3GPP TSG-RAN WG2 #58   *R2-071762*
7-11 May 2007
Kobe, Japan

Agenda item: 4.5

Source: QUALCOMM Europe

Title: Scheduling of D-BCH

Document for: Discussion

1.   Introduction

At RAN2#57, it was agreed that the system information is divided into scheduling units, with SU-1 carrying scheduling information for the other SUs. This document attempts to take some steps towards understanding how this agreement can be implemented in practice.

2.   Discussion

2.1. Simple approaches to scheduling

Once the UE has read the primary (and secondary, if applicable) BCH and camped on a cell, it needs to receive the rest of the system information from the D-BCH. It is fairly clear how this happens in general terms: TS 36.300 specifies that SU-1 includes "[s]cheduling information of the other Scheduling Units", so presumably the UE monitors the scheduling channel for an indicator of SU-1, then reads further scheduling information from there.

In the simplest interpretation, this text could mean that SU-1 contains explicit pointers to the resource blocks that contain the other SUs. In this case, a separate L1/2 control channel for the D-BCH might not actually be needed (with the consequence that when *any* SU changed, all affected UEs would need to read the P-BCH and then SU-1 to find the updated version). This situation is shown in Figure 1.


Figure 1: Explicit scheduling in SU-1

However, it should be remembered that some SUs, such as the container for the positioning information, could be delivered over a *very* long period (*e.g.* tens of seconds to minutes), and it is not clear that it would be realistic to expect the scheduler to know this far in advance exactly how resource blocks would be assigned. Therefore it could be impractical to expect SU-1 to carry a complete set of explicit scheduling for all the other SUs.

On the other hand, if such a complete set is not present in SU-1, the question arises of how exactly the UE *does* receive scheduling for such a long-period SU. The simplest approach is to assume that the "scheduling information" in SU-1

just consists of a pointer to a control channel for the D-BCH, and other SUs are scheduled independently using that control channel, as shown in Figure 2. (A "hybrid" approach could also be used, where some SUs were indicated explicitly in SU-1, but longer-period ones whose scheduling could not be known in advance would be scheduled independently using the control channel.)


Figure 2: Waiting for long-period system information

Obviously, this mode of scheduling is not optimal; as the figure indicates, it forces the UE to "busy-wait" on the scheduling channel for a long period for the scheduling of SU-2. It also raises the question of how the UE distinguishes the scheduling for different SUs; does each SU have its own BCCH-RNTI?

Neither of the straightforward approaches considered here is really satisfactory; either the scheduler must have the ability to reserve resource blocks outlandishly far in advance, or the UE has an unreasonable duty cycle while it waits for all SUs to arrive. We attempt to address these deficiencies in the next section.

2.2. Possible solutions

In order to take advantage of the flexibility of the DL-SCH structure, and to avoid the need for rigid advance scheduling, we assume that the D-BCH has an associated control channel, which is used to schedule the SUs. (SU-1 could be an exception, scheduled only on the P-BCH; the figures below make this assumption for simplicity.) The "scheduling information" referred to in TS 36.300 then indicates not the exact location of the SUs themselves, but the point where they are scheduled (Figure 3).


Figure 3: SU-1 indicates scheduling of SU-2

This approach may be acceptable. However, in the case of a very-long-period SU, it still requires advance information about the scheduling. If the SUs can be scheduled according to rigid repetition cycles, this information is easy for the

network to determine, but then the SUs must be scheduled compatibly to prevent collisions in the SDCCH slots, and if the number of SUs is large this could present a combinatorial problem. Therefore it might be desirable to allow a more flexible transmission schedule, *e.g.*, by allowing SU-1 to indicate that a *later* instance of SU-1 will give the scheduling for SU-2 (Figure 4).


Figure 4: SU-1 indicates when scheduling of SU-2 will be available

Here the scheduling of SU-2 is slightly irregular as compared to the SU-1 cycle; the second instance of SU-2 is scheduled one slot "late", and the location of that scheduling information is included in the fourth instance (only) of SU-1. The second and third instances of SU-1 just contain pointers into the future, to the fourth instance where the scheduling information will be available.

The first alternative is conceptually straightforward, but raises some practical complexities. Although the constraint problem is not as bad as with explicit scheduling of SU-2 within SU-1, the scheduler still needs to be able to guarantee in advance that the timing of the scheduling for different SUs does not collide. To some extent this problem could be addressed by having the timing sent as a lower bound ("SU-2 will be scheduled 23873 slots from now, or shortly thereafter") rather than as an exact offset ("SU-2 will be scheduled exactly 23873 slots from now") or a fixed cycle ("SU-2 is scheduled every 24000 slots, starting 23873 slots from now").

The second approach does not have this minor complexity, but it is a rather roundabout approach to the ultimate goal of reading SU-2, and it may be preferable to use a logically simpler solution and find a way to resolve any issues that arise.

In both cases, the control channel needs to have the ability to indicate *which* SU is being scheduled. This could be achieved either with extra information in the scheduling channel itself (space permitting), or by using multiple BCCH-RNTIs.

In practice, the differences between these two approaches to scheduling do not appear to be very great; none of the issues noted are showstoppers, and either approach will work. Indeed, if the long-term constraints on the scheduler discussed in Section 2.1 are acceptable, the problem does not need to be solved, since SU-1 can simply carry an extremely foresighted schedule for SU-2 itself. It would be premature to take a decision at this level of detail, but it may be a good time to consider whether there is a group preference.

3. Conclusion

This document has identified some potential issues with the scheduling of the D-BCH and possible solutions. While a final decision is not needed at this stage, we suggest that RAN2 should begin discussion of these issues.