

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**

WILEY



Samsung Electronics Co., Ltd. v. Asus Tech Licensing
EX-1005, 1 of 15

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

WILEY
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.

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.

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

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

18.3 Power Control

18.3.1 Overview

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

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

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

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

[7] 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 ±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[8] 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

[8] 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:

Power per resource block = basic open-loop operating point + 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 α.

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

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

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

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

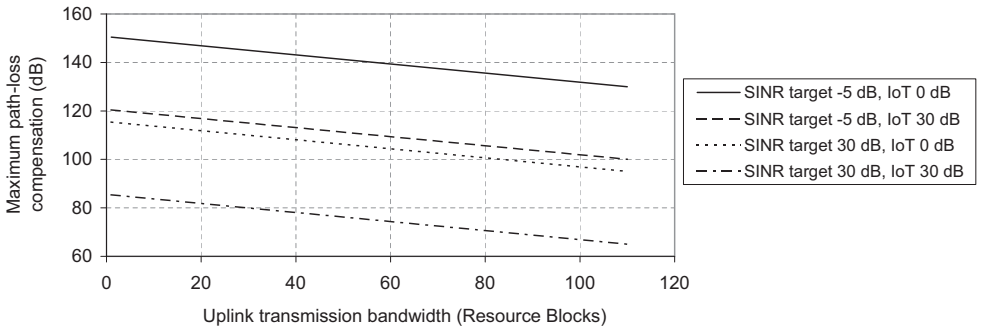


Figure 18.3: Maximum path-loss compensation in typical scenarios for a 23 dBm UE.

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

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

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

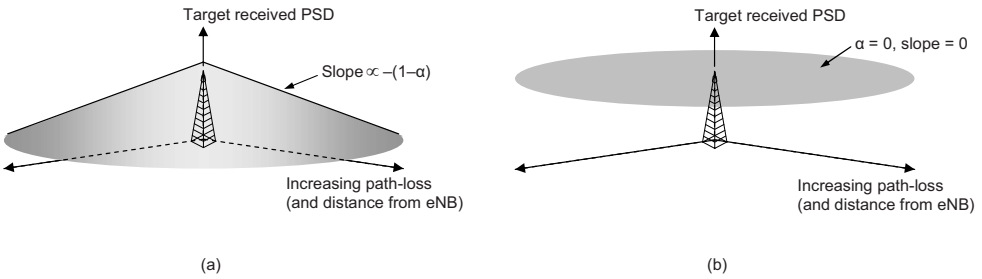


Figure 18.4: (a) Fractional and (b) non-fractional power control.

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

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

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

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

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

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

18.3.2.2 Dynamic Offset

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

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

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

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

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

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

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

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

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

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

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

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

Transport-format-dependent power control is also particularly relevant for the PUCCH, as the PUCCH bandwidth for a UE does not vary depending on the amount of information to be transmitted in a given subframe (ranging from a single bit for a scheduling request or ACK/NACK, to 22 bits for combined dual-codeword ACK/NACK and CQI together – see Section 16.3). For the PUCCH, the magnitude of the power offset for each combination of control information can be adjusted semi-statically by the eNodeB, in order to set a suitable error-rate operating point for each PUCCH format (see Table 16.2). This is analogous to the different power offsets which may be set in HSDPA[10] 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

[9] 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).

[10] 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.[11] The set of offsets which can be signalled by absolute TPC commands is {−4, −1, +1, +4} dB. Thus the absolute power control mode can only control the power within a range of ±4 dB from the semi-static operating point, but a relatively large power step can be triggered by a single command (up to ±8 dB). This mode is therefore suited to scenarios where the scheduling of the UE's uplink transmissions may be intermittent; an absolute TPC command enables the UE's transmission power to be adjusted to a suitable level in a single step after each transmission gap. Absolute TPC commands can also be useful for dynamic frequency-domain inter-cell interference coordination.

The timing of the closed-loop TPC commands is illustrated in Figure 18.5. The transmission power change resulting from a TPC command usually takes effect at the fourth uplink subframe after the TPC command is received.[12] Transmission power changes have to be completed within 20 μs 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:

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

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

[11] 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.

[12] 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].

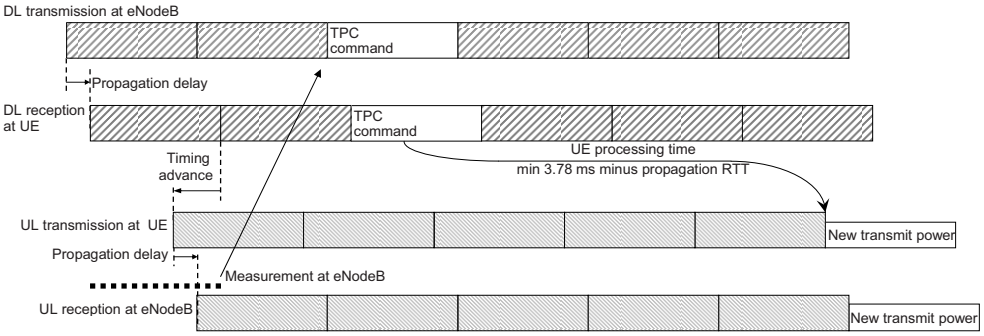


Figure 18.5: Timing of the uplink power control loop.

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

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

References[13]

[1] Nokia, 'R1-063377: UL Timing Control Accuracy and Update Rate', www.3gpp.org, 3GPP TSG RAN WG1, meeting 47, Riga, Latvia, November 2006.

[2] Texas Instruments, 'R1-072841: Simulation of Uplink Timing Error Impact on PUSCH', www.3gpp.org, 3GPP TSG RAN 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.

[4] 3GPP Technical Specification 36.101, 'Evolved Universal Terrestrial Radio Access (E-UTRA); User Equipment (UE) radio transmission and reception', www.3gpp.org.

[5] M. P. J. Baker and T. J. Moulsley, 'Power Control in UMTS Release 99' in *Proc. First IEE Int. Conf. on 3G Communications*, March 2000.

[6] 3GPP Technical Specification 36.213, 'Evolved Universal Terrestrial Radio Access (E-UTRA); Physical Layer Procedures (Release 8)', www.3gpp.org.

[7] Ericsson, 'R1-080881: Range and Representation of Delta_MCS', www.3gpp.org, 3GPP TSG RAN WG1, meeting 52, Sorrento, Italy, February 2008.

[13] All web sites confirmed 1st March 2011.

Samsung Electronics Co., Ltd. v. Asus Tech Licensing EX-1005, 14 of 15

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, $P_{\mathrm{CMAX,c}}$, after taking into account the maximum power reductions explained in Section 21.3.1.2. A single MAC control element contains both $P_{\mathrm{CMAX,c}}$ 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. $P_{\mathrm{CMAX,c}}$ 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.