Autonomous Component Carrier Selection: Interference Management in Local Area Environments for LTE-Advanced

Luis G. U. Garcia, Aalborg University
Klaus I. Pedersen and Preben E. Mogensen, Nokia Siemens Networks

ABSTRACT

Low-power base stations such as femtocells are one of the candidates for high-data-rate provisioning in local areas, such as residences, apartment complexes, business offices, and outdoor hotspot scenarios. Unfortunately, the benefits are not without new challenges in terms of interference management and efficient system operation. Due to the expected large number of user-deployed cells, centralized network planning becomes impractical, and new scalable alternatives must be sought. In this article we propose a fully distributed and scalable solution to the interference management problem in local areas, basing our study case on LTE-Advanced. We present extensive network simulation results to demonstrate that a simple and robust interference management scheme, called autonomous component carrier selection, allows each cell to select the most attractive frequency configuration; improving the experience of all users and not just the few best ones, while overall cell capacity is not compromised.

INTRODUCTION

Low-power base stations, which are also commonly referred to as femtocells or home base stations, are low-cost user-deployed cellular base stations using an IP-based wired backhaul such as cable or digital subscriber line (DSL) designed to provide service in local environments similar to existing WiFi access points. In a recent contribution [1], the authors indicated the key benefits of low-power base stations and outlined the many research opportunities as well as technological and business challenges associated with femtocells. In [2] an interesting analysis of the financial impact of home base stations indicates that current macrocellular network deployment becomes less economically viable for increasing data rates.

In this light, low-power base stations have recently reemerged as a promising technology component, and many believe it will definitely be one of the next steps in the evolutionary path of cellular wireless systems. Dense deployment of low-power base stations offers significantly higher capacity per area than macrocells, arising from using smaller cell sizes and more efficient spatial reuse. On the other hand, installation of many low-power base stations also poses new challenges in terms of interference management and efficient system operation. The latter is especially the case for local areas where end users start installing home base stations without any prior network planning or carefully considering where other people in the immediate surroundings have installed other home base stations.

The vast majority of previous contributions in the literature focused on solutions for cases where the user-deployed cells use the same frequency band employed by macrocells, in which case capacity and coverage gains can dwindle away if macro/femtocell co-channel interference is left unchecked. Nonetheless, in [3] the authors point out that femto-to-femto interference also becomes an important issue for indoor performance, especially when femtocells are densely deployed. Therefore, we pay special attention to the nuances of interference footprint in local area deployments, and do not address the complementary and equally interesting case of co-channel interference to/from macrocells in overlaid networks.

As demonstrated in [4], the interference footprint is significantly different in such local area environments from nicely planned macrocell scenarios, which consequently calls for new self-adjusting interference management techniques. Early work found in [5, 6] also highlights the need for the ability to self-scale and self-adjust, leading to a new autonomic paradigm with fully "robotic" base stations. The optimal sharing of radio resources between low-power base stations depend on many factors such as the mutual interference coupling among them and the offered traffic for individual access nodes. Finding the optimal division of frequency resources between low-power base stations in a highly dynamic and partly chaotic environment is, in general, a nonlinear non-convex NP-hard optimization problem. Several interesting contributions are available in the literature, where decomposition of this challenging problem into subproblems and the use of heuristic algorithms are proposed [7–8].

As a case study, we base our investigations on LTE-Advanced, an evolved version of Long Term Evolution (LTE) Release 8, offering downlink peak data rates in excess of 1 Gb/s in a bandwidth of 100 MHz [9]. LTE-Advanced is currently in the study item phase in the Third Generation Partnership Project (3GPP), and design targets and new technology features for this system are also aimed at for local area scenarios. We propose a fully distributed and scalable solution based on minimal information exchange and negotiation between base stations akin to [10] where each individual low-power base station autonomously makes decisions without involving any centralized network control. The latter is considered to be the most attractive solution, especially for femto-type cells due to the expected large number of deployed cells. Our scheme mainly relies on measurements collected as a by-product of normal system operation, producing useful statistics for interference conditions in the network. In this way each base station gathers knowledge about the surrounding environment and uses this information in the decision making process. We present network simulation results to further demonstrate that a simple and robust interference management scheme, called autonomous component carrier selection, is possible for LTE-Advanced, providing attractive performance results in local area environments. Although the developed scheme is equally applicable for uplink and downlink, and for frequency-division duplex (FDD) and time-division duplex (TDD), we mainly present it for downlink TDD in this study.

