbour cells (i.e. signal strengths of a number of intra-frequency neighbour cells)). The MDT log is suspended whenever the UE transitions to idle mode, changes RAT or exits its Registered PLMN, and it is resumed when the UE re-enters connected mode in its registered PLMN under the RAT in which it received its configuration.

The reporting of the log to the eNodeB can be done when the UE is in connected mode even if the logging period has not expired, provided that the UE is in a cell of the RAT in

¹³Public Land Mobile Network.

¹⁴Global Navigation Satellite System.

which it received the configuration. The UE first provides only an indication that MDT data is available. It is then up to the network to decide whether to retrieve the log at this point in time or later. The network thus keeps control of the MDT log reporting.

Dedicated RRC signalling is used to retrieve MDT logs, using both integrity protection and ciphering (see Section 4.2.3). Once successively delivered, the UE clears the data from memory. At the end of a logging period, the UE should keep the MDT log for a minimum period of 48 hours so that the network can retrieve it. Once retrieved, the MDT log is transferred to an MDT collection entity within the network where, in general, it is possible to associate the log with the UE identity. However, special care is given to privacy and data protection in Release 10, such that MDT trace information can only be collected with prior user consent, which users can give or revoke at any time. In addition, processing of management-based MDT data has to be handled in such a way that it does not allow identification of an individual user or UE and that it protects the privacy of individual user locations.

31.3.2 Immediate MDT

To perform immediate MDT, the eNodeB selects UEs that are in connected mode within the area where the MDT log is required and that have the necessary capabilities. Immediate MDT reuses the existing Rel-8/9 measurement configuration and reporting procedures for Radio Resource Management (RRM) purposes.

The UE reports the measurements to the network every time a configured reporting condition is met. The measurements include:

- RSRP and RSRQ measurements, either provided periodically according to a configured reporting interval and number of reports, or provided with Event A2 (see Section 3.2.5.2) or on Radio Link Failure (RLF) (see Section 22.7);
- The UE's power headroom (see Sections 18.3.3 and 28.3.5.3);
- The uplink signal strength measured by the eNodeB.

In addition to these measurements, the overall report may additionally include location information similar to that for logged MDT, but it does not contain time information. RLF reports are stored by the UE even if it transitions to idle mode.

31.4 Machine-Type Communications

Machine-Type Communication (MTC) is a form of data communication which involves one or more entities that do not necessarily need human interaction. MTC devices are UEs equipped for MTC, which communicate through a Public Land Mobile Network (PLMN) with MTC servers and/or other MTC devices.

Applications for MTC are many and diverse. The following are some examples:

- MTC devices can be used by utility companies to read smart meters remotely at short but regular intervals.
- Automatic earthquake monitoring can provide vital public warnings.
- In the transport sector, MTC can be used for toll collection and road-pricing schemes, for monitoring the location of vehicles and for making automatic emergency calls in case of accidents.

- In the future, many consumer electronic devices are likely to incorporate MTC capability, allowing automatic remote maintenance as well as monitoring of heating, alarms and other applications.

An MTC device and its communication requirements are often different from normal human communications. Their traffic can typically be characterized as small, delay-tolerant data packets which are sent infrequently. Other considerations such as the fact that the traffic may originate predominantly from the device itself, low mobility, and requirements for low cost and low battery consumption may also be relevant depending on the type of application. The data packets are sometimes timer-controlled or they may be triggered by certain events such as earthquakes. There could potentially be a large number of MTC devices in a cell.

These characteristics of MTC devices motivate some special handling compared to human-generated traffic. One of the primary considerations addressed first in 3GPP is the potential for overload that could be caused by a large number of MTC devices. Since the MTC traffic can be timer-controlled or event-triggered, there is a possibility of a sudden surge in MTC traffic in a cell in which there is a large number of MTC devices. It is important to ensure that MTC traffic, which is delay-tolerant, does not disrupt (or at least causes only minimal additional delay for) human-generated traffic.

Release 10 support for MTC therefore focuses on overload protection for the Core Network (CN). This is achieved by exploiting the delay-tolerant nature of the MTC traffic by rejecting connection attempts from MTC devices as necessary. Specific CN nodes can provide overload indications to eNodeBs to request eNodeBs to prevent MTC access to these CN nodes. When an MTC device sends an RRC Connection Establishment request including a 'delay tolerant' cause code towards such a CN node, the eNodeB rejects the RRC connection request. In addition, the eNodeB provides an 'extended wait timer' in the RRC 'Connection Reject' message which prevents the MTC device from accessing the network again for delay-tolerant reasons for the duration of the timer. An example message flow for this type of connection rejection is shown in box A in Figure 31.8.

Depending on the type of access, identification of the CN nodes serving the MTC device is not always possible from the RRC 'Connection Request' message. For such cases, the CN node ID is provided in the RRC Connection Setup Complete message. Since it is then not possible to reject the connection request, the extended wait timer is also included in the RRC Connection Release message. This flow is shown in box B in Figure 31.8.

Additional overload-prevention mechanisms such as a longer periodic tracking area update timer are defined in the CN specifications [9].

While Release 10 does not support dynamic change of access type from MTC devices, some exceptions such as emergency calls are still allowed with higher priority.

The Release 10 support for MTC is expected to further optimized in Release 11 and beyond. Several other areas could be addressed. Mechanisms for handling RAN overload are expected to be included in Release 11 and are likely to include the possibility of barring MTC devices from even performing the RACH access for connection establishment during periods of overload, in order to further minimize impact on normal-priority traffic. Optimizing signalling overhead will also be a priority for Release 11, as the normal signalling connection establishment procedure that is required before a device can send any data causes significant overhead in the case of the small and infrequent packets sent by MTC devices. Other prioritized work areas will include handling device addressing for large numbers of

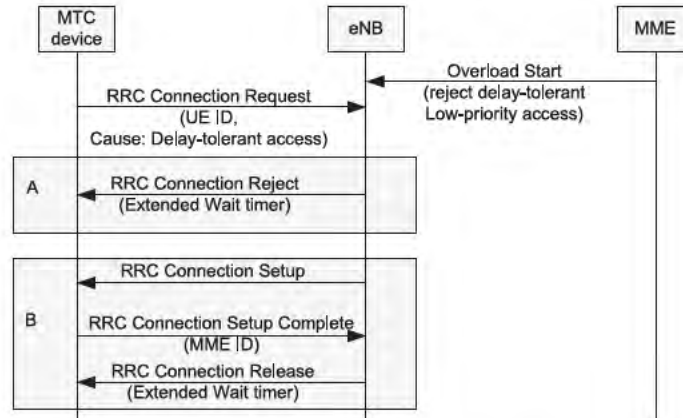


Figure 31.8: Mechanisms for handling CN overload.

devices without a Mobile Station International Subscriber Directory Number (MSISDN), power saving and optimizations for efficient handling of groups of similar MTC devices.

References¹⁵

- [1] 3GPP Technical Report 36.814, 'Evolved Universal Terrestrial Radio Access (E-UTRA); Further advancements for E-UTRA Physical Layer Aspects (Release 9)', www.3gpp.org.
- [2] Alcatel-Lucent and Alcatel-Lucent Shanghai Bell, 'Co-Channel Control Channel Performance for Hetnet', www.3gpp.org, 3GPP TSG RAN WG1, meeting 60bis, Beijing, China, April 2010.
- [3] Alcatel-Lucent, 'R1-093340: Blank Subframes for LTE', www.3gpp.org, 3GPP TSG RAN WG1, meeting 58, Shenzhen, China, August 2009.
- [4] NTT DOCOMO, 'R4-110224: Simulation results for e-ICIC RSRP/RSRQ measurements', www.3gpp.org, 3GPP TSG RAN WG4, meeting 57AH, Austin, USA, January 2011.
- [5] 3GPP Technical Specification 36.213, 'Evolved Universal Terrestrial Radio Access (E-UTRA); Physical Layer Procedures (Release 10)', www.3gpp.org.
- [6] 3GPP Technical Specification 32.422, 'Telecommunication management; Subscriber and equipment trace; Trace control and configuration management (Release 10)', www.3gpp.org.
- [7] 3GPP Technical Specification 37.320, 'Universal Terrestrial Radio Access (UTRA) and Evolved Universal Terrestrial Radio Access (E-UTRA); Radio measurement collection for Minimization of Drive Tests (MDT); Overall description; Stage 2 (Release 10)', www.3gpp.org.
- [8] Open Mobile Alliance, 'DiagMon Functions Supplemental Specification OMA-TS-DiagMon_Functions-V1', www.openmobilealliance.org.
- [9] 3GPP Technical Specification 24.301, 'Non-Access Stratum (NAS) Protocol for Evolved Packet System (EPS) (Release 10)', www.3gpp.org.

¹⁵All web sites confirmed 1st March 2011.

LTE-Advanced Performance and Future Developments

Takehiro Nakamura and Tetsushi Abe

LTE Release 8 incorporates the latest advances in radio technology and has attracted global industry support. Intensive field trials conducted by equipment vendors and network operators have confirmed the high performance of LTE, and commercial services using LTE were first launched in December 2009. Many leading operators worldwide have committed to deploying LTE, and the roll-out of LTE is expanding and accelerating.

LTE Release 9 brings enhancements and performance improvements in a number of aspects, including Home eNodeBs, MBMS, location services and MIMO. For Release 10 and beyond (LTE-Advanced), the standardization activities in 3GPP were aligned with the IMT-Advanced standardization process in ITU-R¹ to meet future demands for continuous improvement of system performance and satisfy the constantly increasing expectations and demands of the consumers. It was confirmed that LTE-Advanced fully satisfies the requirements set by ITU-R, and LTE-Advanced was therefore accepted as an IMT-Advanced technology. This chapter provides a summary of the system performance achievable with LTE-Advanced using the Release 10 version of the specifications, followed by a discussion of future developments beyond Release 10.

32.1 LTE-Advanced System Performance

As listed in Table 27.1, ITU-R specified eight items as the minimum requirements for IMT-Advanced:

¹International Telecommunication Union, Radio Communication Sector.

1. Peak spectral efficiency;
2. Cell spectral efficiency;
3. Cell-edge user spectral efficiency;
4. Bandwidth;
5. Latency;
6. Mobility;
7. Handover interruption time;
8. VoIP capacity.

Some of the results of the assessment of LTE-Advanced conducted by 3GPP against these criteria are presented below.

The peak spectral efficiencies for the uplink and downlink are shown in Table 32.1. Details of the assumptions, including overheads, can be found in [1]. The table shows that the ITU-R requirements can be satisfied by using 4-layer spatial multiplexing (based on LTE Release 8 Single-User MIMO (SU-MIMO)) for the downlink and 2-layer spatial multiplexing (based on Release 10 (LTE-Advanced) SU-MIMO) for the uplink.

The more stringent peak spectral efficiency requirements set by 3GPP for LTE-Advanced (see Table 27.2) can be satisfied using up to 8-layer spatial multiplexing in the downlink and up to 4-layer spatial multiplexing in the uplink, according to the Release 10 specifications for SU-MIMO.

Table 32.1: Peak spectral efficiency for LTE-Advanced Release 10.

	Downlink	Uplink
ITU-R requirements (bps/Hz)	15	6.75
LTE-Advanced peak spectral efficiency (bps/Hz)	16.3 (4 MIMO layers) 30.6 (8 MIMO layers)	8.4 (2 MIMO layers) 16.8 (4 MIMO layers)

The average cell spectral efficiencies and cell-edge user spectral efficiencies for the downlink and uplink in the ITU-R test environments are shown in Figures 32.1 and 32.2 respectively. The downlink performance was evaluated using LTE-Advanced Multi-User MIMO (MU-MIMO) and LTE Release 8 SU-MIMO, assuming four transmit and two receive antennas (4×2).

Figure 32.1 shows that in the downlink, LTE Release 8 SU-MIMO can satisfy the ITU-R requirements in the Indoor Hotspot (InH) and Rural Macrocell (RMa) deployment scenarios, but not in the Urban Microcell (UMi) and Urban Macrocell (UMa) scenarios.² On the other hand, LTE-Advanced MU-MIMO can satisfy the requirements in all the test environments and provides substantial gain over LTE Release 8 SU-MIMO.

The uplink performance was evaluated using LTE-Advanced SU-MIMO with two transmit and four receive antennas (2×4) and LTE Release 8 single-antenna transmission with four receive antennas (1×4).

²See Section 20.3.6.1 for details of the ITU-R test environments.




Figure 32.1: Downlink performance of LTE-Advanced in ITU-R deployment scenarios.




Figure 32.2: Uplink performance of LTE-Advanced in ITU-R deployment scenarios.

Figure 32.2 shows that in the uplink, LTE Release 8 single-antenna transmission satisfies the requirements in all of the environments. Additionally, LTE-Advanced SU-MIMO provides even further improved spectral efficiency.

In addition to these spectral efficiency evaluations, the evaluations conducted by 3GPP show that the other aspects of the IMT-Advanced requirements, namely in relation to bandwidth, latency, mobility, handover interruption time and the number of VoIP users accommodated, are fully satisfied by LTE Release 8 and LTE-Advanced [1].

32.2 Future Developments

Compared to previous mobile radio systems such as those based on HSPA, the achievable system performance is greatly improved by LTE Releases 8 and 9 followed by LTE-Advanced. Nevertheless, expectations from the market side continue to increase inexorably. In recent years, unprecedented market trends have been observed: the proliferation of high-specification terminals, especially smartphones, is bringing video delivery/streaming and other innovative applications within the reach of ever more consumers. In the future, mobile data traffic will grow at a pace that far outstrips that previously experienced. Meanwhile, it is becoming more challenging to achieve revenue growth as a result of the introduction of flat-rate data tariffs by many network operators around the world. Further reduction of the cost per bit will be a necessity in the future more than ever before in the history of mobile communication technology.

From another perspective, ecological concerns are steadily attracting more attention. Mobile communications needs innovative solutions to decrease the power consumed.

A further global concern is the so-called ‘digital divide’ – the regional differences in service availability, for example between rural and urban areas. In the future, comparable mobile services need to be provided for all areas in a cost-effective way.

Following these global trends, the main requirements for future releases of LTE can be broadly identified as follows:

1. **Increased capacity and spectral efficiency.** It is necessary to improve system capacity in order to support the tremendous growth of data traffic. In particular, further improvement of spectrum utilization efficiency is very important, since the available spectrum bands that are suitable for mobile communications are limited.
2. **Improvement of throughput experienced by the user.** For video delivery/streaming and new content-rich applications, the throughput actually experienced by the individual user is paramount and needs further improvement, not just the system capacity. In indoor scenarios, non-voice traffic is becoming dominant, and this will become even more obvious in the future. Further optimization of the system specifications for the indoor environment is therefore important.
3. **Fairness of throughput provision.** HSPA system design is primarily best-effort-based, which results in a throughput gap between cell-centre and cell-edge users. In both LTE and LTE-Advanced, potential solutions to reduce the throughput gap have been studied. To further improve fairness among users in future systems, improvements of cell-edge performance are required without sacrificing cell-centre performance. Fairness between cells and different geographical regions is also important from the perspective of the digital divide.

4. **Reduction of cost per bit.** Growth in data traffic is out-pacing growth in revenues. Besides, further network expansion and performance improvement in rural areas are key to addressing the digital divide.
5. **Energy saving.** Solutions to reduce energy consumption need to be considered at both system and device levels, in order to address both environmental and operational costs.

These requirements will need to be considered by future releases of the LTE specifications. Prime candidates for near-term system improvements of LTE in Release 11 include features that were partially studied but not specified within the timeframe of Release 10. In addition, further enhancements of features that were specified in Release 10 are potential aspects for consideration. Some of the main topics to be addressed in the short-term are therefore as follows:

- **Coordinated Multi-Point Transmission/Reception (CoMP).** As discussed in Section 29.5, some initial studies were performed into coordinated scheduling, coordinated beamforming and joint transmission techniques during the feasibility study for LTE-Advanced. Intra-eNodeB CoMP is already possible in LTE in a manner independent of standardization, but standardized inter-eNodeB coordination is not specified. Further study of the feasibility and performance gains of CoMP during the Release 11 timeframe is expected to result in a clear understanding of these aspects and of the potential use cases and deployment scenarios.
- **Downlink and uplink MIMO enhancement.** As MIMO deployments become more common, additional features for codebook and feedback enhancements, and higher-order MIMO are likely areas for improved performance to be achieved. MIMO performance in new deployment scenarios, including small cells, will also be studied with a view to identifying whether enhancements can deliver useful performance improvements in such conditions.
- **Carrier aggregation enhancements.** As spectrum availability changes for different network operators around the world, new combinations of adjacent and non-adjacent carriers and bands will continue to become relevant. RF requirements will be developed for such combinations and typically introduced in a release-independent manner so that user equipment of any release may support them; this will impose ever-increasing demands on equipment hardware capabilities and complexity. In addition, multiple timing advance controls are likely to be introduced for carrier aggregation cases where the end-to-end propagation delay on one aggregated carrier is significantly different from that on another, for example where one carrier makes use of repeaters. Control signalling improvements may also be considered to reduce overhead, and enhanced transmit diversity for the uplink HARQ acknowledgement feedback is expected to be introduced. Additional types of component carrier that are not backward-compatible (e.g. without cell-specific reference signals or synchronization signals) may also be considered.
- **Inter-Cell Interference Coordination (ICIC) enhancement.** ICIC has been supported since the beginning of LTE and was enhanced in Release 10 (see Section 31.2). In future releases, further improvements are expected. In particular, new features for the X2 interface may be further investigated. Current market conditions set high

expectations for low-cost dense networks and heterogeneous networks, including deployment of both femtocells and picocells. Such network deployments have the potential to provide attractive solutions to enhance system performance with little impact on the radio interface and terminals, and enhanced ICIC will play an important role in achieving this.

