CHAPTER

7 UPLINK PHYSICAL-LAYER PROCESSING

This chapter provides a description of the basic physical-layer functionality related to the LTE uplink. It essentially follows the same outline as the corresponding downlink description provided in the previous chapter, with detailed descriptions regarding transport-channel processing (Section 7.1), the reference-signal structure (Section 7.2), multi-antenna transmission (Section 7.3), and the uplink L1/L2 control-channel structure (Section 7.4). The chapter ends with an overview of uplink power control and the uplink timing-alignment procedure in Sections 7.5 and 7.6, respectively. Physical aspects related to some specific uplink functions and procedures such as random access are provided in later chapters.

7.1 TRANSPORT-CHANNEL PROCESSING

This section describes the physical-layer processing applied to the uplink shared channel (UL-SCH) and the subsequent mapping to the uplink physical resource in the form of the basic OFDM time–frequency grid. As mentioned before, the UL-SCH is the only uplink transport-channel type in LTE[1] and is used for transmission of all uplink higher-layer information—that is, for both user data and higher-layer control signaling.

7.1.1 PROCESSING STEPS

Figure 7.1 outlines the different steps of the UL-SCH physical-layer processing in case of transmission on a single carrier. Similar to the downlink, in the case of uplink carrier aggregation the different component carriers correspond to separate transport channels with separate physical-layer processing.

The different steps of the uplink transport-channel processing are summarized below. Most of these steps are very similar to the corresponding steps for DL-SCH processing outlined in Section 6.1. For a more detailed overview of the different steps, the reader is referred to the corresponding downlink description.

[1]Strictly speaking, the LTE Random-Access Channel is also defined as a transport-channel type, see Chapter 4. However, RACH only includes a layer-1 preamble and carries no data in form of transport blocks.

4G, LTE-Advanced Pro and The Road to 5G. http://dx.doi.org/10.1016/B978-0-12-804575-6.00007-8

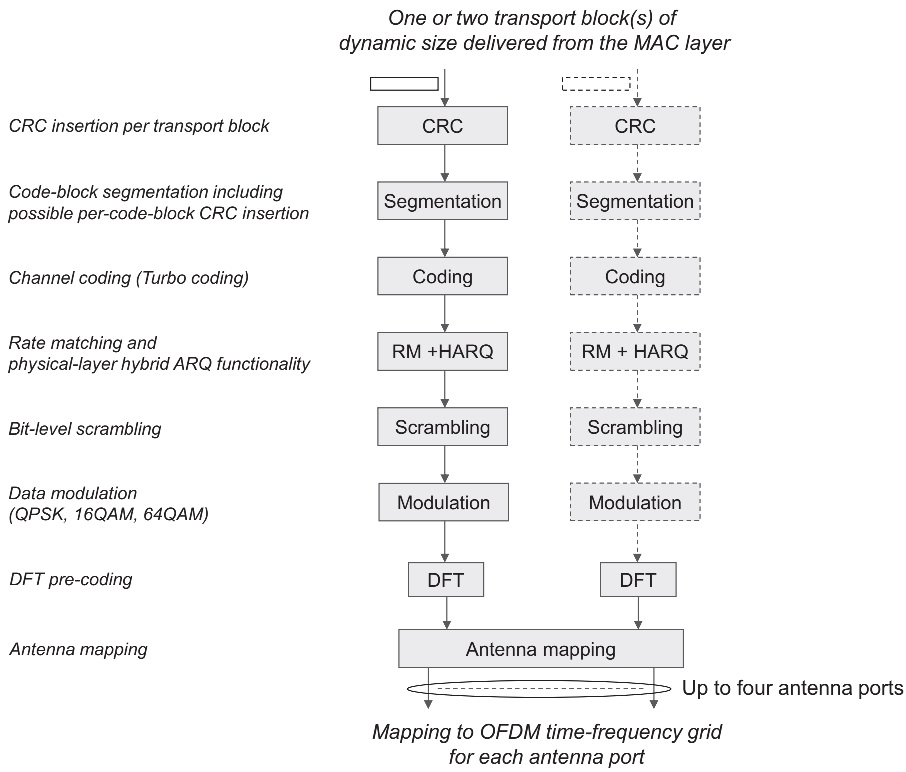

FIGURE 7.1

Physical-layer processing for the uplink shared channel (UL-SCH).