The rest of the article is organized as follows. We present the system model and outline the basic assumptions for autonomous component carrier selection. We include more detailed algorithm descriptions and brief comments on the key distinguishing aspects of TDD and FDD deployments. System-level simulation results are presented for an extended local area residential scenario. Finally, the article is closed with concluding remarks and an outlook on future studies.

System Model

The 100 MHz LTE-Advanced bandwidth consists of five component carriers, each with a bandwidth of 20 MHz. The numerology of each component carrier is in coherence with LTE Release 8. The LTE-Advanced spectrum could also be less than 100 MHz, and therefore consist of less than five component carriers. The frequency band and spectrum allocation expressed via the number of component carriers and their bandwidth are configurable and known a priori by all base stations, hereafter denoted eNBs to follow 3GPP terminology. An LTE-Advanced terminal (user equipment [UE]) can be jointly scheduled on multiple component carriers at the same time (i.e., using carrier aggregation) or on a single component carrier as in LTE Release 8.

We assume that each eNB always has one active component carrier, denoted the primary component carrier (PCC). The PCC is automatically selected by the eNB when it is first switched on, and is assumed to provide full cell coverage as it will be used by the terminals to camp, set up new calls, and so on. Depending on the offered traffic in the cell and mutual interference coupling with surrounding cells, transmission and/or reception on all component carriers may not always be the best solution, especially for cell edge users. It is therefore proposed that each cell dynamically selects additional component carriers for transmission/reception as well (i.e., a second step after having selected the PCC). The latter is referred to as selection of secondary component carriers (SCCs). All component carriers not selected are assumed to be completely muted (uplink/downlink) and not used by the cell.

The proposed scheme uses a distributed and fully scalable approach. That is, selection of primary and secondary carriers is done locally by each cell. Hence, in the proposed concept there is no need for centralized network control.

The proposed scheme uses a distributed and fully scalable approach. That is, selection of primary and secondary carriers is done locally by each cell. Hence, in the proposed concept there is no need for centralized network control. The suggested interference coordination mechanism is part of a hierarchical resource management process. The (re-)selection of component carriers is fairly slow and occurs over a longer time span than fast packet scheduling, which is free to operate within the restrictions imposed by the carrier selection process. Our three fundamental premises are:

- Absolute priority of primary over secondary component carriers; avoidance of PCC reselection, while SCCs can be reselected on a faster basis.
- When the offered traffic for an eNB requires more bandwidth, a cell may augment its cell capacity by allocating SCCs.
- An eNB is only allowed to allocate SCCs provided it does not result in excessive interference to the surrounding cells, as explained later.

The last item is a policy preventing a so-called greedy eNB from using all the available component carriers for its own sake, even when this results in intolerable interference to the neighboring eNBs. Hence, the proposed scheme for autonomous component carrier selection effectively provides an automatic frequency reuse scheme at component carrier resolution. This approach ensures protection of both traffic and control channels.

We assume that the allocation of PCC and SCCs is signaled among eNBs (either over the backhaul or over the air) periodically and/or whenever the allocation is changed, so eNBs know which component carriers neighboring eNBs are currently using. This information is of critical importance and is summarized in what we refer henceforth as the Radio Resource Allocation Table (RRAT). Essentially, such tables make femtocells aware of the existence of other femtocells. Finally, it is assumed that local eNB measurements are available, as well as terminal measurements for selection of the component carriers. The next section on selection of the PCC deals with the first premise, whereas the SCC selection scheme described later embodies the other two assumptions.