- **Self-Optimizing Networks (SON) and Minimization of Drive Tests (MDT).** SON was initially specified in Release 8. It enables the cost of network deployment and maintenance to be reduced, especially for dense networks. Further enhancements are to be investigated in future releases as reduction of the cost of network optimization is an increasingly important part of minimizing the cost per bit. Reducing the number of drive tests that have to be conducted is another effective way to enable operators to reduce the network cost; some MDT features are included in the Release 10 specifications, and more will be investigated for further practical use cases, with consideration being paid to the tradeoffs between benefits and network complexity.
- **Energy saving.** Energy saving is important for base stations in particular. Reduction of power consumption of equipment is essentially an implementation matter, but the possibility of designing efficient implementations should be taken into account when specifying any new features for the radio interface.
- **Machine-Type Communications.** As discussed in Section 31.4, Machine-Type Communications (MTC) are proliferating in cellular networks for numerous applications. Such communications have certain particular characteristics, typically including low mobility, low frequency of use, a need for long battery life, and potentially a very large number of terminals. Future releases of LTE may contain specific radio access network features to provide optimized support for communications with these characteristics, if such features are identified.

Besides the above aspects, further improvement of features already included in Releases 8, 9 and 10, such as MBMS, UE location technologies and relaying (for example relay backhaul enhancements or improved support for mobile relays on high-speed trains) are likely to be considered for future releases of LTE.

Market requirements will continue to evolve; thus mid- and long-term system improvements of LTE will need to keep pace with market growth. For long-term improvements, the most useful developments will be brand-new features that deliver very substantial benefits. Potential candidates may include features related to new multiple access schemes and advanced multi-antenna solutions for the physical layer, features for new architectures and network deployment scenarios, and new performance specifications for the mobile terminal with more advanced receivers.

References³

- [1] 3GPP Technical Report 36.912, 'Feasibility Study for Further Advancements for E-UTRA (LTE-Advanced) (Release 9)', www.3gpp.org.

³All web sites confirmed 1st March 2011.

Index

- 16QAM, 193
- 1xRTT
 - pilot strength, 516
- 3rd Generation Partnership Project (3GPP),
 - 3–6
 - channel model, 439–440
- 3rd Generation Partnership Project 2 (3GPP2),
 - 4
- 3rd-order intercept point, 490
- 3rd-order intermodulation products, 471
- 328-bit transport block size, 347
- 5th-order intermodulation products, 471
- 64QAM, 193, 194
- 802.16, 4
- 802.16e, 547, 549

- ABS, *see* Almost Blank Subframe
- absolute TPC commands, 416
- AC, *see* Access Class
- access barring, 83
- Access Class (AC), 83
- Access Control List (ACL), 586
- access link, 674
- access network, 26, 30–31
- Access Point Name (APN), 28
- access security management entity, 63
- Access Stratum (AS), 29, 30, 33, 57
 - base-key, 63, 93
 - derived keys, 63, 92
- accessibility
 - verification, 83
- accumulative TPC commands, 416
- ACIR, *see* Adjacent Channel Interference Ratio
- ACK, 243, 351–363
- ACK_SN, 107
- ACK/NACK, 364–367
- ACK/NACK bundling, 549–551
- ACK/NACK multiplexing, 549–551
- ACK/NACK repetition, 367
 - carrier aggregation, 640
- ACK/NACK transmission, 200, 243, 352–363
 - in TDD operation, 548–551
 - PUSCH, 667
- Acknowledged Mode (AM), 34, 93, 98,
 - 101–105
 - retransmission, 101
- Acknowledged Mode Data (AMD)
 - PDU format, 106
 - PDU Segment, 106–107
- ACL, *see* Access Control List
- ACLR, *see* Adjacent Channel Leakage Ratio
- ACS, *see* Adjacent Channel Selectivity
- adaptive HARQ, 109, 368
- Adaptive Modulation and Coding (AMC),
 - 217–218, 280, 478
- Additional MPR (A-MPR), 465
- additional SEMs, 469
- adjacent channel interference, 493, 533–544
- Adjacent Channel Interference Ratio (ACIR),
 - 484, 535–537, 539, 540, 543
- Adjacent Channel Leakage Ratio (ACLR),
 - 468–470, 535, 544
 - base station, 575, 576
 - Home eNodeB, 575, 576
- Adjacent Channel Selectivity (ACS),
 - 482–485, 535, 544
 - base station, 578
 - Home eNodeB, 578
- ADSL, *see* Asymmetric Digital Subscriber Line
- advanced receiver, 602
- Advanced TDMA Mobile Access (ATDMA),
 - 438
- aggregation level, 202

- A-GNSS, *see* Assisted Global Navigation Satellite System
- Alamouti space-time code, 257, 265
- Alliance for Telecommunications Industry Solutions (ATIS), 5
- Allocation and Retention Priority (ARP), 34, 35
- Almost Blank Subframe (ABS), 703–710
 - bitmap, 705
 - configuration, 709–710
 - femtocell, 571, 573
 - invoke indication, 705
 - measurement subset, 705
 - pattern, 705
 - status, 705
- Almost Regular Permutation (ARP), 236
- AMC, *see* Adaptive Modulation and Coding
- AMD PDU, *see* Acknowledged Mode Data PDU
- amplitude imbalance, 496–497
- A-MPR, *see* Additional MPR
- AMPS, *see* Analogue Mobile Phone System
- Analogue Mobile Phone System (AMPS), 1
- analogue to digital converter, 479, 494
- Angle of Arrival (AoA), 424, 432–433, 435
- Angular Spread (AS), 442–444
- ANRF, *see* Automatic Neighbour Relation Function
- antenna
 - array, 663
 - configuration, 447
 - constellations, 663
 - correlation, 663
 - coverage, 674
 - cross-polarized, 446, 663
 - donor, 674
 - gain imbalance, 493
 - number of, 190–192
 - port, 145, 167–168
 - selection, 206, 367
 - PUSCH, 367–368
 - SRS, 368
 - selection indication, 367–368
 - selection mask, 368
 - tilting, 606
 - virtualization, 659
- AoA, *see* Angle of Arrival
- aperiodic CSI
 - reporting, 219–221
 - Release 10, 664
 - triggering, 204, 208
 - carrier aggregation, 641–642
- APN, *see* Access Point Name
- a posteriori bit probability, 225
- ARIB, *see* Association of Radio Industries and Businesses
- ARP, *see* Allocation and Retention Priority
- ARP, *see* Almost Regular Permutation
- ARQ, *see* Automatic Repeat reQuest
- array gain, 16, 251
- assignment
 - downlink, 203–211
- Assisted Global Navigation Satellite System (A-GNSS), 424–426
 - UE-assisted, 426
- Association of Radio Industries and Businesses (ARIB), 5
- Asymmetric Digital Subscriber Line (ADSL), 5, 124
- asynchronous HARQ, 109, 241
- ATIS, *see* Alliance for Telecommunications Industry Solutions
- authentication, 30
- Authentication Centre (AuC), 28, 63
- autocorrelation, 155–162, 328, 555
- Automatic Neighbour Relation (ANR), 509
 - Function (ANRF), 50, 582–584
 - inter-frequency, 583–584
 - inter-RAT, 583–584
 - intra-LTE, 582–583
 - neighbour detection function, 583
 - neighbour removal function, 583
 - NRT management function, 583
- available composite load, 588–589
- backhaul link, 674
 - HARQ, 698–699
 - physical channels, 691–696
 - radio link failure, 690
 - random access, 690
 - reference signals, 691–692
 - subframe configuration, 696–697
- backoff indicator, 374
- backward compatibility, 616, 618–619, 624
- backward handover, 69
- bands, 459–460
- bandwidth
 - channel, 13, 140, 323, 350, 458, 460–461

- efficiency, 324
 - parts, 222
 - system, 13, 140, 190, 197, 323, 350, 458, 461
 - transmission, 460–461
- base RS sequence, 329–330, 363
- base station
 - ACLR, 575, 576
 - ACS, 578
 - blocking, 578
 - demodulation performance, 578
 - dynamic range, 578
 - frequency error, 575, 576
 - maximum output power, 464–465, 575
 - Noise Figure (NF), 576
 - reference sensitivity, 576, 577
 - spectrum emission mask, 576
 - spurious emissions, 575, 576
- Base station Colour Code (BCC), 511
- Base Station Identity Code (BSIC), 511
 - initial identification, 511, 512
 - re-confirmation, 511, 512
 - verification, 512
- base-station to base-station interference, 536–539
- baseband equivalent filter, 493, 494
- BCC, *see* Base station Colour Code
- BCJR algorithm, 227
- beamforming, 171, 254–255, 270–272, 558, 561
 - coordinated, 669, 670
 - dual-layer, 171–174, 270–272
 - opportunistic, 669
 - reciprocity-aided, 658
- bearer, 25, 26
 - Bearer Management, 43
 - dedicated, 37
 - default, 37, 73
 - EPS, 34–38
 - establishment, 34, 37–38
 - Guaranteed Bit Rate (GBR), 28, 34, 35
 - ID, 35
 - management, 43–44
 - modification, 34
 - non-Guaranteed Bit Rate, 34
 - re-establishment, 29
 - Setup, 43
 - Setup Request, 37, 44
 - Setup Response, 44
- Belief Propagation (BP), 231
- Bessel function, 134, 176
- bias
 - handover offset, 702
- Binary Phase Shift Keying (BPSK), 200, 226, 474
- BLER, *see* BLock Error Rate
- blind decoding, 192, 211–212
 - carrier aggregation, 633
- blind handover, 69
- block code, 245
- BBlock Error Rate (BLER), 217
- blocking specifications, 480–487
 - base station, 578
 - Home eNodeB, 578
- BM-SC, *see* Broadcast/Multicast Service Centre
- BPSK, *see* Binary Phase Shift Keying
- broadband, 128, 142
- broadcast, 11, 140, 279, 293–314
- Broadcast CHannel (BCH), 111, 189
- Broadcast Control CHannel (BCCH), 59, 110, 312, 511
- Broadcast/Multicast Service Centre (BM-SC), 295, 305, 309
- BSIC, *see* Base Station Identity Code
- BSR, *see* Buffer Status Report
- buffer, 279
 - circular, 238
 - HARQ, 240, 242–243
 - carrier aggregation, 648
 - status, 287
- Buffer Status Report (BSR), 111, 114
 - carrier aggregation, 631
 - long, 114
 - padding, 114
 - periodic, 114
 - regular, 114, 115
 - short, 114
- bundling
 - ACK/NACK, 549–551
 - spatial, 550, 638, 639
 - subframe, 323, 347
 - TTI, 323, 347
- Bussgang's theorem, 494
- byte alignment, 88
- C-RNTI, *see* Cell Radio Network Temporary Identifier, 112
- capacity, 283
 - delay-limited, 283, 285

- ergodic, 283–285
- MIMO, 253
- Shannon, 283
- system, 599–608
- VoIP, 608, 616, 617
- zero-outage, 285
- carrier activation
 - activation of SCell, 630
 - deactivation of SCell, 630
- carrier aggregation, 618, 620, 623–650, 719
 - ACK/NACK, 637–640
 - band combinations, 649–650
 - Buffer Status Report (BSR), 631
 - control signalling, 631–642
 - cross-carrier scheduling, 624, 632–635
 - CSI feedback, 640–642
 - Discontinuous Reception (DRX), 628
 - HARQ, 628
 - logical channel prioritization, 631
 - measurements, 625–627
 - mobility, 625–627
 - PCell, 624
 - change, 625
 - PHICH, 636
 - power control, 642–644
 - PUCCH, 636–641
 - radio link failure, 627
 - random access procedure, 628
 - RRC connection establishment, 624
 - SCell, 624
 - addition, 625
 - scenarios, 649–650
 - scheduling, 628
 - Semi-Persistent Scheduling (SPS), 628
 - serving cell, 624
 - SRS, 642
 - timing advance, 642
 - uplink power control, 642–644
 - user plane, 628–631
- carrier frequency, 459
- Carrier Frequency Offset (CFO), 133–135
- Carrier Indicator Field (CIF), 632–633
- Case 1 deployment scenario, 599, 603–608
- Case 3 deployment scenario, 599, 603–608
- categories of user equipment, 18
- cavity resonator filters, 539
- CAZAC, *see* Constant Amplitude Zero AutoCorrelation
- CB, *see* Circular Buffer
- CC, *see* Component Carrier
- CCCH, *see* Common Control Channel
- CCE, *see* Control Channel Element
- CCO, *see* Cell Change Order
- CCSA, *see* China Communications Standards Association
- CDD, *see* Cyclic Delay Diversity
- CDL, *see* Clustered Delay Line
- CDM, *see* Code Division Multiplexing
- CDMA, *see* Code Division Multiple Access
- CDMA2000, 4, 13
 - 1xEV, 238
 - 1xRTT pilot strength, 516
 - HRPD pilot strength, 516
 - inter-working, 28, 38
 - mobility, 74
- cell, 1
 - acceptable, 79
 - access restriction, 83
 - black-list, 68
 - capacity class, 589
 - Closed Subscriber Group (CSG), 84
 - femto, 563–580
 - hybrid, 84
 - ID, 152, 155, 332
 - enhanced, 424, 431–434
 - positioning, 424, 430–433
 - identification delay, 505, 507, 508
 - identification procedure, 152, 155
 - listed, 76
 - pico, 563–580
 - priority, 69
 - range, 10–11, 409
 - ranking, 82–83
 - reselection, 57, 68–72, 79–84, 513–516, 518
 - criteria, 520, 521
 - inter-frequency, 520–521
 - inter-RAT, 521–522
 - intra-frequency, 519–520
 - paging interruption, 522
 - search, 151–161, 505–512, 524
 - E-UTRAN, 505–508
 - GSM, 511–512
 - inter-frequency, 507–508
 - intra-frequency, 505–507
 - UTRAN, 510–511
 - selection, 57, 68–72, 78–84, 518
 - any cell selection, 84
 - criterion, 519
 - serving, 624

- suitable, 79
 - throughput, 10
 - white-list, 68, 76
- Cell Change Order (CCO), 73
- Cell Radio Network Temporary Identifier (C-RNTI), 72, 112, 374
- Cell Range Expansion (CRE), 702–703
- cell-specific Reference Signals (RS), 527
- cell-specific Reference Signals (RSs), 167–171, 193
- central limit theorem, 131
- CFI, *see* Control Format Indicator
- channel
 - bandwidth, 13, 140, 323, 350, 457, 458, 460–461
 - coding, 215–248
 - control channels, 244–245
 - coherence, 169, 218, 555
 - delay spread, 125, 136, 139, 140, 142, 175, 300, 440, 441, 444
 - estimation, 165–187
 - frequency-domain, 178–181
 - IFFT, 178–179
 - Least Squares (LS), 179
 - linear, 179–180
 - MMSE, 179–181
 - spatial-domain, 184–185
 - time-domain, 181–184
 - Impulse Response (CIR), 136, 139, 165
 - model, 438–454
 - 3GPP, 439–440
 - correlation matrix based, 442
 - geometry-based, 442
 - implementation, 454–455
 - ITU, 439
 - primary, 451
 - module
 - extension, 451
 - sounding, 658
 - spacing, 485
- Channel Quality Indicator (CQI), 215–223, 272, 286, 348, 352–356, 364–366, 662
- channel selection, 637
- Channel State Information (CSI), 261, 281, 283, 352, 362–363, 554, 558–561, 640–642, 662–665
 - multicell, 670
 - Reference Signal (CSI-RS), 653–657
 - Release 10, 662–665
 - restricted measurement, 708–710
- Channel Transfer Function (CTF), 178
- channel vector quantization, 274, 663
- Chase combining, 239
- China Communications Standards Association (CCSA), 5
- CIF, *see* Carrier Indicator Field
- ciphering, 89, 91–93
 - NULL algorithm, 64, 73
- Circuit-Switched FallBack (CSFB), 39, 68, 512
- Circuit Switching (CS), 25
- Circular Buffer (CB), 238–240
- Circular Viterbi Algorithm (CVA), 224
- clipping, 131
- closed access, 565
- closed-loop
 - antenna selection transmit diversity, 367–368, 560
 - power control, 415–417
 - rank 1 precoding, 270
 - spatial multiplexing, 193, 267
 - switched antenna diversity, 367–368, 560
 - timing advance, 408–411
 - uplink precoding, 666
- Closed Subscriber Group (CSG), 19, 84, 563, 565, 585, 669
 - handover, 72
 - Identifier (ID), 565
 - manual selection, 84
 - member, 566
 - subscription list, 565
- Clustered Delay Line (CDL), 444
- clustered DFT-S-OFDM, 645–647
- CM, *see* Cubic Metric
- CMAS, *see* Commercial Mobile Alert Service
- CN, *see* Core Network
- co-channel interference, 493
- Code Division Multiple Access (CDMA), 3, 251, 317, 407, 412
- Code Division Multiplexing (CDM) PUCCH, 351
- Code Division Testbed (CODIT), 438
- code rate, 216, 219
 - greater than one, 347
- codebook, 263, 266–268
 - constant modulus, 267
 - DFT, 267, 274–275
 - Householder, 268

