



# LTE

## The UMTS Long Term Evolution

FROM THEORY TO PRACTICE

Edited by: Stefania Sesia • Issam Toufik • Matthew Baker

SECOND EDITION

Including Release 10 for LTE-Advanced



# LTE – The UMTS Long Term Evolution

# LTE – The UMTS Long Term Evolution

From Theory to Practice

Second Edition

**Stefania Sesia**

*ST-Ericsson, France*

**Issam Toufik**

*ETSI, France*

**Matthew Baker**

*Alcatel-Lucent, UK*



A John Wiley & Sons, Ltd., Publication

This edition first published 2011  
© 2011 John Wiley & Sons Ltd.

*Registered office*

John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex, PO19 8SQ,  
United Kingdom

For details of our global editorial offices, for customer services and for information about how to apply for permission to reuse the copyright material in this book please see our website at [www.wiley.com](http://www.wiley.com).

The rights of the authors to be identified as the authors of this work have been asserted in accordance with the Copyright, Designs and Patents Act 1988.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, except as permitted by the UK Copyright, Designs and Patents Act 1988, without the prior permission of the publisher.

Photograph on cover courtesy of Alcatel-Lucent, from the *ngConnect* LTE-equipped car.  
3GPP website reproduced by permission of © 3GPP™.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

Designations used by companies to distinguish their products are often claimed as trademarks. All brand names and product names used in this book are trade names, service marks, trademarks or registered trademarks of their respective owners. The publisher is not associated with any product or vendor mentioned in this book. This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold on the understanding that the publisher is not engaged in rendering professional services. If professional advice or other expert assistance is required, the services of a competent professional should be sought.

*Library of Congress Cataloging-in-Publication Data*

Sesia, Stefania.

LTE—the UMTS long term evolution : from theory to practice / Stefania Sesia, Issam Toufik, Matthew Baker. — 2nd ed.

p. cm.

Includes bibliographical references and index.

ISBN 978-0-470-66025-6 (hardback)

1. Universal Mobile Telecommunications System. 2. Long-Term Evolution (Telecommunications)  
I. Toufik, Issam. II. Baker, Matthew (Matthew P.J.) III. Title.

TK5103.4883.S47 2011

621.3845'6—dc22

2010039466

A catalogue record for this book is available from the British Library.

Print ISBN: 9780470660256 (H/B)

ePDF ISBN: 9780470978511

oBook ISBN: 9780470978504

epub ISBN: 9780470978641

Printed in Great Britain by CPI Antony Rowe, Chippenham, Wiltshire.

# 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.

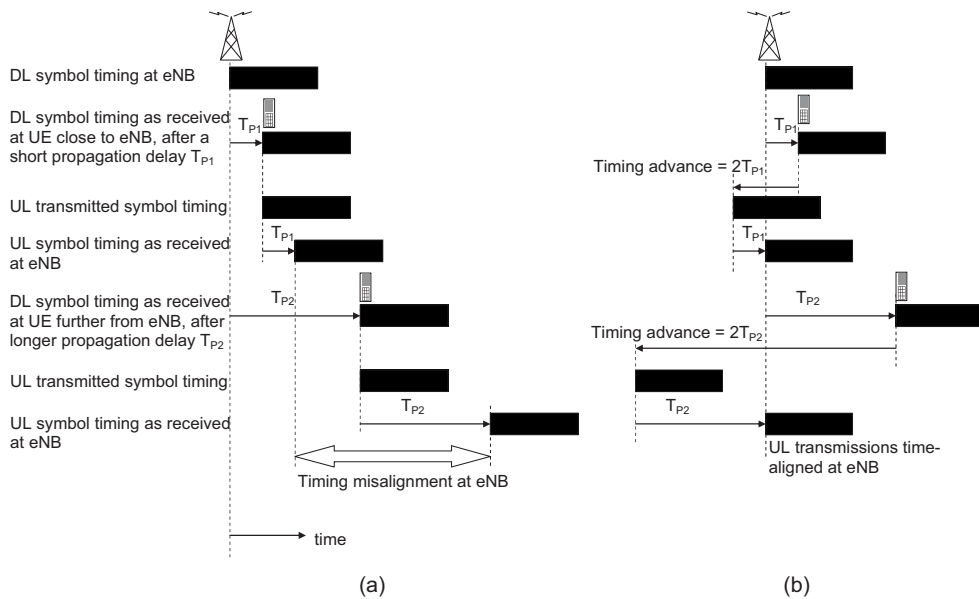


Figure 18.1: Time alignment of uplink transmissions:  
 (a) without timing advance; (b) with timing advance.