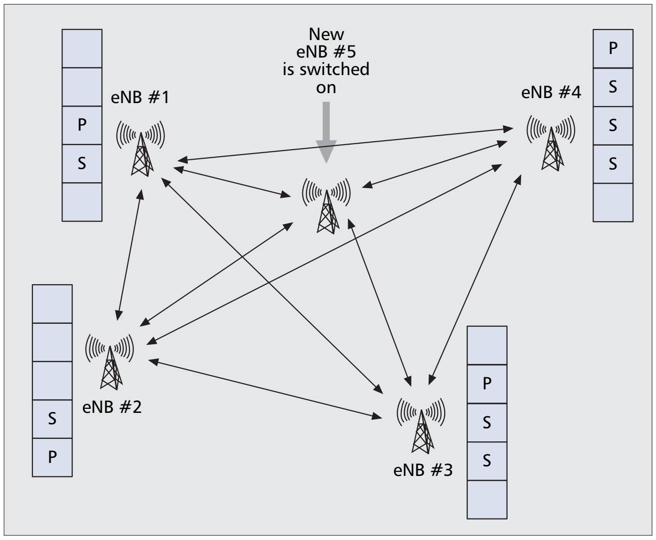

■ **Figure 1.** *Simple illustration of the autonomous component carrier concept. All eNBs announce their existence and current resource allocation. Additionally, eNBs that are being switched off could signal their leaving.*

Primary Component Carrier Selection

The proposed autonomous component carrier selection scheme is illustrated in Fig. 1 with a simple example. Here there are four existing eNBs, while a new eNB, #5, is being switched on, and hence is ready for first selecting its PCC. The current selection of PCC and SCCs is illustrated for each eNB with P and S, respectively. Component carriers not allocated as PCC or SCC are completely muted, and not used to carry any traffic.

As the eNB is being initialized, it clearly cannot rely on UE assisted mechanisms; therefore, in addition to the information available in the RRAT, we propose new inter-eNB measurements based on reference signal received power levels for the purpose of estimating the path loss between neighboring eNBs. In FDD systems this implies that eNBs are able to listen to the downlink band as well. Conversely, in TDD systems, this is not an additional requirement, since uplink and downlink use the same band. It is proposed that the new eNB carry out the measurements on the PCCs of the surrounding cells and that knowledge of their corresponding reference symbol transmit power is available (signaled between eNBs) so that the inter-eNB path loss can be estimated. Notice that these inter-eNB path loss measurements need not be frequent as they are only required by new eNBs when they are switched on.

Given the aforementioned information, a matrix for initial PCC selection is formed as illustrated in Fig. 2, where the eNBs are sorted according to the path loss experienced from the new eNB. As depicted in Fig. 2, only the neighboring eNBs within a certain path loss threshold are considered relevant. Neighboring eNBs with higher path loss are not taken into account as there is marginal interference coupling with those. Based on this matrix, we propose the following procedure for initial primary component carrier selection:

1 If there are row entries in the matrix with no selections, the corresponding component carrier is selected. (If there are multiple of such rows, either select randomly, or select the component experiencing the lowest uplink received interference power.) Otherwise, go to 2.
2 If there are row entries without P, select one of those for primary. Select the row entry with the lowest number of S if there are multiple rows without P.
3 If all row entries include P, select the component carrier for primary with maximum path loss to the neighboring eNB having the same component carrier as its primary.
4 When there are multiple candidate component carriers for primary according to the above rules, select the component carrier with the lowest experienced uplink interference, based on eNB measurements of wideband uplink received interference power.

The above rules essentially assume priority of primary over secondary component carriers, as each eNB should always have one PCC with full cell coverage. The inter-eNB path loss measurements are used to ensure that only eNBs with the largest possible path loss separation select the same component carrier for primary.

It is worth mentioning that the proposed method, solely relying on what the eNBs sense, was found to be sensitive to the order in which eNBs are turned on in case of a very limited number of component carriers from which to choose. However, with five component carriers, the sensitivity was rather small.