- nested, 267
 - Release 8, 266–268
 - Release 10, 663–664
 - two-stage, 664
 - uplink precoding, 666
- codebook subset restriction, 273
- Coded OFDM (COFDM), 124
- codeword, 266, 267
 - PCFICH, 245
- codeword-to-layer mapping, 660
- coding, 215–248
- coherence
 - bandwidth, 169
 - time, 218, 555
- coherent demodulation, 161–163, 165, 327
- coherent detection, 161–163, 165, 327
- collision detection, 402
- Commercial Mobile Alert Service (CMAS), 311–312
- Commercial Mobile Alert System (CMAS), 48
- Common Control CHannel (CCCH), 59, 110
 - integrity protection, 59
- common phase rotation, 498
- Common Pilot CHannel (CPICH), 510
 - E_c/N_0 , 515
 - RSCP, 515
- Common Reference Signals (CRSs), 167
- common search space, 211, 212
- Common Subframe Allocation (CSA), 307–308
- CoMP, *see* Coordinated MultiPoint
- Component Carrier (CC), 618, 623–650
 - activation, 630
 - backward compatibility, 618, 624
 - bandwidth, 623
 - deactivation, 630
 - Primary (PCC), 624
 - Secondary (SCC), 624
- compression point, 477
- concatenation, 99–100
- conformance testing, 454
- connection
 - control, 62–73
 - establishment, 63–67
 - release, 64–67
- Constant Amplitude Zero AutoCorrelation (CAZAC), 155
- constant modulus, 267
- constellation diagrams, 194
- contention-based
 - random access, 372–376, 394, 525
- contention-free, 233
 - decoding, 234–238
 - interleaving, 234–238
 - random access, 372, 376, 525
- contention resolution, 375–376
- context
 - information, 58
 - UE, 29, 31
- contiguous resource block allocations, 209
- continuous-wave, 482
- control channel, 196–212, 348–364
 - region, 198
 - uplink, 348–352
- Control Channel Element (CCE), 202, 203, 211
 - aggregation level, 202
 - relation to PUCCH resource, 360
- Control Format Indicator (CFI), 198
- control logical channel, 110
- Control PDU Type (CPT), 107
- control plane, 33–34, 40–42, 49, 56–86
 - latency, 616, 617
- control signalling overhead, 602–603
- controller entity, 109
- convolution, 127, 131, 137
 - aperiodic, 127
 - circular, 127, 131, 137
 - linear, 131, 137
- convolutional codes, 190, 216, 223–225, 244
- cooperative MIMO, 668–671
- coordinated beamforming, 669, 670
- Coordinated MultiPoint (CoMP), 619, 668–671, 719
 - coordinated beamforming, 670
 - coordinated scheduling, 670
 - downlink, 669–671
 - feedback, 670
 - Joint Transmission (JT), 670
 - uplink, 669
- coordinated scheduling, 669, 670
- co-phasing, 663
- Core Network (CN), 25–30, 295–296
- correlation matrix based
 - channel model, 442
- COST, 438
- COUNT, 92, 93, 97
- counting procedure, 295, 310–311
- coupling, 536

- covariance matrix, 178–180
- coverage, 608–610
 - hole detection, 591–592
- coverage antenna, 674
- CP, *see* Cyclic Prefix
- CPICH, *see* Common Pilot CHannel
- CQI, *see* Channel Quality Indicator
- CRC, *see* Cyclic Redundancy Check
- CRE, *see* Cell Range Expansion
- Create Dedicated Bearer Request, 37
- cross-carrier scheduling, 618, 624, 632–635
- cross-correlation, 156, 157, 160, 328
- cross-polarized antenna, 446, 663
- CRS, *see* Common Reference Signal
- CS, *see* Circuit Switching
- CSA, *see* Common Subframe Allocation
- CSFB, *see* Circuit-Switched FallBack
- CSG, *see* Closed Subscriber Group
- CSI, *see* Channel State Information
- Cubic Metric (CM), 317, 322, 396, 473
 - clustered DFT-S-OFDM, 645–646
 - group, 395
- cut-set bound, 680
- Cyclic Delay Diversity (CDD), 264, 268–270
- Cyclic Prefix (CP), 152, 165, 298, 300, 301, 345
 - EVM computation, 464
 - extended, 140
 - normal, 140, 141
 - OFDM downlink, 127, 131, 135–137, 140, 142
 - SC-FDMA uplink, 323
- Cyclic Redundancy Check (CRC)
 - code block, 243
 - CSI on PUSCH, 366
 - data channels, 243–244
 - PBCH, 192
 - PDCCH, 206
 - transport block, 243, 244
- cyclic shift, 330–335, 337–339, 351, 368, 657
 - hopping, 333–335
 - remapping, 362
- D/C field, 97, 106
- DAB, *see* Digital Audio Broadcasting
- DAI, *see* Downlink Assignment Index
- data-associated control signalling, 348
- data non-associated control signalling, 348
- Data Radio Bearer (DRB), 58, 67, 73
 - establishment, 67–68
- d.c., *see* direct current
- DCCH, *see* Dedicated Control Channel
- DCI, *see* Downlink Control Information
- DCS1800, 487
- deciphering, 89
- Dedicated Control CHannel (DCCH), 59, 110, 516
 - ciphering, 59
 - integrity protection, 59
- Dedicated Physical Control CHannel (DPCCCH), 524
- Dedicated Traffic CHannel (DTCH), 67, 110
- delay-limited capacity, 283, 285
- delay spread, 125, 136, 139, 140, 142, 175, 300, 440, 441, 444
- delay tolerant, 713
- delay-temporal separability, 176
- delta signalling, 73
- Demodulation Reference Signal (DM-RS)
 - clustered DFT-S-OFDM, 647
 - downlink, 171–174, 652–653
 - uplink, 327, 335–337, 343, 657–658
- demultiplexing
 - entity, 109
- deployment scenario, 599
 - Indoor Hotspot (InH), 450
 - ITU, 616
 - Rural Macro (RMa), 450
 - Urban Macro (UMa), 450
 - Urban Micro (UMi), 450
- de-rating, 471
- desensitization, 539
- de-spreading, 361
- detection
 - coherent, 161–163
 - non-coherent, 163
- DFT, *see* Discrete Fourier Transform
- DFT-Spread OFDM (DFT-S-OFDM), 320–321, 344
 - clustered, 645–647
- Differentiated Services (Diffserv), 43
- Digital Audio Broadcasting (DAB), 15, 124, 313
- Digital Mobile Broadcasting (DMB), 15, 313, 314
- digital to analogue converter, 492
- Digital Video Broadcasting (DVB), 15, 124, 313–314
 - Handheld (DVB-H), 313–314
 - S2 (DVB-S2), 232

- Terrestrial (DVB-T), 313
- Direct Channel Feed-Back (DCFB), 561
- direct conversion (zero IF) architectures, 324
- direct current (d.c.)
 - offset, 493, 494
 - subcarrier
 - OFDM, 142
 - SC-FDMA, 324
- direct link, 674
- Direction Of Arrival (DOA), 271
- Discontinuous Reception (DRX), 84, 115–117
 - carrier aggregation, 628
 - cycle, 58, 115, 503, 506–508, 511, 513, 515, 516, 519
 - on duration, 115
 - period, 115
- Discontinuous Transmission (DTX), 355
- Discrete Fourier Transform (DFT), 124, 127, 141, 268, 344
 - codebook, 267, 274–275
- dispersive channel, 134
- distributed
 - mapping, 194
 - transmission, 321
- diversity, 250, 257–258
 - gain, 16, 251
 - multi-user, 282, 284, 287
 - receive, 474
 - transmit, 264–274
 - uplink, 667–668
- DMB, *see* Digital Mobile Broadcasting
- DM-RS, *see* Demodulation Reference Signal
- donor antenna, 674
- donor eNodeB, 674
- Doppler, 140
 - frequency, 134, 139, 140, 175
 - shift, 134, 169, 301
 - spectral density, 134
 - spectrum, 134, 441
 - spread, 134, 135, 169, 301
- downlink, 121–314
 - assignment, 112, 203–211
- Downlink Assignment Index (DAI), 550
- Downlink Control Information (DCI), 202–207, 367
 - Format 0, 204, 659
 - Format 1, 204
 - Format 1A, 204
 - Format 1B, 204
 - Format 1C, 205, 306
 - Format 1D, 205
 - Format 2, 205
 - Format 2A, 205
 - Format 2B, 205
 - Format 2C, 206, 662
 - Format 3, 206
 - Format 3A, 206
 - Format 4, 206, 659
 - Formats, 212
- Downlink Pilot TimeSlot (DwPTS), 547
- Downlink Shared CHannel (DL-SCH), 111, 149
- down-sampling, 494
- DPCCH, *see* Dedicated Physical Control Channel
- DRB, *see* Data Radio Bearer
- drive test, 710–712, 720
- drop, 442
- DRX inactivity timer, 116
- DRX Retransmission Timer, 116
- DRX Short Cycle Timer, 116
- DTCH, *see* Dedicated Traffic Channel
- DTX, *see* Discontinuous Transmission
- dual-layer beamforming, 171–174, 270–272
- duplex modes, 12–13, 532–533
 - discovery, 152
 - relaying, 674
- duplex spacing, 533
- duplexer, 460, 532, 533
- duplicate detection, 100–101, 104
- DVB, *see* Digital Video Broadcasting
- DwPTS, *see* Downlink Pilot TimeSlot
- Dynamic Offset, 415
- dynamic range, 491–492
 - base station, 578
 - Home eNodeB, 578
- dynamic scheduling, 112
- E911 requirements, 424
- Earthquake and Tsunami Warning System (ETWS), 48, 311–312
 - primary notification, 312
 - secondary notification, 312
- ECGI, *see* E-UTRAN Cell Global Identifier
- ECM, *see* EPS Connection Management
- ECM-CONNECTED, 29, 65
- ECM-IDLE, 28, 29, 65
- EDGE, 3, 13

- eICIC, *see* enhanced Inter-Cell Interference Coordination
- Electronic Programme Guide (EPG), 310
- emergency call, 30, 64, 79, 423
- EMM, *see* EPS Mobility Management
- EMM-Deregistered, 65, 78
- EMM-Registered, 65, 78
- Encapsulating Security Payload (ESP), 91
- energy saving, 596–597, 720
- enhanced cell-ID, 424, 431–434
- enhanced Inter-Cell Interference Coordination (eICIC), 701–710
- eNodeB, 13, 26, 30, 31
 - ACLR, 575, 576
 - ACS, 578
 - blocking, 578
 - class
 - Home eNodeB, 464
 - Local Area, 464
 - Wide Area, 464
 - Configuration Transfer, 586
 - Configuration Update, 585, 587
 - demodulation performance, 578
 - dynamic range, 578
 - frequency error, 575, 576
 - maximum output power, 575
 - Noise Figure (NF), 576
 - receiver requirements, 576–579
 - reference sensitivity, 576, 577
 - self-configuration, 584–587
 - spectrum emission mask, 576
 - spurious emissions, 575, 576
 - transmitter requirements, 575–576
- eNodeB-configured sub-band, 221
- entry condition, 76
- EPA, *see* Extended Pedestrian A
- EPC, *see* Evolved Packet Core
- EPDU, *see* External Protocol Data Unit
- EPG, *see* Electronic Programme Guide
- ePHR, *see* extended Power Headroom Report
- EPS, *see* Evolved Packet System
- equalization, 125, 128, 131, 300, 463
- E-RAB, *see* Evolved-UTRAN Radio Access Bearer
 - Bearer
- ergodic capacity, 283–285
- Error Vector Magnitude (EVM), 462–464
- E-SMLC, *see* Evolved Serving Mobile Location Centre
- ESP, *see* Encapsulating Security Payload
- estimation
 - IFFT, 178–179
 - Least Squares (LS), 179
 - linear, 179–180
 - MMSE, 179–181
- ETSI, *see* European Telecommunications Standards Institute
- ETU, *see* Extended Typical Urban
- ETWS, *see* Earthquake and Tsunami Warning System
- European Telecommunications Standards Institute (ETSI), 2, 5
- E-UTRAN Cell Global Identifier (ECGI), 509, 582
 - reporting, 508–510
 - inter-frequency, 510
 - intra-frequency, 509–510
- EVA, *see* Extended Vehicular A
- EV-DO, 4
- event
 - A1, 76, 517, 625
 - A2, 76, 517, 625
 - minimization of drive tests, 712
 - A3, 76, 517, 626
 - A4, 76, 517
 - A5, 76, 517, 626
 - A6, 626
 - B1, 76
 - B2, 76, 626
- EVM, *see* Error Vector Magnitude
- Evolved Packet Core (EPC), 3, 25–30
- Evolved Packet System (EPS), 3, 25, 26, 38
 - bearer, 26, 34–38, 67
 - establishment, 78
 - release, 78
 - Connection Management (ECM), 65
 - Mobility Management (EMM), 65
- Evolved Serving Mobile Location Centre (E-SMLC), 27, 29, 423, 426, 433–435
- Evolved-UTRAN (E-UTRAN), 26, 30–31
 - Radio Access Bearer (E-RAB), 35
- excess delay, 438
- extended CP, 140
- extended ITU
 - channel model, 440–441, 447–449
- Extended Pedestrian A (EPA), 440–442
- extended Power Headroom Report (ePHR), 644
- Extended Typical Urban (ETU), 440–442

- Extended Vehicular A (EVA), 440–442
- extended wait timer, 713
- extension channel module, 451
- External Protocol Data Unit (EPDU), 434
- extrinsic information, 235

- fading, 251
 - multipath, 439
 - shadow, 439
- fair scheduling, 281
- fairness, 282, 285
- Fast Fourier Transform (FFT), 124, 127, 130, 131, 141, 142, 154, 268
- FCC, *see* Federal Communications Commission
- FDMA, *see* Frequency Division Multiple Access
- Feature Group Indicator (FGI), 18
- Federal Communications Commission (FCC), 423
- feedback, 215–223, 272–275, 558, 559
 - CoMP, 670
 - CQI, 215–223, 662
 - CSI, 640–642
 - Release 10, 662–665
 - explicit, 663
 - implicit, 274, 662–663
 - multicell, 670
 - PMI, 272–274, 662–665
 - RI, 272–274, 662, 665
 - sub-band, 665
 - wideband, 665
- feedforward, 558, 559
- femtocell, 563–580
 - access control, 566, 567
 - closed access, 565
 - handover
 - inbound, 566–568
 - outbound, 566–568
 - hybrid access, 566
 - interference, 569–574
 - membership verification, 566, 567
 - mobility, 566–568
 - paging, 566
 - RF requirements, 574–579
 - sensing, 569
- FFT, *see* Fast Fourier Transform
- FGI, *see* Feature Group Indicator
- FI, *see* Framing Info
- Fibre-To-The-Home (FTTH), 5
 - fifth-order intermodulation products, 471
- File Transfer Protocol (FTP), 26
- Finite Impulse Response (FIR), 182
- FIR, *see* Finite Impulse Response
- First Generation system, 1
- First Missing SDU (FMS), 98
- flash-light effect, 272
- flat channel, 134
- forward handover, 69
- forward security, 63
- forward-backward, 227
- Fourier transform, 130
- fractional path-loss compensation, 414
- fractional power control, 413–414
- Frame Structure Type 1, 147
- Frame Structure Type 2, 147
- Framing Info (FI), 106
- frequency
 - accuracy, 409
 - bands, 459–460, 620
 - diverse scheduling, 346–347
 - diversity, 194
 - domain channel estimation, 178–181
 - drifts, 133
 - error, 141, 409, 493, 494
 - base station, 575, 576
 - Home eNodeB, 575, 576
 - hopping, 346–347
 - offset, 135, 140
 - reuse, 287–291, 600
 - CSI-RS, 655
 - partial, 289, 669
 - soft, 289
 - selective channel, 123, 125, 128
 - selective scheduling, 345–346
 - semi-selective scheduling, 337
 - shift, 171
 - synchronization, 127
- Frequency Division Duplex (FDD), 4, 147, 281, 505, 532, 535, 536, 539, 543, 545, 554
- Frequency Division Multiple Access (FDMA), 3, 137, 251
- Frequency Switched Transmit Diversity (FSTD), 264–265, 659
- FSTD, *see* Frequency Switched Transmit Diversity
- FTTH, *see* Fibre-To-The-Home
- full buffer traffic, 286
- full-duplex, 532

- relay, 674
- Galileo, 425
- Gateway Mobile Location Centre (GMLC), 28
- Gaussian Minimum Shift Keying (GMSK), 485
- GBR, *see* Guaranteed Bit Rate
- GCL, *see* Generalized Chirp-Like sequences
- General Packet Radio Service (GPRS), 3
 - Tunnelling Protocol (GTP), 32, 42, 297, 586
 - extension header, 53
 - User plane (GTP-U), 42
- Generalized Chirp-Like (GCL) sequences, 155
- generator sequences, 224
- geometry-based channel models, 442
- GERAN, 6, 293
- Global Navigation Satellite System (GNSS), 425
- Global Positioning System (GPS), 298, 423, 425
- Global System for Mobile Communications (GSM), 1–3, 5, 13, 533
 - cell range, 408
 - coherent demodulation, 166
 - GSM850, 487
 - inter-working, 38
- GLOBAL'naya NAVigatsionnaya Sputnikovaya Sistema (GLONASS), 425
- GLONASS, *see* GLOBAL'naya NAVigatsionnaya Sputnikovaya Sistema
- GMLC, *see* Gateway Mobile Location Centre
- GMSK, *see* Gaussian Minimum-Shift Keying
- GNSS, *see* Global Navigation Satellite System
- Gold code, 148, 170, 334, 344
- GPRS, *see* General Packet Radio Service
- GPS, *see* Global Positioning System
- grid of beams, 663
- group delay distortion, 498
- GSM, *see* Global System for Mobile Communications
- GTP, *see* General Packet Radio Service Tunnelling Protocol
- Guaranteed Bit Rate (GBR), 28, 34, 35
- guard band, 123, 487
- guard period, 147, 544, 546, 547
- Half-Duplex, 532
 - FDD (HD-FDD), 533, 544, 545, 551–552
 - relay, 674
- handover, 32–33, 63, 68–72, 89, 93–95, 513, 515, 516, 518, 522
 - asynchronous, 72
 - backward, 69
 - bias, 702
 - blind, 69, 522, 524
 - cancel, 53
 - command, 71, 74
 - failure, 73
 - forward, 69
 - Handover Command, 45
 - Handover Notify, 45
 - Handover Request, 52
 - Handover Request ACK, 52, 53
 - Handover Required, 47
 - hard, 69
 - inter-RAT, 38
 - inter-RAT, 524–525
 - lossless, 52–53, 89, 93–95
 - non-blind, 522, 524
 - parallel, 73
 - parameter negotiation, 589
 - parameter optimization, 589
 - Preparation Failure, 47
 - preparation request, 71, 74
 - quality based, 514
 - Release Resource, 45, 53
 - report, 592, 593
 - Request, 567
 - robustness, 590–595
 - S1 interface, 45, 46, 50
 - Status Transfer, 45, 46
 - seamless, 52, 93
 - soft, 31
 - Status Transfer, 53
 - to an inappropriate cell, 592–593
 - to an unprepared cell, 594
 - to CDMA2000 1xRTT, 524–525
 - to E-UTRAN, 523–524
 - to GSM, 524
 - to UTRAN, 524
 - too-early, 592
 - too-late, 591
 - unnecessary inter-RAT, 594–595
 - X2 interface, 45, 50–53
 - Setup Request, 50