## 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  $0.67 \text{ ms}$ ,<sup>1</sup> corresponding to a cell radius

<sup>1</sup>This 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.<sup>2</sup> 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  $\mu\text{s}$  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  $\mu\text{s}$ ). 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  $\mu\text{s}$  does not cause significant degradation in system performance due to increased interference. Thus the granularity of 0.52  $\mu\text{s}$  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  $\mu\text{s}$ <sup>3</sup> (except if the downlink system bandwidth is 1.4 MHz, in which case the UE initial transmission timing accuracy is relaxed to 0.79  $\mu\text{s}$ ) (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.<sup>4</sup>

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

<sup>2</sup>In 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.

<sup>3</sup>This is equal to  $12T_s$ .

<sup>4</sup>In 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.<sup>5</sup>

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.

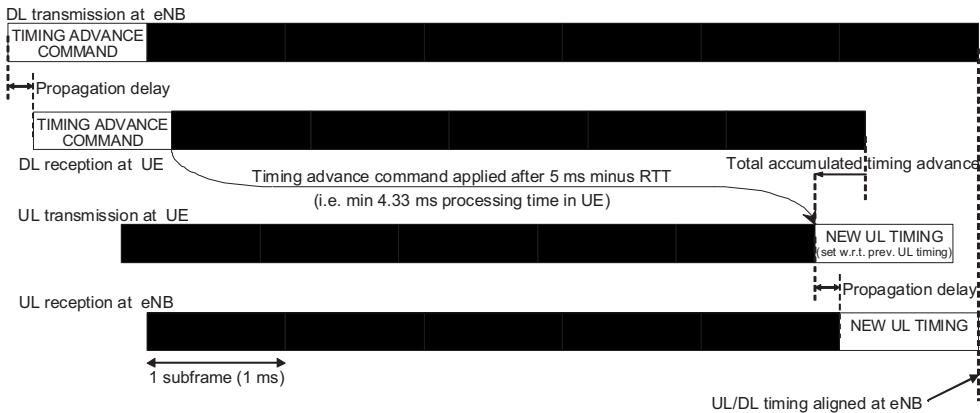


Figure 18.2: Application of timing advance update commands.

<sup>5</sup>Transmission 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  $\mu\text{s}$ ; 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  $\mu\text{s}$ ). 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.<sup>6</sup> 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  $\mu\text{s}$  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,<sup>7</sup> 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,

---

<sup>6</sup>Note that loss of uplink synchronization is possible without leaving RRC\_CONNECTED state (see Chapter 3).

<sup>7</sup>The 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  $\pm 1$  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 PSD<sup>8</sup> 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

<sup>8</sup>In 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 (Block 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  $\alpha$ .

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  $\alpha$  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  $\alpha$  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  $\Delta_{TF}$ , 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.<sup>9</sup>

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 HSDPA<sup>10</sup> 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

<sup>9</sup>In 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).

<sup>10</sup>High 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.<sup>11</sup> 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  $\pm 4$  dB from the semi-static operating point, but a relatively large power step can be triggered by a single command (up to  $\pm 8$  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.<sup>12</sup> Transmission power changes have to be completed within 20  $\mu$ 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  $\Delta_{\text{TPC}}$  denotes a TPC command,  $f(\cdot)$  represents accumulation in the case of accumulative TPC commands, and  $M$  is the number of allocated RBs.

<sup>11</sup>The 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.

<sup>12</sup>For 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].

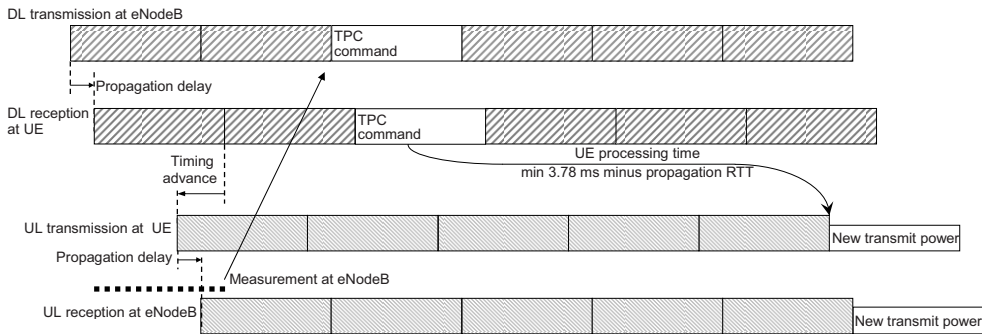


Figure 18.5: Timing of the uplink power control loop.