- *CRC insertion per transport block.* A 24-bit CRC is calculated for and appended to each transport block.
- *Code-block segmentation and per-code-block CRC insertion.* In the same way as for the downlink, code-block segmentation, including per-code-block CRC insertion, is applied for transport blocks larger than 6144 bits.
- *Channel coding*. Rate-1/3 Turbo coding with QPP-based inner interleaving is also used for the uplink shared channel.
- *Rate matching and physical-layer hybrid-ARQ functionality.* The physical-layer part of the uplink rate-matching and hybrid-ARQ functionality is essentially the same as the corresponding downlink functionality, with subblock interleaving and insertion into a circular buffer followed by bit selection with four redundancy versions. There are some important differences between the downlink and uplink hybrid-ARQ protocols, such as asynchronous versus synchronous operation as described in Chapter 8. However, these differences are not visible in terms of the physical-layer processing.

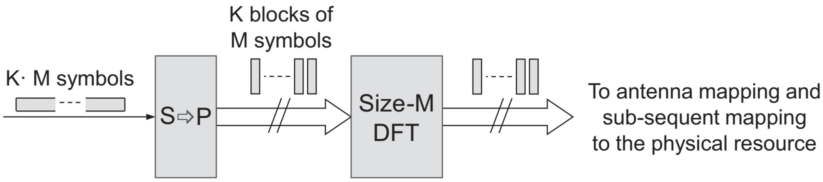

FIGURE 7.2

DFT precoding of K blocks, each consisting of M modulation symbols.

- *Bit-level scrambling*. The aim of uplink scrambling is the same as for the downlink—that is, to randomize, in this case, the *uplink* interference to ensure that the processing gain provided by the channel code can be fully utilized.
- *Data modulation*. Similar to the downlink, QPSK, 16QAM, and 64QAM modulation can also be used for uplink shared-channel transmission, while 256QAM is not supported.
- *DFT precoding*. As illustrated in Figure 7.2, the modulation symbols, in blocks of M symbols, are fed through a size-M DFT, where M corresponds to the number of subcarriers assigned for the transmission. The reason for the DFT precoding is to reduce the cubic metric [10] for the transmitted signal, thereby enabling higher power-amplifier efficiency. From an implementation complexity point of view the DFT size should preferably be constrained to a power of 2. However, such a constraint would limit the scheduler flexibility in terms of the amount of resources that can be assigned for an uplink transmission. Rather, from a flexibility point of view all possible DFT sizes should preferably be allowed. For LTE, a middle-way has been adopted where the DFT size, and thus also the size of the resource allocation, is limited to products of the integers 2, 3, and 5. Thus, for example, DFT sizes of 60, 72, and 96 are allowed but a DFT size of 84 is not allowed.[2] In this way, the DFT can be implemented as a combination of relatively low-complex radix-2, radix-3, and radix-5 FFT processing.
- *Antenna mapping*. The antenna mapping maps the output of the DFT precoder to one or several uplink antenna ports for subsequent mapping to the physical resource (the OFDM time–frequency grid). In the first releases of LTE (releases 8 and 9), only single-antenna transmission was used for the uplink.[3] However, as part of LTE release 10, support for uplink multi-antenna transmission by means of antenna precoding with up to four antennas ports was introduced. More details about LTE uplink multi-antenna transmission are provided in Section 7.3.

[2]As uplink resource assignments are always done in terms of resource blocks of size 12 subcarriers, the DFT size is always a multiple of 12.