After the new eNB has selected its PCC, the cell is configured, and it is ready to transmit and carry traffic. In parallel, the eNB shall constantly monitor the quality of the PCC to make sure that it continues to have the desired quality and coverage. If poor quality is detected, recovery actions will be triggered to improve the situation. Such actions can be understood as additional defensive measures, not allowing potentially erroneous SCC allocations to catastrophically interfere with neighboring base stations. Recovery actions are the subject of ongoing investigations and out of the scope of this contribution; nonetheless, they may range from interference reduction requests toward neighboring cells where the same component carrier is used as an SCC, to the selection of a new PCC with better quality.

Secondary Component Carrier Selection

As stated earlier, our scheme imposes certain constraints for selection of SCCs, which basically implies that eNBs have to take the interference created toward other cells into account. The goal is a flexible yet simple and efficient sharing of the spectral resources that will not prevent one cell from using the entire spectrum when this is a sensible choice. Granting eNBs the ability to "learn" what sensible means is the key aspect here.

One of the design targets is to maximize the cell throughput for each eNB, but always ensuring that the experienced signal-to-interference-plus-noise ratio (SINR) on PCC and SCC equals at least $(C/I)_{PCC}$ and $(C/I)_{SCC}$, which represent minimum SINR targets expressed in decibels for the PCC and SCCs, respectively. These are considered as configurable parameters that could come from operations and maintenance (O&M), for example. Without loss of generality, we assume that $(C/I)_{PCC}$ is higher than $(C/I)_{SCC}$ as the PCC is assumed to always have full cell coverage while the SCCs may have reduced coverage (i.e., use less transmit power).

Once it is detected that the capacity offered by the PCC alone is not sufficient to carry the offered traffic, the eNB will use two information sources to autonomously decide whether it can allocate additional SCCs. The first source is the aforementioned RRAT, which provides real-time information on the usage of component carriers by neighboring eNBs. The second piece is the background interference matrix (BIM), which essentially expresses the interference coupling between cells. Now, unlike the selection of the PCC, UE assistance comes into the picture during the creation and maintenance of BIMs.

Each active UE connected to a cell performs downlink measurements of reference signal received power levels which are reported to its serving eNB. These measurements are conducted both towards the serving cell and the surrounding cells (e.g., for handover purposes). Given these UE measurements, the serving eNB calculates a ratio expressed in decibels of own to other cell received signal power. We call it a conditional C/I sample. That essentially allows eNBs to produce an estimate of potential signal quality as perceived by their served UE. Each time a certain (quantized) value is calculated, an occurrence counter is incremented. Eventually, given enough samples, empirical C/I distributions are generated locally by each eNB, one for each detected neighbor. A matrix is then built; we call it the *incoming BIM*.

The C/I value stored in the BIM for each neighboring cell is the value corresponding to a certain outage probability of , say, 95 percent. The C/I value is a measure of mutual interference coupling between a pair of cells. Therefore, each cell maintains local information on all potential *interfering* cells and a corresponding C/I value. In this example only 5 percent of users are likely to experience C/I values in the downlink lower than the value stored in the BIM. Notice that this C/I is *only* realized if the interfered cell and the interfering cell use the same component carrier simultaneously. As component carriers are likely to experience the same path loss conditions, the BIM is component-carrier-independent as it is only based on path loss types of measurement (i.e., it is sufficient for the UE to measure a single component carrier per cell).

Alternatively, in a more dynamic setting the C/I value stored in the BIM for each neighboring cell could correspond to near-real-time conditional C/I values reported by the served UE most severely impacted by that particular neighbor. This approach would better capture the effects of faraway yet strong femtocells that dramatically affect only few UEs (e.g., those near windows in a tall building).