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  $\alpha$ ) 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.

## References<sup>13</sup>

- [1] Nokia, ‘R1-063377: UL Timing Control Accuracy and Update Rate’, [www.3gpp.org](http://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](http://www.3gpp.org), 3GPP TSG RAN WGI, 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](http://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](http://www.3gpp.org).
- [5] M. P. J. Baker and T. J. Mouldsley, ‘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](http://www.3gpp.org).
- [7] Ericsson, ‘R1-080881: Range and Representation of Delta\_MCS’, [www.3gpp.org](http://www.3gpp.org), 3GPP TSG RAN WG1, meeting 52, Sorrento, Italy, February 2008.

<sup>13</sup>All web sites confirmed 1<sup>st</sup> 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).<sup>1</sup> 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

<sup>1</sup>For TDD deployment, the number of CCs and the bandwidth of each CC in uplink and downlink is expected to be the same.

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, NAS<sup>2</sup> 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-RNTI<sup>3</sup>) 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.<sup>4</sup> 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

<sup>2</sup>Non-Access Stratum.

<sup>3</sup>Cell Radio Network Temporary Identifier.

<sup>4</sup>See 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).

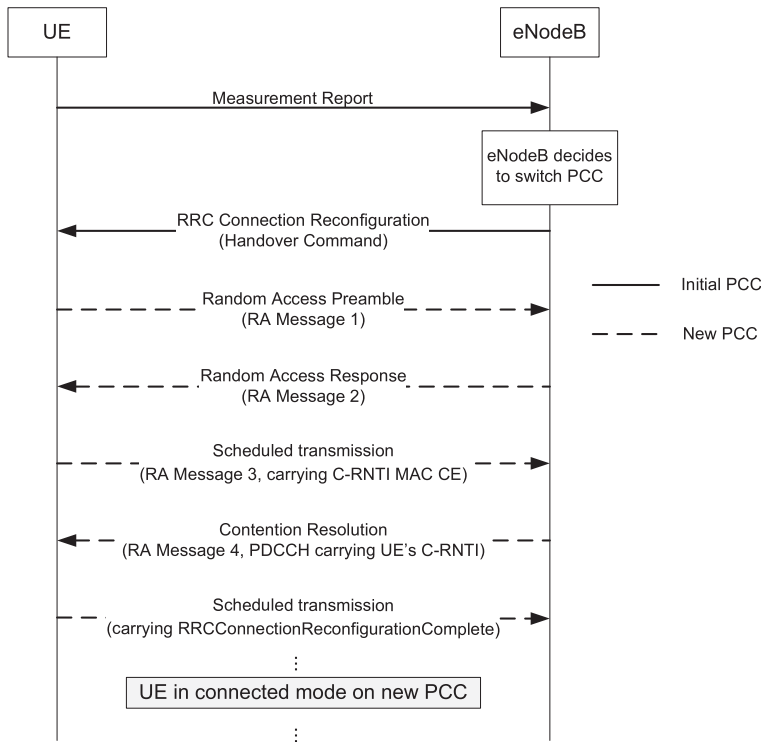


Figure 28.1: Procedure for PCell change with contention-based handover.

- 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

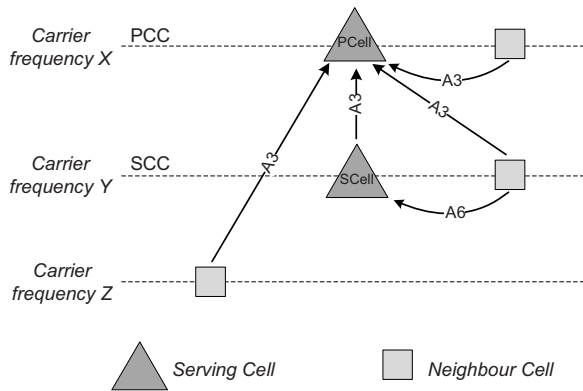


Figure 28.2: Measurement Events A3 and A6.

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.