- Status Transfer, 45
- hard handover, 69
- hardware load, 588
- HARQ, *see* Hybrid Automatic Repeat reQuest
- HD-FDD, *see* Half-Duplex FDD
- header
 - compression, 31, 33, 89–91, 93
 - context, 93
 - decompression, 89
- headerless MAC PDU, 120
- heterogeneous network, 618, 623, 634–635, 701–703
 - co-channel, 701–703
 - CoMP, 670
 - energy saving, 596
 - macro-femto, 569
- HFN, *see* Hyper Frame Number
- High Interference Indicator (HII), 54, 290
- High Rate Packet Data (HRPD)
 - pilot strength, 516
- High Speed Downlink Packet Access (HSDPA), 4, 7
- High Speed Packet Access (HSPA), 4, 7
 - HSPA+, 4
- High Speed Uplink Packet Access (HSUPA), 4, 7
- HII, *see* High Interference Indicator
- Home eNodeB (HeNB), 509, 563–580, 669
 - access control, 566, 567
 - ACLR, 575, 576
 - ACS, 578
 - architecture, 564–565
 - blocking, 578
 - closed access, 565
 - demodulation performance, 578
 - dynamic range, 578
 - frequency error, 575, 576
 - Gateway (GTW), 564–565, 585
 - handover
 - inbound, 566–568
 - outbound, 566–568
 - hybrid access, 566
 - inbound mobility, 509
 - interference, 569–574
 - maximum output power, 464, 575
 - membership verification, 566, 567
 - mobility, 566–568
 - Noise Figure (NF), 576
 - paging, 566
 - power adjustment, 575
 - receiver requirements, 576–579
 - reference sensitivity, 576, 577
 - RF requirements, 574–579
 - sensing, 569
 - spectrum emission mask, 576
 - spurious emissions, 575, 576
 - timing synchronization, 579
 - transmitter requirements, 575–576
- Home Subscriber Server (HSS), 28
- hopping
 - PUSCH, 346–347
- Householder
 - codebook, 268
 - matrix, 268
 - transform, 268
- HSDPA, *see* High Speed Downlink Packet Access
- HSPA, *see* High Speed Packet Access
- HSS, *see* Home Subscriber Server
- HSUPA, *see* High Speed Uplink Packet Access
- hybrid access, 566
- Hybrid Automatic Repeat reQuest (HARQ), 108, 200, 240–242, 280
 - adaptive, 109, 241
 - asynchronous, 109, 241
 - backhaul, 698–699
 - buffer, 240, 242–243
 - carrier aggregation, 648
 - in TDD operation, 548–551
 - information, 204, 205, 241–242
 - non-adaptive, 109, 241
 - process number, 109
 - profile, 68
 - relaying, 698–699
 - Round Trip Time (RTT)
 - timer, 116
 - round trip times, 240, 551
 - synchronous, 109, 241, 345
 - timing, 241
 - downlink, 241
 - uplink, 241
- Hyper Frame Number (HFN), 52, 92
- ICIC, *see* Inter-Cell Interference Coordination
- ICI, *see* Inter-Carrier Interference
- IEEE
 - 802.16, 4
 - 802, 232
- IETF, *see* Internet Engineering Task Force

- IFDMA, *see* Interleaved Frequency Division Multiple Access
- IFFT, *see* Inverse Fast Fourier Transform
- IIP3, *see* Input IP3
- IM3, 490
- impairments, 492–500
- implementation margin, 478
- implicit feedback, 274
- impulse response, 127
- IMSI, *see* International Mobile Subscriber Identity
- IMT, *see* International Mobile Telecommunications
- in-sequence delivery, 89, 93, 94
- in-sync, 526
- Incremental Redundancy (IR), 238
- index remapping, 362
- Indoor Hotspot scenario (InH), 450, 603–608
- information
 - transfer, 78
- InH, *see* Indoor Hotspot
- Input IP3, 490
- integrity protection, 89, 91–93, 97
- inter-RAT
 - cell
 - reselection, 81
 - mobility, 73–74
- Inter-Symbol Interference (ISI), 300
- Inter-Carrier Interference (ICI), 133, 134, 136, 301
- inter-cell interference, 287, 298
- Inter-Cell Interference Coordination (ICIC), 287, 290, 291, 669, 701–710, 719
 - enhanced (eICIC), 701–710
- interference, 287
 - base-station to base-station, 536–539
 - coordination, 54, 278–292
 - semi-static, 291
 - static, 291
 - femtocell, 569–574
 - Home eNodeB, 569–574
 - information, 49
 - inter-OFDM symbol, 128
 - intra-OFDM symbol, 128
 - management, 54, 600–601
 - mobile-to-mobile, 539–543
 - randomization, 333
- Interference Rejection Combining (IRC), 8, 542, 543, 602
- Interleaved Frequency Division Multiple Access (IFDMA), 320, 338
- interleaver
 - contention-free, 234–238
 - turbo code, 229
- intermodulation
 - distortion, 489–491, 493, 499–500
 - products, 471, 489, 495, 499–500
- International Mobile Subscriber Identity (IMSI), 85
- International Mobile Telecommunications (IMT), 2, 12, 437
 - 2000, 439, 615
 - Advanced, 615–617
 - channel model, 449–452
 - spectrum, 620
- International Telecommunication Union (ITU), 531, 615
 - channel model, 439, 449–452
 - deployment scenario, 616
 - Indoor Hotspot (InH), 450
 - Rural Macro (RMa), 450
 - Urban Macro (UMa), 450
 - Urban Micro (UMi), 450
 - Pedestrian A, 440
 - Pedestrian B, 440
 - Radiocommunication sector (ITU-R), 2, 615
 - regions, 620
 - test environments, 449–452, 599, 603–608
 - base coverage urban, 449
 - high speed, 449
 - indoor, 449
 - microcellular, 449
 - Vehicular A, 440
- Internet Engineering Task Force (IETF), 295
- Internet Protocol (IP), 25, 26, 41, 91, 295
 - address, 36, 42
 - header, 36
 - header compression, 91
 - header configuration, 91
 - IPv4, 91
 - IPv6, 91
 - packet, 32
 - version, 41, 42, 49
- Inter-Operability Testing (IOT), 18
- Inter-Site Distance (ISD), 599
- inter-subframe hopping, 346

- Inter-Symbol Interference (ISI), 125, 127, 136, 139, 140, 165, 318
- inter-working, 38–39
- intraFreqReselection, 83
- intra-subframe hopping, 346
- Inverse Fast Fourier Transform (IFFT), 125, 127, 131, 133, 320
 - estimator, 178–179
- invoke indication, 705
- IOT, *see* Inter-Operability Testing
- IP, *see* Internet Protocol
- IP Multimedia Subsystem (IMS), 27, 295
- IP3, 490
- IPSec, 569, 586
- IR, *see* Incremental Redundancy
- IS-95, 4
- ISD, *see* Inter-Site Distance
- ISI, *see* Inter-Symbol Interference
- iterative decoding, 216
- ITU, *see* International Telecommunication Union

- Jain index, 282
- Jake's model, 134
- Japanese Total Access Communication System (J-TACS), 1
- Joint Transmission (JT), 670
- JT, *see* Joint Transmission
- J-TACS, *see* Japanese Total Access Communication System

- Kaiser window, 324
- K_{ASME} , 64
- K_{eNB} , 63
- Kronecker
 - assumption, 447
 - model, 176–178

- large adjacent interferer, 483
- Last Segment Flag (LSF), 107
- latency, 11–12, 505
 - control plane, 616, 617
 - user plane, 616, 617
- layer
 - spatial, 253, 266, 618
- Layer 1, 38
- Layer 2, 38, 87
- Layer 2/Layer 3 message, 375
- Layer 3, 34
- layers, 145
- LB, *see* Long Block

- LBRM, *see* Limited Buffer Rate Matching
- LDPC, *see* Low-Density Parity Check
- leakage
 - adjacent carrier, 469
 - transmit signal, 475–476
- Least Squares (LS), 179
 - estimator, 179
- leaving condition, 77
- Length Indicator (LI), 106
- Limited Buffer Rate Matching (LBRM), 242–243
- limited service, 79
- Line-Of-Sight (LOS), 442, 536
- linear channel estimation, 179–180
- linear interpolation estimator, 178
- Linear Minimum Mean Squared Error (LMMSE), 602
- link adaptation, 215–248, 560
- link budget, 608–610
- LIPA, *see* Local IP Access
- LLR, *see* Log-Likelihood Ratio
- LMMSE, *see* Linear Minimum Mean Squared Error
- LNA, *see* Low Noise Amplifier
- load
 - available composite, 588–589
 - balancing, 47, 587–590
 - exchange, 588–590
 - inter-RAT, 589–590
 - intra-LTE, 588–589
 - hardware, 588
 - indication, 54
 - information, 49
 - management, 47–48, 54
 - Overload, 47
 - overload indicator, 54
 - rebalancing, 48
 - Reporting Request, 590
 - Reporting Response, 590
 - transport, 588
- Local IP Access (LIPA), 568–569
 - bearer, 569
 - Gateway (L-GW), 568, 569
- local oscillators, 133, 134, 488
- localized
 - allocation, 346
 - mapping, 194
 - transmission, 319, 321
- location-based services, 423
- logical channel, 98, 108–110

- group, 114
- prioritization, 117–118, 286
 - carrier aggregation, 631
 - prioritization entity, 109
- Logical Channel ID (LCID), 119, 120
- Log-Likelihood Ratio (LLR), 226, 242
- log-MAP, 228
- Long Block (LB), 335–337
- long BSR, 114
- long DRX cycle, 116
- Long-Term Evolution (LTE)
 - evaluation channel model, 445–447
- Lorentzian spectrum, 497
- LOS, *see* Line-Of-Sight
- lossless handover, 89, 93–95, 98
- Low-Density Parity Check (LDPC), 231–232
- Low Noise Amplifier (LNA), 475
- LPP, *see* LTE Positioning Protocol
- LPPa, *see* LTE Positioning Protocol A
- LTE-Advanced, 19–20, 615–720
 - backward compatibility, 619
 - features, 618
 - frequency bands, 620
 - MIMO, 651–671
 - requirements, 616–617
 - system performance, 715–718
 - UE categories, 621
- LTE Positioning Protocol (LPP), 433–435
- LTE Positioning Protocol A (LPPa), 434–435

- M-sequences, *see* maximum length sequences
- M1 interface, 297
- M2 interface, 297
- M2M, *see* Machine-Type Communications
- M3 interface, 297
- MAC-I, *see* Message Authentication Code for Integrity
- Machine-Type Communications (MTC), 712–714, 720
- macrocell, 9
- Master Information Block (MIB), 60, 190, 509
- matched-filter, 251
- Maximum
 - Likelihood (ML), 162, 225, 245, 257
 - Likelihood Detection (MLD), 602
 - Power Reduction (MPR), 465–466
 - Ratio Combining (MRC), 255, 543, 602
- maximum
 - excess delay, 438, 440
 - length sequences (M-sequences), 148, 158, 332
 - output power, 464–466
 - base station, 464–465, 575
 - Home eNodeB, 575
 - reduction (MPR), 465–466
 - UE, 417, 465–466, 643–644
 - rate scheduling, 283–284
- Maximum Bit Rate (MBR), 34, 35
- max-log-MAP, 229
- MBMS, *see* Multimedia Broadcast/Multicast Service
- MBMS Radio Network Temporary Identifier (M-RNTI), 306
- MBMSCountingRequest message, 311
- MBR, *see* Maximum Bit Rate
- MBSFN, *see* Multimedia Broadcast Single Frequency Network
- MBSFNAreaConfiguration message, 304–306
- MCCH, *see* Multicast Control Channel
- MCE, *see* Multicell Coordination Entity
- MCH, *see* Multicast Channel
- MCH Scheduling Information (MSI), 306, 308
- MCH Scheduling Period (MSP), 306, 308
- MCH Subframe Allocation (MSA), 307
- MCL, *see* Minimum Coupling Loss
- MCS, *see* Modulation and Coding Scheme
- MDT, *see* Minimization of Drive Tests
- Mean Squared Error (MSE), 180, 182, 183
- measurement, 74–77, 512–516
 - carrier aggregation, 625–627
 - CDMA2000, 516
 - configuration, 58, 74–77
 - enhanced inter-RAT, 512
 - E-UTRAN, 513–514
 - event, 76
 - A1, 76, 517, 625
 - A2, 76, 517, 625, 712
 - A3, 76, 517, 626
 - A4, 76, 517
 - A5, 76, 517, 626
 - A6, 626
 - B1, 76
 - B2, 76, 626
 - carrier aggregation, 625
- event-triggered reporting, 75–77, 516
 - E-UTRAN, 517
 - inter-RAT, 517–518
- gap, 70, 72, 75, 113, 626

- inter-frequency, 506, 507
- inter-RAT, 506, 507, 511
- GSM, 516
- identity, 75, 625
- object, 75, 625
- period, 626
- periodic reporting, 516
- reporting, 58, 77, 516–518
 - requirement, 516–518
- restricted, 706–710
 - CSI, 708–710
 - RLM, 707, 709
 - RRC signalling, 708–710
 - RRM, 707–709
 - RSRP, 708
 - RSRQ, 707–708
 - RSSI, 707
- rule, 80–81
- UTRAN, 514–516
- MeasurementReport, 70
- Medium Access Control (MAC), 32, 87, 107–120
 - ciphering, 59
 - control element, 109, 117, 306, 411, 630
 - Buffer Status Report, 119
 - C-RNTI MAC, 119
 - DRX Command, 119
 - MBMS Dynamic Scheduling Information, 119
 - power headroom, 119
 - Timing Advance Command, 119
 - UE Contention Resolution Identity, 119
 - header, 119
 - integrity protection, 59
 - PDU format, 118–120
 - subheader, 119
- Message Authentication Code, 73
- Message Authentication Code for Integrity (MAC-I), 92, 97
- MIB, *see* Master Information Block
- microcell, 9
- midpaths, 444
- MIMO, *see* Multiple-Input Multiple-Output
- Minimization of Drive Tests (MDT), 710–712, 720
 - consent, 712
 - data protection, 712
 - immediate, 710, 712
 - logged, 710–712
 - privacy, 712
- Minimum Coupling Loss (MCL), 539
- Minimum Mean-Squared Error (MMSE), 179–185
 - estimator, 179–181
 - mismatched, 179–181
 - Successive Interference Cancellation (SIC), 266
- mixed PUCCH RB, 362–363
- ML, *see* Maximum Likelihood
- MMSE-SIC, 266
- Mobile Originated (MO)
 - call, 83
 - signalling, 83
- mobile relay, 679
- mobile TV, 11
- mobile-to-mobile interference, 539–543
- mobility, 10–11, 68–72, 504, 512
 - carrier aggregation, 625–627
 - CDMA2000, 74
 - Change Acknowledge, 589
 - Change Failure, 589
 - Change Request, 589
 - control, 68–72
 - Home eNodeBs, 47
 - inter-RAT, 73–74
 - network controlled, 58
 - inter-LTE, 45–47
 - intra-LTE, 45
 - performance, 518–525
 - robustness, 53
 - robustness optimization, 590–595
 - S1 interface, 45–47
 - Settings Change procedure, 589
 - X2 interface, 50–54
- Mobility Management Entity (MME), 27, 28, 30, 296, 433
 - Configuration Update, 585
 - Home eNodeB mobility, 566–568
 - pool, 31
 - self-configuration, 584–587
- MobilityFromEUTRACCommand, 74
- modification period, 61, 62
- modulation, 216–218
- Modulation and Coding Scheme (MCS), 204, 205, 216–218, 663
- modulator image interference, 495
- MPR, *see* Maximum Power Reduction
- MPS, *see* Multimedia Priority Service

- M-RNTI, *see* MBMS Radio Network Temporary Identifier
- MSA, *see* MCH Subframe Allocation
- MSE, *see* Mean Squared Error
- MSI, *see* MCH Scheduling Information
- MSP, *see* MCH Scheduling Period
- MTC, *see* Machine-Type Communications
- MTCH, *see* Multicast Traffic Channel
- multicarrier, 123–124
 - modulation, 123–124
 - transmission, 123–124
- multicast, 293–295, 297, 314
- Multicast Channel (MCH), 111, 299, 304
- Multicast Control Channel (MCCH), 110, 304, 306–307
 - modification period, 306–307
 - update time, 308
- Multicast Traffic Channel (MTCH), 110, 304
- Multicell Coordination Entity (MCE), 295, 296, 305, 308, 309
- multicell MU-MIMO, 670
- Multimedia Broadcast Single Frequency Network (MBSFN), 111, 140, 167, 297–303
 - area, 299, 301–303, 307–308
 - reference signal, 167, 300
 - subframe, 299
 - ABS, 704
 - for relaying, 686
 - PDSCH transmission, 660
 - synchronization area, 301
- Multimedia Broadcast/Multicast Service (MBMS), 11, 99, 196, 293–314, 505
 - architecture, 295
 - channel switching time, 306
 - Common Subframe Allocation (CSA), 307–308
 - content, 313
 - content synchronization, 308–310
 - control signalling, 304, 306–307
 - change, 306–307
 - modification period, 306–307
 - counting procedure, 295, 310–311
 - dedicated carrier, 300
 - Gateway (MBMS GW), 296
 - interfaces, 296–297
 - M1, 297
 - M2, 297
 - M3, 297
 - protocol, 303–311
 - reference signals, 167, 300
 - resource allocation, 307–308
 - scheduling, 306
 - service area, 296, 301, 305
 - service prioritization, 303
 - session start, 304–305
 - session stop, 304–305
- Multimedia Priority Service (MPS), 44
- multipath, 127
- multiple antenna techniques, 249–276, 367–369, 601
- multiple preparation, 53
- Multiple-Input Multiple-Output (MIMO), 16, 249–276, 367–369, 560–561
 - channel model, 441–449
 - cooperative, 668–671
 - LTE-Advanced, 618, 651–671
 - downlink, 659–665
 - uplink, 665–668
 - multi-user, 205
 - Release 10, 661–662
 - network, 670
 - performance, 275–276
- Multiple-Input Single-Output (MISO), 250
- multiplexing
 - ACK/NACK, 549–551
 - entity, 109
 - gain, 251, 253, 257
- Multi-User
 - capacity, 261–262
 - diversity, 282, 284, 287
 - MIMO (MU-MIMO), 205, 250, 258–264, 274–275, 368
 - multicell, 670
 - Release 10, 661–662
 - scheduling, 278–292
- MU-MIMO, *see* Multi-User MIMO
- muting
 - PDSCH for CSI-RS, 657
 - PRS, 429
- mutual information, 283, 285
- N310, 526
- N311, 526
- NACC, *see* Network Assisted Cell Change
- NACK, 243, 351
- NACK_SN List, 107
- narrowband
 - blocking, 485