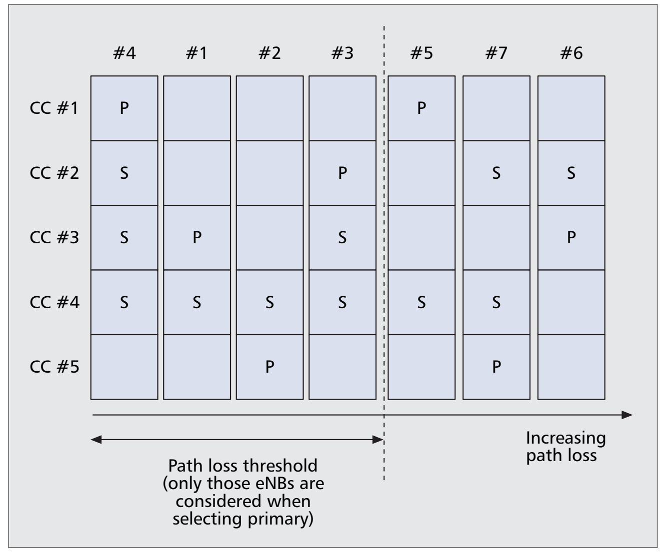

■ **Figure 2.** *Matrix for initial primary component carrier selection.*

In addition to the incoming BIM, eNBs also maintain another BIM table that lists all the potentially *interfered* cells. This BIM is known as the *outgoing BIM*. Basically, it allows a cell to estimate how much interference it generates toward each of its neighbors if it decides to use the same CC the neighboring cell already uses. It is linked to the incoming BIM as follows: At the same time an interfering cell entry (cell 2) is added or modified into the incoming BIM of the interfered cell (cell 1), the corresponding interfered cell (cell 1) is added as an entry into the outgoing BIM of the interfering cell (cell 2). The relation between the incoming and outgoing BIMs is illustrated in Fig. 3.

It is assumed that the reporting of measurements from the UE to the eNBs for the purpose of BIM is fairly slow in order to minimize the control signaling overhead and measurement burden from this. Similarly, the update rate of the local BIM information in each eNB is also anticipated to be rather slow compared to, say, packet scheduling. However, the ideal update rate is the subject of future investigations.

In possession of the information just described, an eNB is now able to decide whether or not the new allocation(s) will jeopardize any existing allocations based on the target SINR values. As explained, we assume a priori knowledge of the minimum SINR targets $(C/I)_{PCC}$ and $(C/I)_{SCC}$ for primary and secondary component carriers, respectively. The process is fairly straightforward, and the interested reader can find a somewhat more formal mathematical description in [11]. In the following we provide a simplified description of the process.

In essence, for each component carrier not yet allocated to the cell, the eNB calculates a set of four differences (in dB). These differences can be understood as neighbor-specific BIM

If for any given neighbor using that particular component carrier either as a PCC or SCC, any of the four margins is found to be negative, that particular component carrier is not taken into use and another component carrier is evaluated.