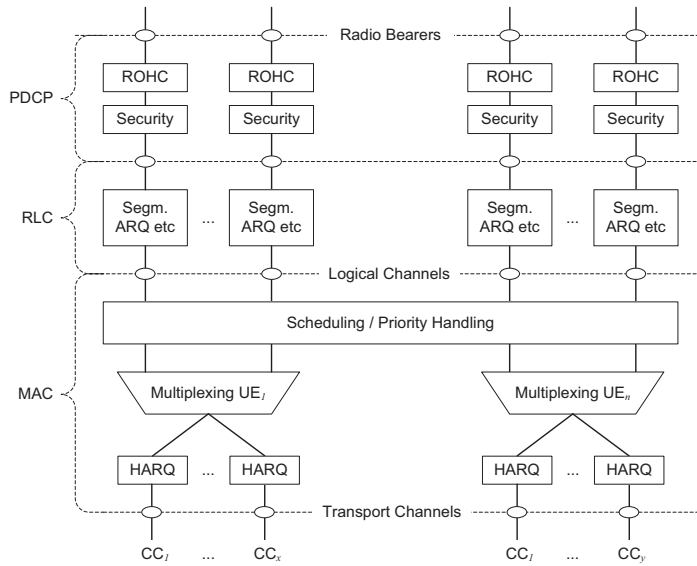


Figure 28.3: Downlink Layer 2 protocol structure for carrier aggregation.  
 Reproduced by permission of © 3GPP.

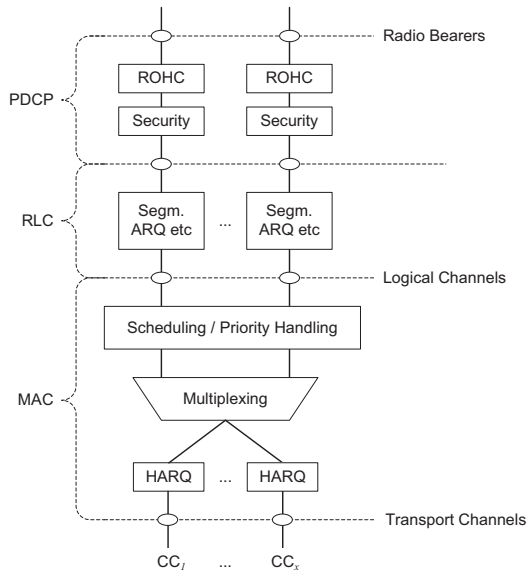


Figure 28.4: Uplink Layer 2 protocol structure for carrier aggregation.  
 Reproduced by permission of © 3GPP.

#### 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.<sup>5</sup> 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  $n + 8$ .<sup>6</sup>

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.

---

<sup>5</sup>Downlink SCell quality is never used to cause a UE to cease transmissions in an SCell.

<sup>6</sup>If 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.<sup>7</sup> 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 PHICH<sup>8</sup> at the start of each subframe, as in Release 8 (see Figure 9.5).

<sup>7</sup>The introduction of uplink MIMO in Release 10 (see Section 29.4) further increases the uplink data rates, with corresponding impact on the BSR.

<sup>8</sup>Physical Control Format Indicator Channel, Physical Downlink Control Channel and Physical Hybrid ARQ Indicator Channel.

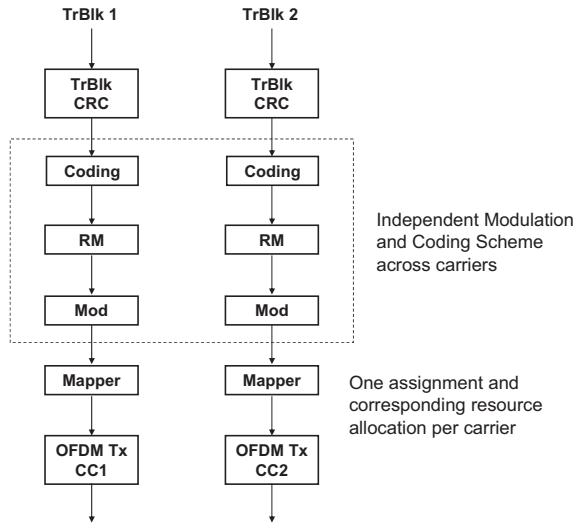


Figure 28.5: MAC to physical layer mapping for carrier aggregation.

### 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.<sup>9</sup> 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).

<sup>9</sup>Note that PDSCH transmissions on the PCell cannot use cross-carrier scheduling – their corresponding PDCCH messages must also be transmitted on the PCell.