- intermodulation requirement, 489
- NAS, *see* Non-Access Stratum
- NAS Node Selection Function (NNSF), 43, 47
- NCC, *see* Network Colour Code
- NCL, *see* Neighbour Cell List
- NDI, *see* New Data Indicator
- NDS, *see* Network Domain Security
- Neighbour Cell List (NCL), 68, 512
- Neighbour cell Relation (NR), 509, 582
- neighbour detection function, 583
- Neighbour Relation Table (NRT), 582–583
 - management function, 583
 - No Handover flag, 583
 - No Remove flag, 583
 - No X2 flag, 583
- neighbour removal function, 583
- network architecture, 25–55
- Network Assisted Cell Change (NACC), 73
- network coding, 684
- Network Colour Code (NCC), 511
- Network Domain Security (NDS), 691
- network energy saving, 596–597, 720
- Network Listen Mode (NLM), 573–574
- network MIMO, 670
- New Data Indicator (NDI), 204, 241
 - uplink MIMO, 666
- Next Generation Mobile Networks (NGMN), 7, 9, 10
- Next Hop (NH), 63
 - Chaining Counter (NCC), 63
- NF, *see* Noise Figure
- NGMN, *see* Next Generation Mobile Networks
- NLM, *see* Network Listen Mode
- NLOS, *see* non-Line-Of-Sight
- NMT, *see* Nordic Mobile Telephone
- NNSF, *see* NAS Node Selection Function
- Node Selection Function (NSF), 43
- noise figure, 537
- Noise Figure (NF), 479
 - base station, 576
 - Home eNodeB, 576
- nominal maximum power, 466
- non-Line-Of-Sight (NLOS), 441
- non-adaptive HARQ, 109, 368
- non-adjacent channel selectivity, 485–486
- non-coherent
 - detection, 163
- non-frequency-selective scheduling, 345–346
- non-Guaranteed Bit Rate, 34
- non-linear distortion, 495, 499–500
- non-linearity, 131–138
- non-selective scheduling, 346–347
- Non-Access Stratum (NAS), 28–30, 57
 - ciphering, 59
 - common information, 58
 - dedicated information, 58
 - Extended Service Request, 39
 - integrity protection, 59
 - states, 65
- Nordic Mobile Telephone (NMT), 1
- normal CP, 140–142
- normalized least-mean-square, 183–184
- NR, *see* Neighbour cell Relation
- NRT, *see* Neighbour Relation Table
- NSF, *see* Node Selection Function
- Nyquist's sampling theorem, 169
- OBPD, *see* Occupied Bandwidth Power
 - De-rating
- Observed Time Difference Of Arrival (OTDOA), 19, 424, 426–430, 435
 - errors, 429
 - performance, 429
- OCC, *see* Orthogonal Cover Code
- occupied bandwidth, 473
- Occupied Bandwidth Power De-rating (OBPD), 473
- OFDM, *see* Orthogonal Frequency Division Multiplexing
- OFDMA, *see* Orthogonal Frequency Division Multiple Access
- OI, *see* Overload Indicator
- on duration, 115, 116
- on-off mask, 466
- OOB, *see* Out-Of-Band
- open access, 566
- Open Subscriber Group (OSG), 563
- open-loop spatial multiplexing, 268–270
- operator service, 79
- opportunistic beamforming, 669
- opportunistic scheduling, 281
- orthogonal block spreading code remapping, 358, 362
- Orthogonal Cover Code (OCC), 652, 657–658
- Orthogonal Frequency Division Multiple Access (OFDMA), 14, 123–143, 600
- Orthogonal Frequency Division Multiplexing (OFDM), 3, 14, 15, 123–143

- symbol, 125, 128, 130, 131
- OSG, *see* Open Subscriber Group
- OTDOA, *see* Observed Time Difference Of Arrival
- out-of-sync, 526
- outage, 10, 703
 - probability, 279
- Out-Of-Band (OOB), 467–471
 - blocking, 486–487
 - emissions, 349, 467–470
- out-of-sync
 - heterogeneous network, 707
- Output IP3, 490
- overhead, 602–603
- Overload Indicator (OI), 290, 414, 419

- P-RNTI, *see* Paging RNTI
- PA, *see* Power Amplifier
- Packet Data Convergence Protocol (PDCP), 32, 87–98
 - control PDU, 90, 97
 - data PDU, 90, 97
 - PDU format, 97–98
 - Sequence Number (NS), 53
 - status report, 95, 98
- Packet Data Network (PDN), 25, 26
 - address, 28
 - GateWay (P-GW), 27, 28, 38
- Packet Data Network GateWay (P-GW), 31
- Packet Switching (PS), 16–17, 25, 503
- Paging
 - RNTI (P-RNTI), 84, 196
- paging, 44–45, 84–85, 196
 - cycle, 85
 - cell-specific, 85
 - UE-specific, 85
 - frame (PF), 85
 - interruption, 522
 - message, 59, 62, 192, 205
 - occasion (PO), 85
 - Paging Request, 44
 - Priority Indicator, 44
- Paging Channel (PCH), 111
- Paging Control CHannel (PCCH), 59, 110
- paired spectrum, 12, 147, 531–562
- PAPR, *see* Peak-to-Average Power Ratio
- Parallel-to-Serial (P/S), 127
- partial frequency reuse, 289, 669
- partial path-loss compensation, 414
- path loss, 439, 445
 - compensation, 413
 - estimation, 413, 643
- Path Switch Request, 568
- PBR, *see* Prioritized Bit Rate
- PCC, *see* Policy Control and Charging, *see* Primary Component Carrier
- PCCH, *see* Paging Control Channel
- PCEF, *see* Policy Control Enforcement Function
- PCell, 624
 - change, 625
- PCFICH, *see* Physical Control Format Indicator Channel
- PCH, *see* Paging Channel
- PCI, *see* Physical Cell Identity
- P_{CMAX} , 466
- P-CPICH, *see* Primary Common Pilot Channel
- PDCCH, *see* Physical Downlink Control Channel
- PDCCH order, 204
- PDSCH, *see* Physical Downlink Shared Channel
- PDU, *see* Protocol Data Unit
- peak rates, 8–10
- Peak-to-Average Power Ratio (PAPR), 15, 317, 473
 - OFDM, 124, 131–133
- P_{EMAX} , 466
- performance
 - MIMO, 275–276
 - system, 599–611
 - LTE-Advanced, 715–718
- performance requirements, 7–12
- periodic CSI reporting, 221–223
 - Release 10, 664–665
- periodicBSR, 114
- PF, *see* Paging Frame
- P-GW, *see* Packet Data Network GateWay
- phase
 - imbalance, 495–497
 - noise, 133, 493, 494, 497–499
 - reference, 193
 - shift, 136
- Phase-Locked Loops (PLL), 498
- PHICH, *see* Physical Hybrid ARQ Indicator Channel
- PHR, *see* Power Headroom Report
- Physical Broadcast Channel (PBCH), 152, 189–192, 201, 264

- decoding, 192
- Physical Cell Identity (PCI), 50, 152, 155, 332, 505, 582
 - automatic configuration, 587
- Physical Control Format Indicator Channel (PCFICH), 198–199
- Physical Downlink Control CHannel (PDCCH)
 - BLER, 527
- Physical Downlink Control Channel (PDCCH), 202–212, 264
 - aggregation level, 202
 - BLER, 526
 - blind decoding, 211–212
 - carrier aggregation, 633
 - channel coding, 244
 - formats, 203–207
 - order, 204
 - search space, 211
 - carrier aggregation, 633–634
- Physical Downlink Shared Channel (PDSCH), 192–196, 264
 - transmission modes, 193
 - transmission mode 9, 659–665
- Physical Hybrid ARQ Indicator CHannel (PHICH), 199–202, 549
 - carrier aggregation, 636
 - groups, 201, 202
 - uplink MIMO, 666
- physical layer
 - cell ID, 152, 155, 332
- Physical Multicast Channel (PMCH), 196, 299–301
- Physical Random Access CHannel (PRACH), 371, 376–396
 - formats, 378
 - guard-band, 387
 - preamble transmitter, 397
 - receiver, 398–404
 - resource configurations, 383–385
 - sequence ordering, 394–396
 - sequence planning, 395
 - slot, 378
 - transmitter, 396–398
- Physical Random Access Channel (PRACH), 343
 - preamble, 526
- Physical Resource Block (PRB), 208
- Physical Uplink Control CHannel (PUCCH), 376
 - Physical Uplink Control Channel (PUCCH), 343, 348–364
 - carrier aggregation, 636–641
 - channel selection, 637
 - format 1, 363–364
 - format 1a/1b, 356–363
 - format 1b with channel selection, 639–640
 - format 2, 352–356
 - format 2a/2b, 353–356
 - format 3, 637–639
 - mixed RB, 362–363
 - periodic CSI reporting, 219
 - Release 10, 664–665
 - region, 348–352
 - transmit diversity, 667–668
 - Physical Uplink Shared CHannel (PUSCH), 376
 - Physical Uplink Shared Channel (PUSCH), 343–348
 - antenna selection, 367–368
 - aperiodic CSI reporting, 219
 - Release 10, 664
 - control signalling, 365–367
 - hopping, 346–347
 - transmission mode, 666
 - transport block size, 347–348
- picocell, 563–580
- PLL, *see* Phase-Locked Loops
- PLMN, *see* Public Land Mobile Network
- PMCH, *see* Physical Multicast Channel
- PMI, *see* Precoding Matrix Indicator
- PMIP, *see* Proxy Mobile Internet Protocol
- P-MPR, *see* Power-management MPR
- PO, *see* paging Occasion
- polarization covariance matrices, 447
- Policy Control and Charging (PCC)
 - Decision Provision, 37
- Policy Control and charging Rules Function (PCRF), 27, 34, 37
- Policy Control Enforcement Function (PCEF), 27
- polling, 103–106
- pool area, 31, 43
- positioning, 19, 31, 423–436
 - A-GNSS, 424–426
 - UE-assisted, 426
 - AoA, 424, 432–433
 - assistance data, 426
 - cell-ID, 424, 430–433

- enhanced cell-ID, 424, 431–434
- occasion, 428
- OTDOA, 424, 426–430
 - errors, 429
 - performance, 429
- protocols, 433–435
- reference signals, 427–429, 435
- RF pattern matching, 436
- RTT, 432–433
- subframe, 427
- UTDOA, 436
- Positioning Reference Signal (PRS), 427–429
 - configuration, 435
 - muting, 429
 - power, 429
- Power Amplifier (PA), 131, 460, 471–473
 - compression point, 477
 - de-rating, 471
 - rated output power, 471
- power class
 - UE, 465
- power control, 411–420, 560
 - carrier aggregation, 642–644
 - command, 204, 205, 416
 - fractional, 413–414
- Power Delay Profile (PDP), 175, 392, 399–401, 438
- power headroom, 279, 473
 - reporting, 419–420
 - carrier aggregation, 644
 - prohibit timer, 420
- Power Headroom Report (PHR), 419–420
 - extended, 644
 - minimization of drive tests, 712
 - prohibit timer, 420
 - virtual, 644
- Power-management MPR (P-MPR), 465–466
- power transient, 466
- PRB, *see* Physical Resource Block
- preamble
 - format, 378, 382, 383
 - format 4, 404–405
 - power ramping, 375
- Precoder Type Indication (PTI), 665
- precoding, 266–268
 - information, 205
 - PUSCH, 666
 - uplink, 666
- Precoding Matrix Indicator (PMI), 272–274, 348, 352, 355, 364–366, 662–665
 - sub-band, 273, 665
 - UE-selected sub-band, 273
 - wideband, 273, 665
- Precoding Resource block Group (PRG), 661
- PRG, *see* Precoding Resource block Group
- primary channel model, 451
- Primary Common Control Physical CHannel (P-CCPCH)
 - RSCP, 515
- Primary Common Pilot CHannel (P-CPICH), 510
 - E_c/N_0 , 511
 - RSCP, 511
- Primary Component Carrier (PCC), 624
- primary scrambling code, 510
- Primary Synchronization CHannel (P-SCH), 510
- Primary Synchronization Code (PSC), 510
- Primary Synchronization Signal (PSS), 151–155, 157–158
- Prioritized Bit Rate (PBR), 67, 117, 118
- propagation, 437–456
 - delay, 408–411, 439
- Proportional Fair Scheduling (PFS), 284–285
- protocol, 32–34
 - for MBMS, 303–311
 - for positioning, 433–435
 - tunnelling, 32
- Protocol Data Unit (PDU), 88
 - in-sequence delivery, 89
 - reordering, 89
- Proxy Mobile Internet Protocol (PMIP), 32, 38
- PRS, *see* Positioning Reference Signal
- PS, *see* Packet Switching
- P/S, *see* Parallel-to-Serial
- PSC, *see* Primary Synchronization Code
- P-SCH, *see* Primary Synchronization Channel
- PSS, *see* Primary Synchronization Signal
- PTI, *see* Precoder Type Indication
- Public Land Mobile Network (PLMN), 31, 57, 582
 - equivalent, 79
 - high quality, 79
 - Registered (R-PLMN), 79
 - selection, 78–84
- Public Warning System (PWS), 19, 311–312
- PUCCH, *see* Physical Uplink Control Channel
- pulse shaping, 324–325
- PUSCH, *see* Physical Uplink Shared Channel

- PWS, *see* Public Warning System
- QAM, *see* Quadrature Amplitude Modulation
- QCI, *see* QoS Class Identifier
- QoS, *see* Quality of Service
- QPP, *see* Quadratic Permutation Polynomial
- QPSK, *see* Quadrature Phase Shift Keying
- Quadratic Permutation Polynomial (QPP), 236–239, 243
- Quadrature Amplitude Modulation (QAM), 193, 194, 216–218
- quadrature error, 493, 494
- Quadrature Phase Shift Keying (QPSK), 193, 194, 216–218
- Quality of Service (QoS), 25, 26, 34–38, 44, 279, 281–283, 285, 286
 - Class Identifier (QCI), 34–35, 44
- quantity configuration, 75
- quantization, 263, 663
- queue status, 287
- RACH, *see* Random Access Channel
- radiation pattern, 446
- radio
 - controller, 31
 - frame, 145
- Radio Access Network (RAN)
 - energy saving, 596–597, 720
- Radio Frequency (RF), 457
 - band, 461
 - impairments, 492–500
 - receiver requirements, 473–492
 - transmitter requirements, 462–473
- Radio Interface Technology (RIT), 615
 - FDD, 617
 - Set of (SRIT), 617
 - TDD, 617
- Radio Link Control (RLC), 32, 59, 87, 98–107
 - Control PDU, 103, 105
 - Data PDU, 103, 105
 - entities, 98
 - mode configuration, 34, 67
- Radio Link Failure (RLF), 73, 526, 527, 591
 - backhaul link, 690
 - carrier aggregation, 627
 - heterogeneous network, 707
 - Indication, 591, 592, 594
 - minimization of drive tests, 712
- Radio Link Monitoring (RLM), 526–528
 - restricted measurement, 707, 709
- radio propagation, 437–456
- Radio Resource Control (RRC), 34, 58–78
 - ciphering, 63
 - connection
 - control, 58, 62–73
 - establishment, 64–67
 - mobility control, 525–526
 - re-establishment, 72–73, 525
 - release, 64–67
 - connection establishment for carrier aggregation, 624
 - Connection Reconfiguration, 37
 - context information, 58
 - dedicated message, 59
 - integrity protection, 63
 - relaying, 689
 - state, 57
- Radio Resource Management (RRM), 503–530
 - restricted measurement, 707–709
- Rake receiver, 8, 10
- RAN, *see* Radio Access Network
- RAN Information Management (RIM), 590, 595
- random access, 72, 371–406, 525–526
 - backhaul link, 690
 - backoff indicator, 374
 - contention-based, 73, 372–376, 394
 - contention-free, 372, 376
 - contention resolution, 375–376
 - control entity, 109
 - dedicated preamble, 376
 - Layer 2/Layer 3 message, 375
 - preamble, 372, 373, 376–396
 - bandwidth, 385–387
 - cyclic prefix duration, 381–383
 - format, 378
 - guard time duration, 381–383
 - power ramping, 375
 - sequence duration, 378–381
 - sequence length, 387–388
 - transmission, 373–374
 - procedure, 115, 372–376
 - carrier aggregation, 628
- Random Access CHannel (RACH), 111, 371
 - Self-Optimization, 595–596
 - WCDMA, 371–372
- Random Access Preamble IDentifier (RAPID), 120
- Random Access Radio Network Temporary Identifier (RA-RNTI), 374