[3]Uplink multi-antenna transmission in form of *antenna selection* has been part of the LTE specification since release 8. However, it is an optional device feature with limited implementation in commercially available devices.

7.1.2 MAPPING TO THE PHYSICAL RESOURCE

The scheduler assigns a set of resource-block pairs to be used for the uplink transmission, more specifically for transmission of the physical uplink shared channel (PUSCH) that carries the UL-SCH transport channel. Each such resource-block pair spans 14 OFDM symbols in time (one subframe).[4] However, as will be described in Section 7.2.1, two of these symbols are used for uplink *demodulation reference signals* (DM-RS) and are thus not available for PUSCH transmission. Furthermore, one additional symbol may be reserved for the transmission of *sounding reference signals* (SRS), see Section 7.2.2. Thus, 11 or 12 OFDM symbols are available for PUSCH transmission within each subframe.

Figure 7.3 illustrates how $K \cdot M$ DFT-precoded symbols at the output of the antenna mapping are mapped to the basic OFDM time–frequency grid, where K is the number of available OFDM symbols within a subframe (11 or 12 according to the text in the preceding paragraphs) and M is the assigned bandwidth in number of subcarriers. As there are 12 subcarriers within a resource-block pair, $M = N \cdot 12$ where N is the number of assigned resource-block pairs.

Figure 7.3 assumes that the set of assigned resource-block pairs are contiguous in the frequency domain. This is the typical assumption for DFTS-OFDM and was strictly the case for LTE releases 8 and 9. Mapping of the DFT-precoded signal to frequency-contiguous resource elements is preferred in order to retain good cubic-metric properties of the uplink transmission. At the same time, such a restriction implies an additional constraint for the uplink scheduler, something which may not always be desirable. Therefore, LTE release 10 introduced the possibility to assign partly frequency-separated resources for PUSCH transmission. More specifically, in release 10 the assigned uplink resource may consist of a maximum of two frequency-separated *clusters* as illustrated in Figure 7.4, where each cluster consists of a number of resource-block pairs (N_1 and N_2 resource-block pairs, respectively). In the case of such *multi-cluster transmission*, a single DFT

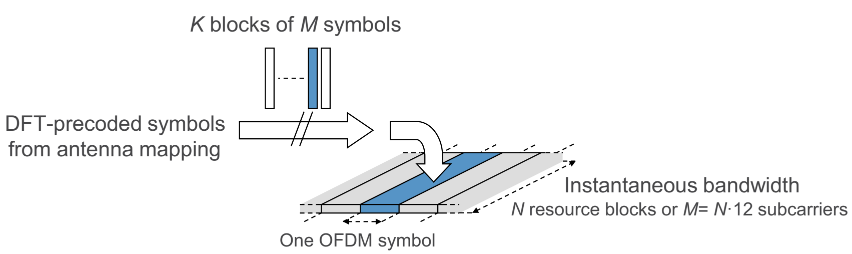

FIGURE 7.3

Mapping to the uplink physical resource.

[4]Twelve symbols in the case of extended cyclic prefix.

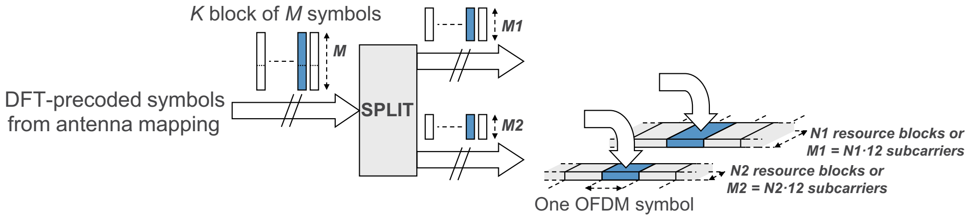

FIGURE 7.4

Uplink multi-cluster transmission.

precoding spans the overall assigned resource in the frequency domain—that is, both clusters. This means that the total assigned bandwidth in number of subcarriers ($M = M1 + M2$) should be aligned with the restrictions on available DFT sizes described in the preceding paragraphs.