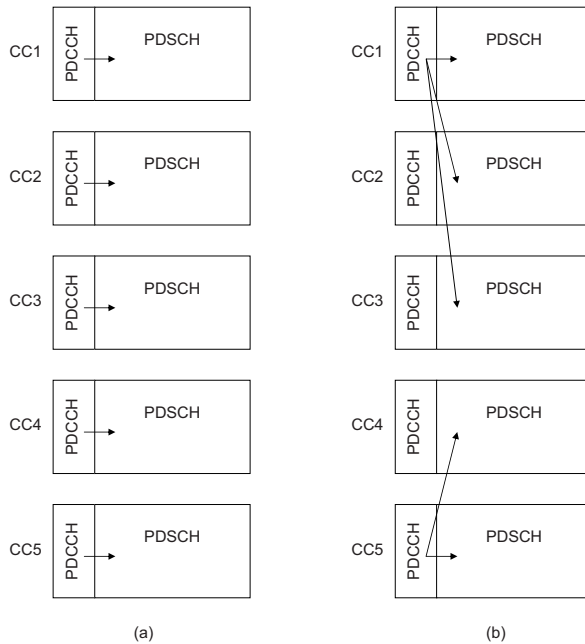


Figure 28.6: Examples of semi-statically configured mappings from PDCCH scheduling messages to CCs for data transmission: (a) without cross-carrier scheduling; (b) with cross-carrier scheduling.

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).<sup>10</sup>

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.<sup>11</sup>

For any downlink CC where the UE monitors PDCCH without the CIF being configured, the search space is the same as in Release 8.

<sup>10</sup>In 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.

<sup>11</sup>A 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.



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.<sup>12</sup> This effectively provides ICIC for the PDCCH, while the Release 8 ICIC mechanisms may be utilized for PDSCH data.<sup>13</sup>




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 OFDM<sup>14</sup> 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

<sup>12</sup>These 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.

<sup>13</sup>For 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.

<sup>14</sup>Orthogonal 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-OFDM<sup>15</sup> 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

---

<sup>15</sup>Discrete Fourier Transform Spread Orthogonal Frequency Division Multiplexing.

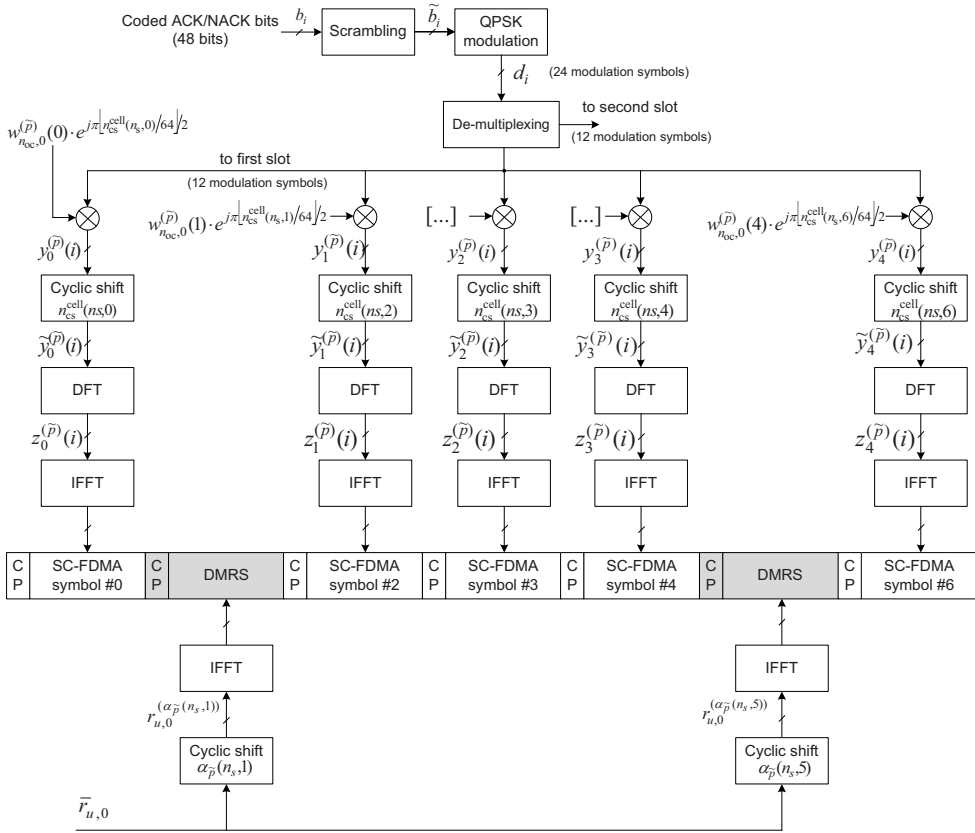


Figure 28.9: Structure of PUCCH format 3.