- Random Access Response (RAR), 120, 374–375
 - window, 374
- range expansion, 702–703
- rank, 266
 - channel, 266
 - transmission, 266
- Rank Indicator (RI), 272–274, 348, 352, 364–366, 662, 665
- RAR, *see* Random Access Response
- RA-RNTI, *see* Random Access Radio Network Temporary Identifier
- rate matching, 238–240
- rated output power, 471
- RB, *see* Resource Block
- RBG, *see* Resource Block Group
- R-criterion, 82
- RE, *see* Resource Element
- Real-time Transport Control Protocol (RTCP), 96
- Real-time Transport Protocol (RTP), 91
 - header compression, 91
- reassembly, 100–101
- receive diversity, 190, 474
- receive spatial signature, 252, 258
- Received Signal Strength Indicator (RSSI), 511, 514
 - GSM Carrier, 516
 - restricted measurement, 707
 - UTRAN FDD, 515
 - UTRAN TDD, 515
- receiver
 - advanced, 602
 - IRC, 602
 - LMMSE, 602
 - MLD, 602
 - MRC, 602
 - noise figure, 478–479
 - sensitivity, 479
 - SIC, 602
- reciprocal mixing, 499
- reciprocity, 552–562
 - aided beamforming, 658
 - applications, 558–561
- Redundancy Version (RV), 204, 233, 238–241
- Reed–Muller code, 353, 638
- reference sensitivity, 479–482
 - base station, 576, 577
 - Home eNodeB, 576, 577
- Reference Signal (RS), 165–185
 - backhaul link, 691–692
 - base sequence, 329–330
 - cell-specific, 167–171, 527
 - Channel State Information (CSI-RS), 167, 653–657
 - common, 167–171
 - DeModulation (DM-RS), 167
 - downlink, 652–653
 - uplink, 327, 335–337, 657–658
 - downlink, 167–174, 651–657
 - frequency shift, 171
 - initialization, 170, 173
 - MBSFN, 167, 300
 - Positioning (PRS), 167, 427–429
 - sequence, 169
 - sounding (SRS)
 - uplink, 327, 337–340, 658–659
 - time offset tracking, 136
 - UE-specific, 167, 171–174
 - Release 8, 171
 - Release 9, 171–174
 - Release 10, 652–653
 - uplink, 327–341, 657–659
 - demodulation, 327, 335–337, 647
 - sounding (SRS), 327, 337–340, 658–659
- Reference Signal Received Power (RSRP), 505, 513–514, 519, 520
 - for positioning, 432
 - map, 432
 - minimization of drive tests, 711, 712
 - restricted measurement, 708
- Reference Signal Received Quality (RSRQ), 505, 514, 519, 520
 - for positioning, 432
 - map, 432
 - minimization of drive tests, 711, 712
 - restricted measurement, 707–708
- Reference Signal Time Difference (RSTD), 426
- Reference Signals (RS), 513
- Reference Symbol, 168–171, 300, 301
- REFSENS, 479
- REG, *see* Resource Element Group
- region 1, 620
- region 2, 620
- region 3, 620
- regularized least squares, 179

- Relative Narrowband Transmit Power (RNTP), 290
- relaying, 619, 673–699
 - access link, 674
 - architecture, 687
 - backhaul HARQ, 698–699
 - backhaul link, 674
 - backhaul subframe configuration, 696–697
 - bearers, 688
 - capacity, 680–683
 - control plane, 688
 - cut-set bound, 680
 - deployment scenarios, 677–679
 - direct link, 674
 - inband, 674
 - initialization, 689–690
 - network coding, 684
 - outband, 674, 685
 - radio link failure, 690
 - random access, 690
 - repeater, 673–674
 - RRC, 689
 - security, 690–691
 - system information, 689
 - theory, 679–684
 - user plane, 688
- Relay Node (RN), 619, 673–699
 - deployment scenarios, 677–679
 - indicator, 689
 - initialization, 689–690
 - Layer 1, 677
 - Layer 2, 677
 - Layer 3, 677
 - mobile, 679
 - reception window, 686
 - Type 1, 684, 696
 - Type 1a, 684–685
 - Type 1b, 684–685
 - Type 2, 684–685
- Relay Physical Downlink Control Channel (R-PDCCH), 692–694
 - cross-interleaved, 692–693
 - non-cross-interleaved, 692–694
 - search space, 694
- Release 6, 7, 8, 10
- Release 7, 4
- Release 8, 4
- Release 9, 4, 19, 293
- Release 10, 4, 19–20, 615–718
- Release 11, 597, 719–720
- Release 99, 4
- Remote Radio Head (RRH), 670
- reordering, 89, 98, 100–101
 - timer, 100
- repeater, 673–674
- replay attack, 92
- reportAmount, 77
- reporting
 - configuration, 75
- reportInterval, 77
- requirements, 6–13
 - control channel, 196–198
 - RF, 462–492
- resegmentation, 102–103, 106
- Resegmentation Flag, 106
- resource allocation, 207–211, 279, 280, 283–285
 - clustered DFT-S-OFDM, 647, 666
 - opportunistic, 282
 - Type 0, 208
 - Type 1, 210
 - Type 2, 210
 - type bit, 647
- Resource Block (RB), 139, 146, 193–195, 290, 322
 - assignment, 204, 205
 - grant, 204
- Resource Block Group (RBG), 208, 209
- Resource Element (RE), 146, 322
- Resource Element Group (REG), 198–203, 207
- Resource Indication Value (RIV), 210
- resource scheduling, 280, 281, 286–287
- Resource Status Request, 588
- Resource Status Response, 588
- Resource Status Update, 588
- restricted measurement, 706–710
 - CSI, 708–710
 - RLM, 707, 709
 - RRC signalling, 708–710
 - RRM, 707–709
 - RSRP, 708
 - RSRQ, 707–708
 - RSSI, 707
- retransmission, 102–103
 - selective, 53
- retxBSR, 114
- RF, *see* Radio Frequency
- RF pattern matching, 436

- RFC 3095, 90, 91
- RFC 3843, 91
- RFC 4815, 91
- RFC 4995, 90, 91
- RFC 4996, 91
- RFC 5225, 91
- RI, *see* Rank Indicator
- RIM, *see* RAN Information Management
- Rise over Thermal (RoT), 412
- RIT, *see* Radio Interface Technology
- RIV, *see* Resource Indication Value
- RLF, *see* Radio Link Failure
- RMa, *see* Rural Macrocell
- RN, *see* Relay Node
- roaming, 1, 31–32
- RObust Header Compression (ROHC), 90–91
 - feedback, 97
- ROHC, *see* RObust Header Compression
- Root-Raised-Cosine, 324
- Round Trip Time (RTT), 431–433
- R-PDCCH, *see* Relay Physical Downlink Control Channel
- RRC_CONNECTED, 12, 57, 58, 371–372
 - cell search, 505
 - measurement, 513–516
 - mobility, 68–72, 503, 522–525
- RRC_IDLE, 12, 57, 58
 - measurement, 513–516
 - mobility, 68–72, 503, 518–522
- RRC signalling
 - restricted measurement, 708–710
- RRConnectionReconfiguration, 67, 71, 75, 78
- RRConnectionReconfigurationComplete, 67
- RRConnectionReestablishmentRequest, 73, 118
- RRConnectionRequest, 66
- RRConnectionSetup, 66
- RRConnectionSetupComplete, 66, 78
- RRH, *see* Remote Radio Head
- RRM, *see* Radio Resource Management
- RS, *see* Reference Signal
- RSRP, *see* Reference Signal Received Power
- RSRQ, *see* Reference Signal Received Quality
- RSSI, *see* Received Signal Strength Indicator
- RSTD, *see* Reference Signal Time Difference
- RTCP, *see* Real-time Transport Control Protocol
- RTT, *see* Round Trip Time
- Rural Macrocell scenario (RMa), 450, 603–608
- RV, *see* Redundancy Version
- S-criterion, 626
- S1 interface, 30, 35, 40–48, 584
 - application protocol, 41
 - bearer, 35
 - connection, 65
 - control plane, 41–43
 - flex, 31, 43
 - Home eNodeB Gateway, 565
 - handover, 45, 46, 50, 586
 - Status Transfer, 45, 46
 - mobility, 45–47
 - Setup, 43
 - Setup procedure, 585
 - user plane, 42–43
- S1-MME interface
 - Home eNodeB, 564
- S5 interface, 568
- S5/S8 interface, 32, 35
 - bearer, 35
- SAE-Temporary Mobile Subscriber Identity (S-TMSI), 29, 44
- sampling
 - frequency, 141
 - rate, 324, 498
- SAP, *see* Service Access Point
- SAW, *see* Stop-And-Wait
- scattering, 171–176
- SCC, *see* Secondary Component Carrier
- SCell, 624
 - activation, 630
 - addition, 625
 - deactivation, 630
- scenario
 - deployment, 599
- SC-FDMA, *see* Single-Carrier Frequency Division Multiple Access
- scheduling, 111–114, 194, 212, 279–287
 - coordinated, 669
 - dynamic, 112
 - fair, 282, 285
 - fast packet, 16
 - frequency selective, 337, 345–346
 - grant, 369
 - information transfer, 114
 - maximum sum rate, 284
 - maximum-rate constant-power, 284

- non-frequency selective, 346
- opportunistic, 282, 285
- Proportional Fair (PFS), 284
- PUSCH, 345–346
- queue-aware, 287
- Request (SR), 114, 348, 363, 372
- Semi-Persistent (SPS), 68, 113, 413
- system information, 61
- Scheduling Request (SR), 115
- SCID, *see* SCrambling IDentity
- SCM, *see* Spatial Channel Model
- SCME, *see* Spatial Channel Model – Extension
- scrambling, 148–149
 - IDentity (SCID), 662
 - secondary synchronization signal, 158
 - sequence, 148, 199
 - sequence generator, 148
 - UE-specific RS, 661
- S-criterion, 80, 81, 519
- SCTP, *see* Stream Control Transmission Protocol
- SDO, *see* Standards Development Organization
- SDU, *see* Service Data Unit
- seamless handover, 93
- search space, 211, 212
 - carrier aggregation, 633–634
 - common, 211
 - UE-specific, 211
- Second Generation system, 1
- Secondary Component Carrier (SCC), 624
- Secondary Synchronization CHannel (S-SCH), 510
- Secondary Synchronization Code (SSC), 510
- Secondary Synchronization Signal (SSS), 151–155, 158–161
- secret
 - key, 63
- Secure User Plane Location (SUPL), 423, 433
- security, 31, 63, 91–93
 - forward, 63
 - key, 30, 63–64
 - relaying, 690–691
- Security Gateway (SeGW), 565, 568
- SecurityModeCommand, 66
- SecurityModeComplete, 67
- SecurityModeFailure, 67
- segmentation, 99–100
 - Offset (SO), 107
- selective retransmission, 53
- Self-Optimizing Network (SON), 581–598, 720
 - Automatic Neighbour Relation (ANR), 582–584
 - energy saving, 596–597
 - eNodeB self-configuration, 584–587
 - handover robustness, 590–595
 - load-balancing, 587–590
 - MME self-configuration, 584–587
 - mobility robustness, 590–595
 - PCI configuration, 587
 - RACH optimization, 595–596
- Semi-Persistent Scheduling (SPS), 68, 113, 198, 413, 601
 - C-RNTI (SPS-C-RNTI), 113
 - carrier aggregation, 628
 - release, 113
- semi-static interference coordination, 291
- sequence
 - grouping, 329–330
 - hopping, 330, 333, 334
- sequence-group, 330, 331
 - hopping, 332–333
 - planning, 333
- Sequence Number (SN), 52, 53, 92, 106, 107
- Serial-to-Parallel (S/P), 125
- Service Access Point (SAP), 98
- Service Data Unit (SDU), 88, 108
- Service Specific Access Control (SSAC), 83
- serving cell
 - carrier aggregation, 624
- Serving GateWay (S-GW), 27, 28, 30, 38
 - pool, 31
- Session Initiation Protocol (SIP), 295
- session start, 304–305
- session stop, 304–305
- Set of RITs (SRIT), 617
- SFBC, *see* Space-Frequency Block Code
- SGi interface, 568, 569
- SGs interface, 39
- S-GW, *see* Serving GateWay
- shadowing, 439
- Shannon capacity, 261, 283
- Short Block (SB), 335
- short BSR, 114
- short DRX cycle, 116
- SI, *see* System Information
- SIB, *see* System Information Block
- SIC, *see* Successive Interference Cancellation

- Signal-to-Interference plus Noise Ratio (SINR)
 - target, 414
- Signalling Radio Bearer (SRB), 58, 59
 - SRB0, 59
 - SRB1, 59, 65, 73
 - SRB2, 59, 65, 67, 73
- Signal-to-Interference plus Noise Ratio (SINR), 215–217, 317
 - requirements, 475, 478
 - target, 414
- SIMO, *see* Single-Input Multiple-Output
- sinc pulse-shaping, 320–321
- single frequency network, 11, 293, 297
- Single Radio Voice Call Continuity (SRVCC), 39–40
- Single-Carrier Frequency Division Multiple Access (SC-FDMA), 317–326, 343, 600
- Single-Input Multiple-Output (SIMO), 250
 - channel model, 438–441
- Single-Input Single-Output (SISO)
 - channel model, 438–441
- Single-User MIMO (SU-MIMO), 250, 253–259, 264–274
 - PUSCH, 666–667
 - uplink, 666–667
- Single-Carrier Frequency Division Multiple Access (SC-FDMA), 14, 15
- Singular-Value Decomposition (SVD), 254
- singular vector beamforming, 254
- SINR, *see* Signal-to-Interference plus Noise Ratio
- SintraSearch, 81
- SI-RNTI, *see* System Information Radio Network Temporary Identifier
- SISO decoders, 225–229
- slot, 145
- SLP, *see* SUPL Location Platform
- small adjacent interferer, 483
- SN, *see* Sequence Number
- SnonintraSearch, 81, 521
- Snonintrasearch, 521
- SO, *see* Segmentation Offset
- soft buffer
 - size, 242–243
 - statistical management, 242
- soft frequency reuse, 289
- SON, *see* Self-Optimizing Network
- SORTD, *see* Space Orthogonal-Resource Transmit Diversity
- Sounding Reference Signal (SRS), 286, 327, 328, 337–340, 343, 561, 658–659
 - antenna selection, 368
 - aperiodic, 658–659
 - bandwidth, 340
 - carrier aggregation, 642
 - configuration, 341
 - parameters, 341
 - periodicity, 337
 - SRS-comb, 339
 - type 0, 658
 - type 1, 658
- S/P, *see* Serial-to-Parallel
- SPA, *see* Sum-Product Algorithm
- Space-Frequency Block Code (SFBC), 264–265, 659
- Space Orthogonal-Resource Transmit Diversity (SORTD), 667–668
- Space-Time Block Code (STBC), 257, 265, 666
- space-time processing, 250
- spatial
 - correlation, 175, 177, 438, 447–449
 - modelling, 442
 - domain channel estimation, 184–185
 - layer, 16, 253, 266, 618
 - multiplexing, 255–257
 - gain, 16, 251
 - schemes, 265–270
 - spatial bundling, 550, 638, 639
- Spatial Channel Model (SCM), 442–445
 - SCM-A, 180
- Spatial Channel Model – Extension (SCME), 442, 444–445
- Spatial Division Multiple Access (SDMA)
 - uplink, 330, 369
- special subframe, 147, 547–548
- spectral
 - efficiency, 8–10, 250, 298, 299, 603–608, 616, 617, 716
 - MBMS, 297
 - OFDM, 137, 139, 140
 - leakage, 132
- spectrum, 2, 12–13
 - emission mask, 468
 - base station, 576
 - Home eNodeB, 576
 - global, 620

- IMT, 620
 - paired, 531–562
 - refarming, 12
 - unpaired, 531–562
- spreading code
 - PUCCH, 358
- SPS, *see* Semi-Persistent scheduling
- SPS-C-RNTI, *see* Semi-Persistent Scheduling-C-RNTI
- spurious emissions, 467, 470–471, 487–489
 - base station, 575, 576
 - Home eNodeB, 575, 576
- spurious response, 487
- SR, *see* Scheduling Request
- SRB, *see* Signalling Radio Bearer
- SRIT, *see* Set of RITs
- SRS, *see* Sounding Reference Signal
- SRVCC, *see* Single Radio Voice Call Continuity
- Srxlev, 80
- SS7, 41
- SSAC, *see* Service Specific Access Control
- SSC, *see* Secondary Synchronization Code
- S-SCH, *see* Secondary Synchronization Channel
- SSS, *see* Secondary Synchronization Signal
- standardization, 2, 5–6
- Standards Development Organization (SDO), 5
- starvation, 117
- static interference coordination, 291
- STATUS PDU, 107
- status prohibit, 103–105
- status report, 102–105
- STBC, *see* Space-Time Block Code
- S-TMSI, *see* SAE-Temporary Mobile Subscriber Identity
- Stop-And-Wait (SAW), 108, 240
- Stream Control Transmission Protocol (SCTP), 41, 43, 297, 585, 586
 - Home eNodeB, 565
 - INIT message, 586
- study item, 7, 16
- sub-band
 - eNodeB-configured, 221
 - UE-selected, 221
- subcarrier, 14, 125, 127
 - leakage, 133
 - mapping, 344
 - spacing, 134, 139–141
 - SC-FDMA, 323, 345
 - unused, 142
- subframe, 145, 147, 192
 - bundling, 323, 347
- suburban
 - cell, 140
- successive decoding, 285
- Successive Interference Cancellation (SIC), 257, 266, 602
- SU-MIMO, *see* Single-User MIMO
- Sum-Product Algorithm (SPA), 231
- sum-rate, 261–264
- SUPL, *see* Secure User Plane Location
- SUPL Location Platform (SLP), 423, 433
- SVD, *see* Singular-Value Decomposition
- switched antenna diversity, 261, 560
- switching point alignment, 547
- SYNC protocol, 308–310
- SYNC-UL, 524
- synchronization, 151–163
 - error, 134, 135, 320, 498
 - loss
 - heterogeneous network, 707
 - MBMS content, 308–310
 - MBSFN, 297–298
 - signals, 151–163
 - timing
 - Home eNodeB, 579
- synchronized reconfiguration, 70
- synchronous HARQ, 109, 345
- System Architecture Evolution (SAE), 3, 25
- system bandwidth, 12, 141, 190, 197, 323, 350, 457, 461
- system capacity, 599–608
- System Frame Number (SFN), 61
- System Information (SI), 58–62, 110, 111, 189
 - change notification, 61–62
 - message, 60–61
 - relaying, 689
 - scheduling, 61
 - validity notification, 61–62
 - window, 61
- System Information Block (SIB), 59, 190, 196
 - SIB1, 60, 509
 - SIB2, 60, 83, 624
 - SIB3, 60
 - SIB4, 60
 - SIB5, 60