7.1.3 PUSCH FREQUENCY HOPPING

In Chapter 6 it was described how the notion of *virtual resource blocks* (VRBs) in combination with the mapping from VRBs to *physical resource blocks* (PRBs) allowed for downlink distributed transmission—that is, the spreading of a downlink transmission in the frequency domain. As described, downlink distributed transmission consists of two separate steps: (1) a mapping from VRB pairs to PRB pairs such that frequency-consecutive VRB pairs are not mapped to frequency-consecutive PRB pairs and (2) a split of each resource-block pair such that the two resource blocks of the resource-block pair are transmitted with a certain frequency gap in between. This second step can be seen as frequency hopping on a slot basis.

The notion of VRBs can also be used for the LTE uplink, allowing for frequency-domain-distributed transmission for the uplink. However, in the uplink, where transmission from a device should always be over a set of consecutive subcarriers in the absence of multi-cluster transmission, distributing resource-block pairs in the frequency domain, as in the first step of downlink distributed transmission, is not possible. Rather, uplink distributed transmission is similar to the second step of downlink distributed transmission—that is, a frequency separation of the transmissions in the first and second slots of a subframe. Uplink distributed transmission for PUSCH can thus more directly be referred to as *uplink frequency hopping*.

There are two types of uplink frequency hopping defined for PUSCH:

- subband-based hopping according to cell-specific hopping/mirroring patterns;
- hopping based on explicit hopping information in the scheduling grant.

Uplink frequency hopping is not supported for multi-cluster transmission as, in that case, sufficient diversity is assumed to be obtained by proper location of the two clusters.

7.1.3.1 Hopping Based on Cell-Specific Hopping/Mirroring Patterns

To support subband-based hopping according to cell-specific hopping/mirroring patterns, a set of consecutive subbands of a certain size is defined from the overall uplink frequency band, as illustrated in Figure 7.5. It should be noted that the subbands do not cover the total uplink frequency band, mainly due to the fact that a number of resource blocks at the edges of the uplink frequency band are used for transmission of L1/L2 control signaling on the physical uplink control channel (PUCCH). For example, in Figure 7.5, the overall uplink bandwidth corresponds to 50 resource blocks and there are a total of four subbands, each consisting of 11 resource blocks. Six resource blocks are not included in the hopping bandwidth and could, for example, be used for PUCCH transmission.

In the case of subband-based hopping, the set of VRBs provided in the scheduling grant are mapped to a corresponding set of PRBs according to a cell-specific hopping pattern. The resource to use for transmission, the *PRBs*, is obtained by shifting the VRBs provided in the scheduling grant by a number of subbands according to the hopping pattern, where the hopping pattern can provide different shifts for each slot. As illustrated in Figure 7.6, a device is assigned VRBs 27, 28, and 29. In the first slot, the predefined hopping pattern takes the

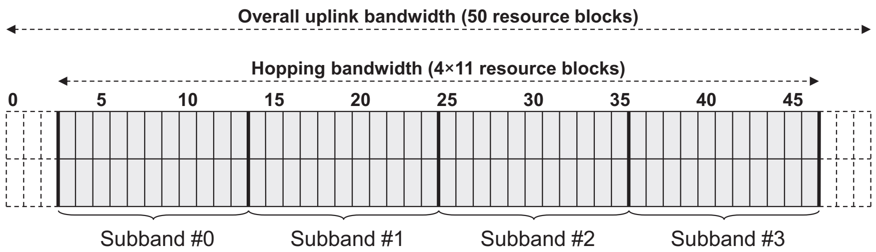

FIGURE 7.5

Definition of subbands for PUSCH hopping. The figure assumes a total of four subbands, each consisting of eleven resource blocks.