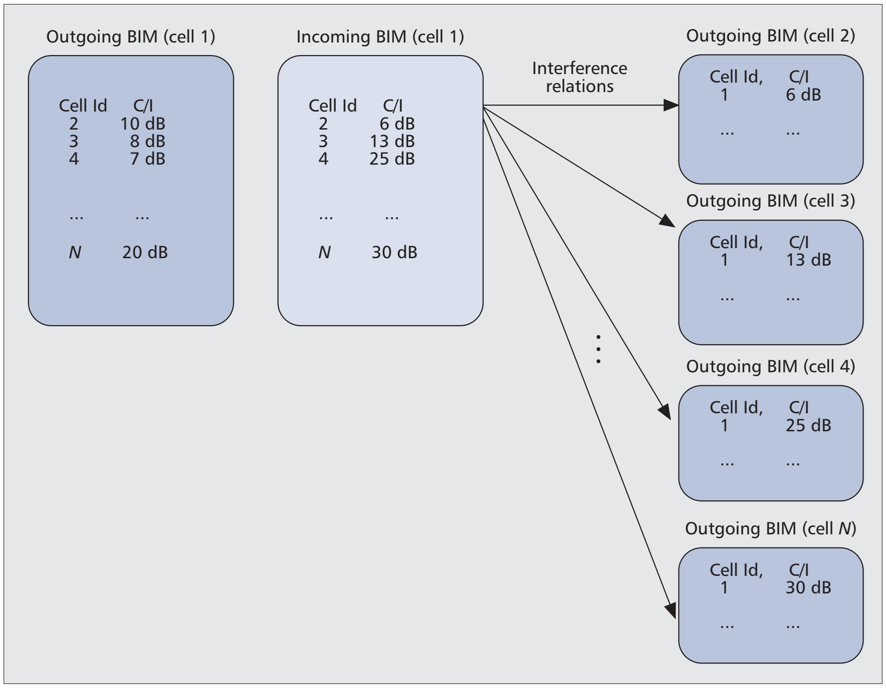

■ **Figure 3.** *Relation between incoming and outgoing BIM entries.*

entry margins with respect to $(C/I)_{SCC}$ in incoming interference evaluations, and with respect to either $(C/I)_{PCC}$ and $(C/I)_{SCC}$ in outgoing interference evaluations, depending on the component carrier usage of the interfered neighbor. If for any given neighbor using that particular component carrier as either a PCC or SCC, any of the four margins is found to be negative, that particular component carrier is not taken into use, and another component carrier is evaluated.

The four differences mentioned earlier correspond in fact to estimated downlink incoming, downlink outgoing, uplink incoming, and uplink outgoing SINR margins. It is important to stress that all uplink estimations are rough approximations of the actual uplink interference situation based on measurements UE has made on the "interfered" side. The rationale behind this is that incoming/outgoing downlink interference propagates through the same path as the outgoing/incoming uplink interference; thus, the downlink C/I estimate contains correlated and useful information. Now, given the hypothetical C/I values in Fig. 3, a simple example illustrates the proposed concept. Let us assume cell 1 is evaluating a component carrier that is currently only in use by cell 3 as its PCC, and $(C/I)_{PCC}$ and $(C/I)_{SCC}$ are set to 10 dB and 8 dB, respectively. Since cell 1 intends to use this component carrier as an SCC, the estimated downlink incoming C/I margin is positive, since 13 dB is above $(C/I)_{SCC}$. However, allocation will be denied because the estimated downlink outgoing C/I margin is negative, for 8 dB is lower than $(C/I)_{PCC}$. Uplink incoming and outgoing SINR margins are calculated similarly.

Performance Results

We study the potential benefits of our proposed autonomous component carrier selection (ACCS) for LTE-Advanced femtocells using system-level simulations. Our system operates at 3.4 GHz carrier frequency with up to 100 MHz bandwidth, the maximum transmission power of eNBs is 200 mW (23 dBm), and 3dBi antenna gain is assumed. Even though our scheme does not preclude other power allocations, for simplicity, there is no downlink power control, and the total transmission power is evenly divided among the component carriers into which the bandwidth is divided; hence, eNBs will only transmit at full power if they employ all component carriers. A simple full-buffer traffic model (i.e., eNBs and UEs always have data to transmit) and a simple round-robin packet scheduler are considered.

Figure 4 depicts the topology of our reference residential scenario. It represents the model for a single indoor floor layout with one eNB (small circle) randomly placed in each 10 m × 10 m four-room residence. The number of uniformly distributed users per residence is fixed to 4. The indoor path loss and slow fading models used are based on A1-type generalized path loss models for the frequency range 2–6 GHz developed in WINNER [12].

The simulation tool relies on series of "snapshots." During each snapshot, path loss, shadowing, and the location of devices remain constant. In practice, various system-level practical aspects such as the effects of achievable bandwidth efficiency, control channel overhead, and receiver algorithms all limit the achievable system-level

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.

WHAT WILL YOU BUILD? | **sales@docketalarm.com** | **1-866-77-FASTCASE**

fastcase®
Smarter legal research.