- SIB6, 60
- SIB7, 60
- SIB8, 60, 525
- SIB9, 60
- SIB10, 60
- SIB11, 60
- SIB12, 60, 312
- SIB13, 60, 304, 306, 307
- System Information Radio Network
 - Temporary Identifier (SI-RNTI), 196
- system performance, 599–611
 - LTE-Advanced, 715–718
- SystemInfoModification, 62

- T300, 66
- T301, 73
- T304, 72, 74
- T310, 526
- T311, 73
- TAC, *see* Tracking Area Code
- TACS, *see* Total Access Communication System
- tail-biting, 224
 - convolutional code, 366
- Tanner graph, 231–232
- Tapped Delay Line (TDL), 444, 446
 - model, 444, 446
- targets, 6–8, 11
- TCP, *see* Transmission Control Protocol
- TDL, *see* Tapped Delay Line
- TDMA, *see* Time Division Multiple Access
- Technical Specification Group (TSG), 5
- TEID, *see* Tunnelling End ID
- Telecommunications Technology Committee (TTC), 5
- Temporary Mobile Group Identifier (TMGI), 310
- test environment, 449–452, 599, 603–608
 - base coverage urban, 449
 - high speed, 449
 - indoor, 449
 - microcellular, 449
- TFT, *see* Traffic Flow Template
- thermal noise, 478, 493
 - figure, 478–479
 - base station, 576
 - Home eNodeB, 576
- Third Generation system, 2, 3
- third-order intercept point, 490
- third-order intermodulation products, 471
- throughput, 8–11, 284, 287, 288
 - requirement, 479
- tilting, 606
- Time Division Duplex (TDD), 4, 8, 147, 281, 505, 531, 561
 - PRACH, 404–405
- Time Division Multiple Access (TDMA), 3, 139, 251
- Time-Division Synchronous Code Division Multiple Access (TD-SCDMA), 4, 547
- Time To First Fix (TTFF), 425, 426
- time-domain channel estimation, 181–184
- time-domain offsets, 498
- Time-Spatial Propagation (TSP)
 - channel model, 450
- time-varying channels, 133–135
- timer
 - barring, 84
 - periodicBSR, 114
 - retxBSR, 114
 - T300, 66
 - T301, 73
 - T304, 72, 74
 - T310, 526
 - T311, 73
- timeToTrigger, 76
- timing
 - advance, 115, 337, 408–411, 431, 546
 - carrier aggregation, 642
 - initial, 408–409
 - Type 1 measurement, 431–432
 - Type 2 measurement, 431–432
 - update, 409–411
 - update commands, 410
 - alignment, 115
 - errors, 409
 - estimation, 385, 386, 393, 402
 - offset, 135–137
 - synchronization, 127, 136, 297
 - Home eNodeB, 579
- TM, *see* Transparent Mode
- TMD, *see* Transparent Mode Data
- TMGI, *see* Temporary Mobile Group Identifier
- TNL, *see* Transport Network Layer
- Total Access Communication System (TACS), 1

- Total Power De-rating (TPD), 473
- TPD, *see* Total Power De-rating
- TPMI, *see* Transmitted Precoding Matrix Indicator
- trace
 - cell traffic, 48
 - collection entity, 48
 - depth, 48
 - Failure Indication, 48
 - function, 48
 - management-based, 48, 711
 - signalling-based, 48, 711
 - identity, 48
 - session, 48
- Tracking Area (TA), 585
 - update, 29
- Tracking Area Code (TAC), 582
- Traffic Flow Template (TFT), 28, 36
 - DownLink (DL TFT), 36
 - UpLink (UL TFT), 36
- traffic logical channel, 110
- Training Sequence Code (TSC), 511
- transient output power, 466
- transmission bandwidth, 460–461
- transmission bandwidth configuration, 460–461
- Transmission Control Protocol (TCP), 41, 91
 - port number, 36
 - transmit window, 96
- transmission mode, 98
 - PDSCH, 193
 - transmission mode 9, 659–665
 - PUSCH, 666
- transmission resources, 145–147, 207–211
- Transmission Time Interval (TTI), 192
 - bundling, 323, 347
- transmit diversity, 193, 264–274
 - antenna switching, 560
 - PUCCH, 667–668
 - SORTD, 668
 - uplink, 667–668
- transmit signal leakage, 475–476
- transmit signal quality, 463
- transmit spatial signature, 252
- transmitted output power, 464–466
- Transmitted Precoding Matrix Indicator (TPMI), 273
- transmitter power control, 411–420
- Transparent Mode (TM), 98–99
- Transparent Mode Data (TMD)
 - PDU format, 105
- transport block, 108, 192, 266
 - size, 347–348
 - 328 bits, 347
- transport channel, 108, 110–111
- transport channels, 149
 - uplink, 344
- transport load, 588
- Transport Network Layer (TNL), 43, 50
- Truncated Normalized Mean Squared Error, 180
- TSC, *see* Training Sequence Code (TSC)
- TSG, *see* Technical Specification Group
- TTC, *see* Telecommunications Technology Committee
- TTFF, *see* Time To First Fix
- TTI, *see* Transmission Time Interval
- Tunnelling End ID (TEID), 42
- turbo codes, 216, 229–232
- Typical Urban (TU) channel model, 440
- UCI, *see* Uplink Control Information
- UDP, *see* User Datagram Protocol
- UE-specific Reference Signals (RSs), 167, 171–174, 193, 270
 - Release 8, 171
 - Release 9, 171–174
 - Release 10, 652–653
 - scrambling, 661
- UE-specific RSs, 561
- UE-specific search space, 211, 212
- UE-to-UE interference, 539–543
- UL Index, 551
- UL-SCH, *see* Uplink Shared Channel
- ULA, *see* Uniform Linear Array
- UM, *see* Unacknowledged Mode
- UMa, *see* Urban Macrocell
- UMi, *see* Urban Microcell
- UMTS, *see* Universal Mobile Telecommunications System
- Un interface, 687
 - reconfiguration, 689
- Unacknowledged Mode (UM), 93, 98–101
- Unacknowledged Mode Data (UMD)
 - PDU format, 105–106
- Uncorrelated-Scattered, 175
- unicast, 294, 505
- Uniform Linear Array (ULA), 663
- Unitary Precoding (UP), 254, 269

- Universal Mobile Telecommunications System (UMTS), 2, 3, 5, 141, 215, 293, 457
 - coherent demodulation, 166
 - inter-working, 38
- Universal Subscriber Identity Module (USIM), 63
- Universal Terrestrial Radio Access Network (UTRAN), 13
- unpaired spectrum, 12, 147, 531–562
- unwanted emissions, 467–471
- uplink, 317–420
 - Control Information (UCI), 348–364
 - PUSCH, 667
 - grant messages, 112, 204, 206
 - index, 551
 - MIMO, 201, 259, 330, 368
 - power control, 411–420
 - carrier aggregation, 642–644
 - reference signals, 327–342, 657–659
 - shared channel, 344–348, 364–367
 - Sounding Reference Signals (SRS), 337–340, 658–659
 - synchronization, 371
 - timing, 371, 407–411
- Uplink Pilot Timeslot (UpPTS), 547, 548, 556
- Uplink Shared Channel (UL-SCH), 111, 344–348, 365–367
- Uplink Time Difference Of Arrival (UTDOA), 436
- UpPTS, *see* Uplink Pilot Timeslot
- Urban Macrocell scenario (UMa), 450, 603–608
- Urban Microcell scenario (UMi), 450, 603–608
- User Datagram Protocol (UDP), 42, 91, 297
 - header compression, 91
- User Equipment (UE), 17–18, 206
 - capability, 17–18, 77
 - transfer, 77
 - category, 17–18
 - LTE-Advanced, 621
 - context, 29, 31, 43, 58, 71
 - Context Release, 43
 - information, 58
 - Initial Context Setup, 43
 - Initial Context Setup Request, 43
 - management, 43
 - historical information, 54
 - identity, 206
 - power class, 465
 - selected sub-band, 221
 - specific search space, 211
- user plane, 32–33, 40, 42–43, 49, 87–120
 - carrier aggregation, 628–631
 - data discard, 89, 95–97
 - latency, 616, 617
 - timeout, 89
- USIM, *see* Universal Subscriber Identity Module
- UTDOA, *see* Uplink Time Difference Of Arrival
- UTRA High Chip Rate TDD, 549
- UTRA Low Chip Rate TDD, 549
- Uu interface, 35
 - bearer, 35
- Virtual Circular Buffer (VCB), 240
- virtual MIMO, 259, 330, 368
- virtual power headroom report, 644
- Viterbi Algorithm (VA), 226
- Voice over IP (VoIP), 10, 26, 91, 99, 113, 194, 198, 347, 413
 - capacity, 10, 608, 616, 617
 - TTI bundling, 323
- Walsh–Hadamard
 - codes, 200, 356–359, 652, 653
 - matrices, 161
 - sequences, 200, 356–359
 - transform, 161
- warning
 - Area List, 48
 - message, 48–49
 - Write-Replace, 48
- water-filling, 254, 255, 284
- waveform power de-rating, 473
- WCDMA, *see* Wideband Code Division Multiple Access
- WFT, *see* Winograd Fourier Transform
- WG, *see* Working Group
- Wide-Sense Stationary Uncorrelated Scattering (WSSUS), 175
- Wideband Code Division Multiple Access (WCDMA), 3, 4, 156, 196, 285, 407, 411, 471
- WiFi, 13
- WiMAX, 4, 13, 547
 - channel coding, 234
 - inter-working, 28, 38
 - phase noise, 497

rate matching, 238
 WINNER, 437, 444
 channel model, 445
 Winograd Fourier Transform (WFT), 124
 Wireless Local Area Network (WLAN), 4,
 124, 139, 169
 WLAN, *see* Wireless Local Area Network
 Working Group (WG), 5
 World Radiocommunication Conference
 (WRC), 437, 615, 620
 WRC, *see* World Radiocommunication
 Conference
 WSSUS, *see* Wide-Sense Stationary
 Uncorrelated Scattering
 X2 interface, 30, 31, 49–55, 584
 ABS signalling, 704–706
 application protocol, 49
 Handover Request, 55
 handover, 45, 50–53, 586
 HeNB–HeNB, 566–568
 Setup Request, 50
 Status Transfer, 45, 52

HeNB–HeNB, 573
 ICIC signalling, 290, 291, 669, 704–706
 initiation, 49–50
 Load Indication procedure, 705
 mobility, 50–54
 Resource Status Reporting Initiation
 procedure, 705
 Setup procedure, 586
 Setup Request, 586, 587
 Setup Response, 587
 Zadoff–Chu (ZC) sequences, 155–158,
 328–330, 385–394
 cyclic shift, 388, 390
 high-speed cell configuration,
 393–394
 high-speed cell restriction, 391–393
 ZC, *see* Zadoff–Chu (ZC) sequences
 Zero Correlation Zone (ZCZ), 156, 389
 zero-forcing, 255–262
 equalizer, 463
 zero-outage capacity, 285

Attachment 3-B

LIBRARY | CATALOG

← 1 of 1 →

BOOK

LTE--the UMTS long term evolution : from theory to practice

[Full Record](#) [MARC Tags](#)

Personal name

Sesia, Stefania.

Main title

LTE--the UMTS long term evolution : from theory to practice / Stefania Sesia, Issam Toufik, Matthew Baker.

Edition

2nd ed.

Published/Created

Chichester, West Sussex, U.K. ; Hoboken, N.J. : Wiley, 2011.

[Request this Item](#)

[LC Find It](#)

Links >

Links

Cover image

<http://catalogimages.wiley.com/images/db/jimages/9780470660256.jpg>

More Information >

LCCN Permalink

<https://lcn.loc.gov/2010039466>

2/5/2021

LC Catalog - Item Information (Full Record)

Description

xl, 752 p. : ill. ; 25 cm.

ISBN

9780470660256 (hardback)

LC classification

TK5103.4883 .S47 2011

Related names

[Toufik, Issam.](#)

[Baker, Matthew \(Matthew P. J.\)](#)

Summary

""Where this book is exceptional is that the reader will not just learn how LTE works but why it works" Adrian Scrase, ETSI Vice-President, International Partnership Projects Following on the success of the first edition, this book is fully updated, covering the latest additions to LTE and the key features of LTE-Advanced. This book builds on the success of its predecessor, offering the same comprehensive system-level understanding built on explanations of the underlying theory, now expanded to include complete coverage of Release 9 and the developing specifications for LTE-Advanced. The book is a collaborative effort of more than 40 key experts representing over 20 companies actively participating in the development of LTE, as well as academia. The book highlights practical implications, illustrates the expected performance, and draws comparisons with the well-known WCDMA/HSPA standards. The authors not only pay special attention to the physical layer, giving an insight into the fundamental concepts of OFDMA-FDMA and MIMO, but also cover the higher protocol layers and system architecture to enable the reader to gain an overall understanding of the system. Key New Features: Comprehensively updated with the latest changes of the LTE Release 8 specifications, including improved coverage of Radio Resource Management RF aspects and performance requirements. Provides detailed coverage of the new LTE Release 9 features, including: eMBMS, dual-layer beamforming, user equipment positioning, home eNodeBs / femtocells and pico cells and self-optimizing networks. Evaluates the LTE system performance. Introduces LTE-Advanced, explaining its context and motivation, as well as the key new features including: carrier aggregation, relaying, high-order MIMO, and Cooperative Multi-Point transmission (CoMP). Includes an accompanying website containing a complete list of acronyms related to LTE and LTE-Advanced, with a brief description of each This book is an invaluable reference for all research and development engineers involved in implementation of LTE or LTE-Advanced, as well as graduate and PhD students in wireless communications. Network operators, service providers and R&D managers will also find this book insightful."-- Provided by publisher.

"LTE -- The UMTS Long Term Evolution, Second Edition builds on the success of its predecessor, offering an updated treatment of the new LTE standard"-- Provided by publisher.

2/5/2021

LC Catalog - Item Information (Full Record)

LC Subjects

[Universal Mobile Telecommunications System.
Long-Term Evolution \(Telecommunications\)](#)

Other Subjects

TECHNOLOGY & ENGINEERING / Telecommunications

Browse by shelf order

[TK5103.4883](#)

Notes

Includes bibliographical references and index.

LCCN

2010039466

Dewey class no.

621.3845/6

Other class no.

TEC041000

Type of material

Book

Item Availability >

Shelf Location

FLM2016 049281

CALL NUMBER

[TK5103.4883 .S47 2011 OVERFLOWJ34](#)

Request in

Jefferson or Adams Building Reading Rooms (FLM2)

Status

Not Charged

Attachment 3-C

LIBRARY CATALOG

← 1 of 1 →

BOOK

LTE--the UMTS long term evolution : from theory to practice

Full Record

MARC Tags

000 03945cam a2200361 a 4500
 001 16458224
 005 20170313085135.0
 008 100914s2011 enka b 001 0 eng
 906 _ |a 7 |b cbc |c orignew |d 1 |e ecip |f 20 |g y-gencatlg
 925 0_ |a acquire |b 1 shelf copy |x policy default
 925 1_ |a acquire |b 2 shelf copies |x policy default |e claim1 2011-10-19
 955 _ |b xj12 2010-09-14 |c xj12 2010-09-14 ONIX (telework) to STM |w rd11 2010-10-19 |a xn02 2011-11-21 2 copies rec'd., to CIP ver. |f xj16 2012-04-10 copy 1 and 2 to BCCD
 010 _ |a 2010039466
 020 _ |a 9780470660256 (hardback)
 040 _ |a DLC |c DLC |d DLC
 042 _ |a pcc
 050 00 |a TK5103.4883 |b .S47 2011
 082 00 |a 621.3845/6 |2 22
 084 _ |a TEC041000 |2 bisacsh
 100 1_ |a Sesia, Stefania.
 245 10 |a LTE--the UMTS long term evolution : |b from theory to practice / |c Stefania Sesia, Issam Toufik, Matthew Baker.
 250 _ |a 2nd ed.
 260 _ |a Chichester, West Sussex, U.K. ; |a Hoboken, N.J. : |b Wiley, |c 2011.
 300 _ |a xl, 752 p. : |b ill. ; |c 25 cm.