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)<sup>16</sup> 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)

<sup>16</sup>See 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.

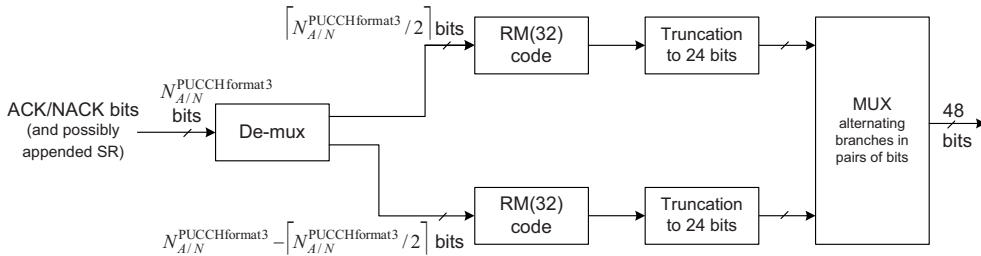


Figure 28.10: Coding and scrambling for PUCCH format 3.

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.<sup>17</sup>

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’<sup>18</sup> 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.

<sup>17</sup>This 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).

<sup>18</sup>‘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 normally<sup>19</sup> 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

<sup>19</sup>The 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  $\alpha$ , 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.<sup>20</sup>

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

---

<sup>20</sup>Note 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 IFFT<sup>21</sup> 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  $x^{\text{th}}$  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.

<sup>21</sup>Inverse 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 1x2 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)<sup>22</sup> 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  $\geq 20$  MHz;
- **Option B:** Multiple RF, and baseband processing with bandwidth  $\leq 20$  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.

---

<sup>22</sup>See Chapter 21.

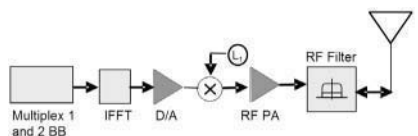
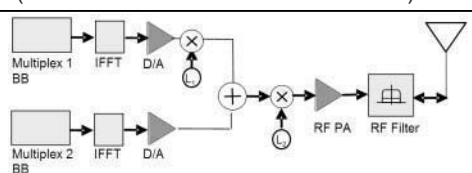
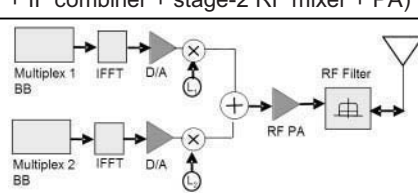
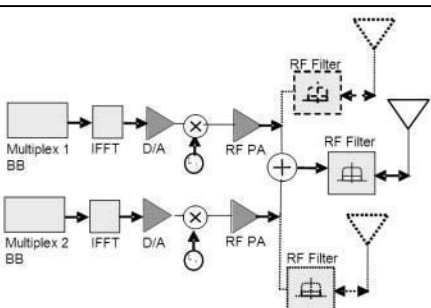
	Transmitter architecture	Aggregation Scenarios
A	 <p>Single (baseband + IFFT + DAC + mixer + PA)</p>	Intra-band contiguous CCs
B	 <p>Multiple (baseband + IFFT + DAC), single (stage-1 IF mixer + IF combiner + stage-2 RF mixer + PA)</p>	Intra-band contiguous CCs Intra-band non-contiguous CCs
C	 <p>Multiple (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.<sup>23</sup> 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.

## References<sup>24</sup>

- [1] 3GPP Technical Report 36.913, 'Requirements for further advancements for Evolved Universal Terrestrial Radio Access (E-UTRA) (LTE-Advanced)', [www.3gpp.org](http://www.3gpp.org).
- [2] 3GPP Technical Report 36.815, 'LTE-Advanced Feasibility Studies in RAN WG4 (Release 9)', [www.3gpp.org](http://www.3gpp.org).
- [3] 3GPP Technical Specification 36.213, 'Physical Layer Procedures (Release 10)', [www.3gpp.org](http://www.3gpp.org).
- [4] 3GPP Technical Report 25.814, 'Physical Layer Aspects for Evolved Universal Terrestrial Radio Access (UTRA) (Release 7)', [www.3gpp.org](http://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](http://www.3gpp.org).

---

<sup>23</sup>See Section 21.2 for the E-UTRA band definitions.

<sup>24</sup>All web sites confirmed 1<sup>st</sup> March 2011.