- 520 — |a ""Where this book is exceptional is that the reader will not just learn how LTE works but why it works" Adrian Scrase, ETSI Vice-President, International Partnership Projects Following on the success of the first edition, this book is fully updated, covering the latest additions to LTE and the key features of LTE-Advanced. This book builds on the success of its predecessor, offering the same comprehensive system-level understanding built on explanations of the underlying theory, now expanded to include complete coverage of Release 9 and the developing specifications for LTE-Advanced. The book is a collaborative effort of more than 40 key experts representing over 20 companies actively participating in the development of LTE, as well as academia. The book highlights practical implications, illustrates the expected performance, and draws comparisons with the well-known WCDMA/HSPA standards. The authors not only pay special attention to the physical layer, giving an insight into the fundamental concepts of OFDMA-FDMA and MIMO, but also cover the higher protocol layers and system architecture to enable the reader to gain an overall understanding of the system. Key New Features: Comprehensively updated with the latest changes of the LTE Release 8 specifications, including improved coverage of Radio Resource Management RF aspects and performance requirements. Provides detailed coverage of the new LTE Release 9 features, including: eMBMS, dual-layer beamforming, user equipment positioning, home eNodeBs / femtocells and pico cells and self-optimizing networks. Evaluates the LTE system performance. Introduces LTE-Advanced, explaining its context and motivation, as well as the key new features including: carrier aggregation, relaying, high-order MIMO, and Cooperative Multi-Point transmission (CoMP). Includes an accompanying website containing a complete list of acronyms related to LTE and LTE-Advanced, with a brief description of each This book is an invaluable reference for all research and development engineers involved in implementation of LTE or LTE-Advanced, as well as graduate and PhD students in wireless communications. Network operators, service providers and R&D managers will also find this book insightful."-
- |c Provided by publisher.
- 520 — |a "LTE -- The UMTS Long Term Evolution, Second Edition builds on the success of its predecessor, offering an updated treatment of the new LTE standard"-- |c Provided by publisher.
- 504 — |a Includes bibliographical references and index.
- 650 _0 |a Universal Mobile Telecommunications System.
- 650 _0 |a Long-Term Evolution (Telecommunications)
- 650 _7 |a TECHNOLOGY & ENGINEERING / Telecommunications |2 bisacsh.
- 700 1_ |a Toufik, Issam.
- 700 1_ |a Baker, Matthew |q (Matthew P. J.)
- 856 42 |3 Cover image |u <http://catalogimages.wiley.com/images/db/jimages/9780470660256.jpg>

[Request this Item](#)
[LC Find It](#)

Item Availability

Shelf Location

FLM2016 049281

<https://catalog.loc.gov/vwebv/staffView?searchId=18997&recPointer=0&recCount=25&searchType=2&bibId=16458224>

2/3

2/5/2021

LC Catalog - Item Information (MARC Tags)

CALL NUMBER

TK5103.4883 .S47 2011 OVERFLOWJ34

Request in

Jefferson or Adams Building Reading Rooms (FLM2)

Status

Not Charged

Attachment 3-D

Search WorldCat

Search

[Advanced Search](#) [Find a Library](#)

[<< Return to Search Results](#)

[Cite/Export](#)

[Print](#)

[E-mail](#)

[Share](#)

[Permalink](#)

[Add to list](#)

[Add tags](#)

[Write a review](#)

Rate this item: 1 2 3 4 5




[Preview this item](#)

LTE - the UMTS long term evolution : from theory to practice

Author: [Stefania Sesia](#); [Matthew Barker](#); [Issam Toufik](#)

Publisher: Oxford : Wiley Blackwell, 2011.

Edition/Format:  Print book : English : Second ed [View all editions and formats](#)

Summary: This book builds on the success of its predecessor, offering the same comprehensive system-level understanding built on explanations of the underlying theory, now expanded to include complete [Read more...](#)

Rating: (not yet rated) [0 with reviews - Be the first.](#)

Subjects: [WIDEBAND CODE-DIVISION MULTIPLE ACCES, WCDMA \(TELECOMMUNICATIONS\)](#)
[UNIVERSAL MOBILE TELECOMMUNICATION SYSTEMS, UMTS \(MOBILE COMMUNICATIONS\)](#)
[MOBILE RADIO NETWORKS + WIRELESS COMMUNICATIONS NETWORKS \(TELECOMMUNICATIONS\)](#)

[View all subjects](#)

More like this [Similar Items](#)

Get a Copy

[Find a copy in the library.](#)

[AbeBooks](#) \$49.10

[Amazon](#) \$74.03

[Better World Books](#) \$64.52

Find a copy in the library

Enter your location:

Submit a complete postal address for best results.

Displaying libraries 1-6 out of 770 for all 42 editions (Washington, DC, USA)

Show libraries holding [just this edition](#)

« First < Prev 1 2 3 Next > Last »

Library	Held formats	Distance	
1. Consumer Financial Protection Bureau Washington, DC 20552 United States	 Book	1 mile MAP IT	Library info Add to favorites
2. Howard University Washington, DC 20059 United States	 Book	1 mile MAP IT	Library info Ask a librarian Add to favorites
3. Research Center, National Academies of Sciences, Engineering, and Medicine Washington, DC 20001 United States	 Book	1 mile MAP IT	Library info Add to favorites
4. US Department of Justice Library Washington, DC 20530 United States	 Book	1 mile MAP IT	Library info Add to favorites

- | | | | | |
|----|--|--|-------------------|---|
| 5. | Federal Communications Commission
Washington, DC 20554 United States |  Book | 2 miles
MAP IT | Library info
Add to favorites |
| 6. | Library of Congress
Washington, DC 20540 United States |  Book | 2 miles
MAP IT | Library info
Ask a librarian
Add to favorites |

« First < Prev 1 2 3 Next > Last »

Details

Document Type: Book

All Authors / Contributors: [Stefania Sesia](#); [Matthew Barker](#); [Issam Toufik](#)

Find more information about:

ISBN: 9780470660256 0470660252

OCLC Number: 745962607

Description: 752 Seiten

Contents: Editors Biographies List of Contributors Foreword Preface Acknowledgements List of Acronyms 1 Introduction and Background 1 Thomas Salzer and Matthew Baker 1.1 The Context for the Long Term Evolution of UMTS 1 1.2 Requirements and Targets for the Long Term Evolution 7 1.3 Technologies for the Long Term Evolution 14 1.4 From Theory to Practice 20 References 21 Part I Network Architecture and Protocols 23 2 Network Architecture 25 Sudeep Palat and Philippe Godin 2.1 Introduction 25 2.2 Overall Architectural Overview 26 2.3 Protocol Architecture 32 2.4 Quality of Service and EPS Bearers 34 2.5 The E-UTRAN Network Interfaces: S1 Interface 40 2.6 The E-UTRAN Network Interfaces: X2 Interface 49 2.7 Summary 55 References 55 3 Control Plane Protocols 57 Himke van der Velde 3.1 Introduction 57 3.2 Radio Resource Control (RRC) 58 3.3 PLMN and Cell Selection 78 3.4 Paging 84 3.5 Summary 86 References 86 4 User Plane Protocols 87 Patrick Fischer, SeungJune Yi, SungDuck Chun and YoungDae Lee 4.1 Introduction to the User Plane Protocol Stack 87 4.2 Packet Data Convergence Protocol (PDCP) 89 4.3 Radio Link Control (RLC) 98 4.4 Medium Access Control (MAC) 108 4.5 Summary of the User Plane Protocols 120 References 120 Part II Physical Layer for Downlink 121 5 Orthogonal Frequency Division Multiple Access (OFDMA) 123 Andrea Ancora, Issam Toufik, Andreas Bury and Dirk Slock 5.1 Introduction 123 5.2 OFDM 125 5.3 OFDMA 137 5.4 Parameter Dimensioning 139 5.5 Summary 142 References 142 6 Introduction to Downlink Physical Layer Design 145 Matthew Baker 6.1 Introduction 145 6.2 Transmission Resource Structure 145 6.3 Signal Structure 148 6.4 Introduction to Downlink Operation 149 References 150 7 Synchronization and Cell Search 151 Fabrizio Tomatis and Stefania Sesia 7.1 Introduction 151 7.2 Synchronization Sequences and Cell Search in LTE 151 7.3 Coherent Versus Non-Coherent Detection 161 References 163 8 Reference Signals and Channel Estimation 165 Andrea Ancora, Stefania Sesia and Alex Gorokhov 8.1 Introduction 165 8.2 Design of Reference Signals in the LTE Downlink 167 8.2.1 Cell-Specific Reference Signals 168 8.3 RS-Aided Channel Modelling and Estimation 174 8.4 Frequency-Domain Channel Estimation 178 8.5 Time-Domain Channel Estimation 181 8.6 Spatial-Domain Channel Estimation 184 8.7 Advanced Techniques 185 References 186 9 Downlink Physical Data and Control Channels 189 Matthew Baker and Tim Mousley 9.1 Introduction 189 9.2 Downlink Data-Transporting Channels 189 9.3 Downlink Control Channels 196 References 214 10 Link Adaptation and Channel Coding 215 Brian Classon, Ajit Nimbalkar, Stefania Sesia and Issam Toufik 10.1 Introduction 215 10.2 Link Adaptation and CQI Feedback 217 10.3 Channel Coding 223 10.4 Conclusions 245 References 246 11 Multiple Antenna Techniques 249 Thomas Salzer, David Gesbert, Cornelius van Rensburg, Filippo Tosato, Florian Kaltenberger and Tetsushi Abe 11.1 Fundamentals of Multiple Antenna Theory 249 11.2 MIMO Schemes in LTE 262 11.3 Summary 276 References 277 12 Multi-User Scheduling and Interference Coordination 279 Issam Toufik and Raymond Knopp 12.1 Introduction 279 12.2 General Considerations for Resource Allocation Strategies 280 12.3 Scheduling Algorithms 283 12.4 Considerations for Resource Scheduling in LTE 286 12.5 Interference Coordination and Frequency Reuse 287 12.6 Summary 291 References 293 13 Broadcast Operation 293 Himke van der Velde, Olivier Hus and Matthew Baker 13.1 Introduction 293 13.2 Broadcast Modes 293 13.3 Overall MBMS Architecture 295 13.4 MBMS Single Frequency Network Transmission 297 13.5 MBMS Characteristics 303 13.6 Radio Access Protocol Architecture and Signalling 304 13.7 Public Warning Systems 312 13.8 Comparison of Mobile Broadcast Modes 312 References 314 Part III Physical Layer for Uplink 315 14 Uplink Physical Layer Design 317 Robert Love and Vijay Nangia 14.1 Introduction 317 14.2 SC-FDMA Principles 318 14.3 SC-FDMA Design in LTE 321 14.4 Summary 325

References 326 15 Uplink Reference Signals 327 Robert Love and Vijay Nangia 15.1 Introduction 327 15.2 RS Signal Sequence Generation 328 15.3 Sequence-Group Hopping and Planning 332 15.4 Cyclic Shift Hopping 333 15.5 Demodulation Reference Signals (DM-RS) 335 15.6 Uplink Sounding Reference Signals (SRS) 337 15.7 Summary 340 References 341 16 Uplink Physical Channel Structure 343 Robert Love and Vijay Nangia 16.1 Introduction 343 16.2 Physical Uplink Shared Data Channel Structure 344 16.3 Uplink Control Channel Design 348 16.4 Multiplexing of Control Signalling and UL-SCH Data on PUSCH 365 16.5 ACK/NACK Repetition 367 16.6 Multiple-Antenna Techniques 367 16.7 Summary 369 References 369 17 Random Access 371 Pierre Bertrand and Jing Jiang 17.1 Introduction 371 17.2 Random Access Usage and Requirements in LTE 371 17.3 Random Access Procedure 372 17.4 Physical Random Access Channel Design 376 17.5 PRACH Implementation 396 17.6 Time Division Duplex (TDD) PRACH 404 17.7 Concluding Remarks 405 References 406 18 Uplink Transmission Procedures 407 Matthew Baker 18.1 Introduction 407 18.2 Uplink Timing Control 407 18.3 Power Control 411 References 420 Part IV Practical Deployment Aspects 421 19 User Equipment Positioning 423 Karri Ranta-aho and Zukang Shen 19.1 Introduction 423 19.2 Assisted Global Navigation Satellite System (A-GNSS) Positioning 425 19.3 Observed Time Difference Of Arrival (OTDOA) Positioning 426 19.4 Cell-ID-based Positioning 431 19.5 LTE Positioning Protocols 433 19.6 Summary and Future Techniques 435 References 436 20 The Radio Propagation Environment 437 Juha Ylitalo and Tommi Jamsa 20.1 Introduction 437 20.2 SISO and SIMO Channel Models 438 20.3 MIMO Channel Models 441 20.4 Radio Channel Implementation for Conformance Testing 454 20.5 Concluding Remarks 455 References 455 21 Radio Frequency Aspects 457 Moray Rumney, Takaharu Nakamura, Stefania Sesia, Tony Sayers and Adrian Payne 21.1 Introduction 457 21.2 Frequency Bands and Arrangements 459 21.3 Transmitter RF Requirements 462 21.4 Receiver RF Requirements 474 21.5 RF Impairments 492 21.6 Summary 500 References 501 22 Radio Resource Management 503 Muhammad Kazmi 22.1 Introduction 503 22.2 Cell Search Performance 505 22.3 Mobility Measurements 513 22.4 UE Measurement Reporting Mechanisms and Requirements 516 22.5 Mobility Performance 518 22.6 RRC Connection Mobility Control Performance 525 22.7 Radio Link Monitoring Performance 526 22.8 Concluding Remarks 528 References 529 23 Paired and Unpaired Spectrum 531 Nicholas Anderson 23.1 Introduction 531 23.2 Duplex Modes 532 23.3 Interference Issues in Unpaired Spectrum 533 23.4 Half-Duplex System Design Aspects 544 23.5 Reciprocity 552 24 Picocells, Femtocells and Home eNodeBs 563 Philippe Godin and Nick Whinnett 24.1 Introduction 563 24.2 Home eNodeB Architecture 564 24.3 Interference Management for Femtocell Deployment 569 24.4 RF Requirements for Small Cells 574 24.5 Summary 580 References 580 25 Self-Optimizing Networks 581 Philippe Godin 25.1 Introduction 581 25.2 Automatic Neighbour Relation Function (ANRF) 582 25.3 Self-Configuration of eNodeB and MME 584 25.4 Automatic Configuration of Physical Cell Identity 587 25.5 Mobility Load Balancing Optimization 587 25.6 Mobility Robustness Optimization 591 25.7 Random Access Channel (RACH) Self-Optimization 595 25.8 Energy Saving 596 25.9 Emerging New SON Use Cases 597 References 598 26 LTE System Performance 599 Tetsushi Abe 26.1 Introduction 599 26.2 Factors Contributing to LTE System Capacity 599 26.3 LTE Capacity Evaluation 603 26.4 LTE Coverage and Link Budget 608 26.5 Summary 610 References 611 Part V LTE-Advanced 613 27 Introduction to LTE-Advanced 615 Dirk Gerstenberger 27.1 Introduction and Requirements 615 27.2 Overview of the Main Features of LTE-Advanced 618 27.3 Backward Compatibility 619 27.4 Deployment Aspects 620 27.5 UE Categories for LTE-Advanced 621 References 622 28 Carrier Aggregation 623 Juan Montojo and Jelena Damnjanovic 28.1 Introduction 623 28.2 Protocols for Carrier Aggregation 624 28.3 Physical Layer Aspects 631 28.4 UE Transmitter and Receiver Aspects 648 28.5 Summary 650 References 650 29 Multiple Antenna Techniques for LTE-Advanced 651 Alex Gorokhov, Amir Farajdana, Kapil Bhattacharya, Xiliang Luo and Stefan Geirhofer 29.1 Downlink Reference Signals 651 29.2 Uplink Reference Signals 657 29.3 Downlink MIMO Enhancements 659 29.4 Uplink Multiple Antenna Transmission 666 29.5 Coordinated MultiPoint (CoMP) Transmission and Reception 669 29.6 Summary 671 References 671 30 Relaying 673 Eric Hardouin, J. Nicholas Laneman, Alexander Golitschek, Hidetoshi Suzuki, Osvaldo Gonsa 30.1 Introduction 673 30.2 Theoretical Analysis of Relaying 679 30.3 Relay Nodes in LTE-Advanced 684 30.4 Summary 699 References 699 31 Additional Features of LTE Release 10 701 Teck Hu, Philippe Godin and Sudeep Palat 31.1 Introduction 701 31.2 Enhanced Inter-Cell Interference Coordination 701 31.3 Minimization of Drive Tests 710 31.4 Machine-Type Communications 712 References 714 32 LTE-Advanced Performance and Future Developments 715 Takehiro Nakamura and Tetsushi Abe 32.1 LTE-Advanced System Performance 715 32.2 Future Developments 718 References 720 Index 721

Responsibility:

edited by Stefania Sesia, Issam Toufik, Matthew Baker.

Reviews**User-contributed reviews**

[Add a review](#) and share your thoughts with other readers. Be the first.

- Tags

[Add tags](#) for "LTE - the UMTS long term evolution : from theory to practice". Be the first.

- Similar Items

Related Subjects: (9)

[WIDEBAND CODE-DIVISION MULTIPLE ACCES, WCDMA \(TELECOMMUNICATIONS\)](#)

[UNIVERSAL MOBILE TELECOMMUNICATION SYSTEMS, UMTS \(MOBILE COMMUNICATIONS\)](#)

[MOBILE RADIO NETWORKS + WIRELESS COMMUNICATIONS NETWORKS \(TELECOMMUNICATIONS\)](#)

[ACCÈS MULTIPLE À RÉPARTITION PAR CODE BANDE LARGE, AMRC LARGE BANDE \(TÉLÉCOMMUNICATIONS\)](#)

[BREITBAND-CODEMEHRFACHZUGRIFF, WCDMA \(NACHRICHTENTECHNIK\)](#)

[SYSTÈMES DE TÉLÉCOMMUNICATION MOBILES UNIVERSELS, UMTS \(TÉLÉCOMMUNICATIONS MOBILES\)](#)

[UNIVERSALE MOBILE TELEKOMMUNIKATIONSSYSTEME, UMTS \(MOBILKOMMUNIKATION\)](#)

[RÉSEAUX DE TÉLÉCOMMUNICATIONS MOBILES + RÉSEAUX DE TÉLÉCOMMUNICATIONS SANS FIL \(TÉLÉCOMMUNICATIONS\)](#)

[MOBILFUNKNETZE + DRAHTLOSE NACHRICHTENNETZE \(NACHRICHTENTECHNIK\)](#)

+ Linked Data