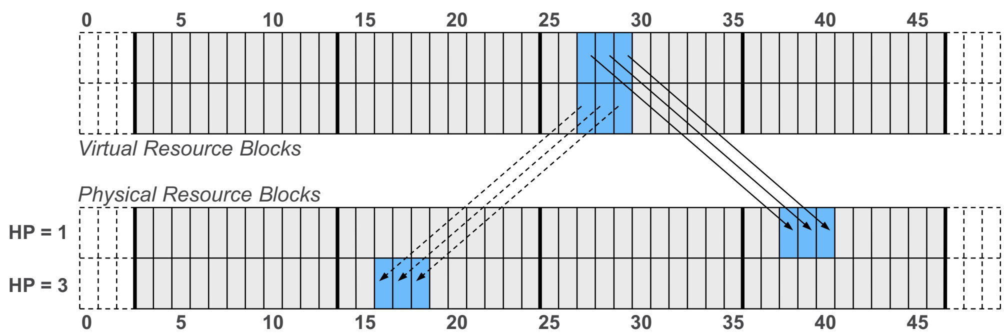

FIGURE 7.6

Hopping according to predefined hopping pattern.

value 1, implying transmission using PRBs one subband to the right—that is, PRBs 38, 39, and 40. In the second slot, the predefined hopping pattern takes the value 3, implying a shift of three subbands to the right in the figure and, consequently, transmission using resource blocks 16, 17, and 18. Note that the shifting "wraps-around"—that is, in the case of four subbands, a shift of three subbands is the same as a negative shift of one subband. As the hopping pattern is cell specific—that is, the same for all devices within a cell—different devices will transmit on nonoverlapping physical resources as long as they are assigned nonoverlapping virtual resources.

In addition to the hopping pattern, there is also a cell-specific *mirroring pattern* defined in a cell. The mirroring pattern controls, on a slot basis, whether or not mirroring within each subband should be applied to the assigned resource. In essence, mirroring implies that the resource blocks within each subband are numbered right to left instead of left to right. Figure 7.7 illustrates mirroring in combination with hopping. Here, the mirroring pattern is such that mirroring is not applied to the first slot while mirroring is applied to the second slot.

Both the hopping pattern and the mirroring pattern depend on the physical-layer cell identity and are thus typically different in neighboring cells. Furthermore, the period of the hopping/mirroring patterns corresponds to one frame.

7.1.3.2 Hopping Based on Explicit Hopping Information

As an alternative to hopping/mirroring according to cell-specific hopping/mirroring patterns as described in the preceding section, uplink slot-based frequency hopping for PUSCH can also be controlled by *explicit hopping information* provided in the scheduling grant. In such a case the scheduling grant includes

- information about the resource to use for uplink transmission in the first slot, exactly as in the nonhopping case;
- additional information about the offset of the resource to use for uplink transmission in the second slot, relative to the resource of the first slot.

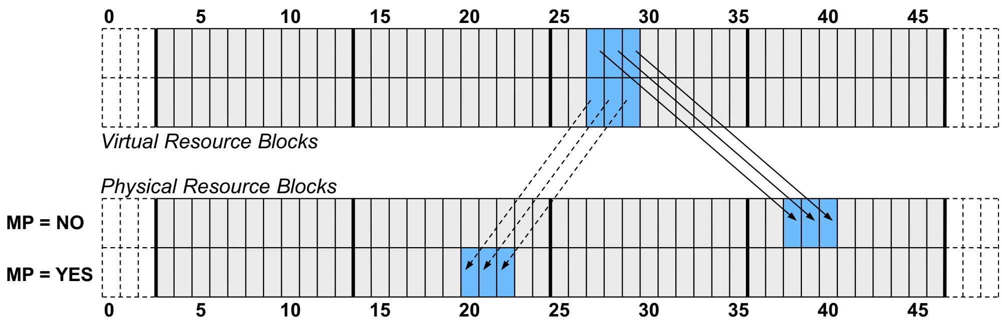

FIGURE 7.7

Hopping/mirroring according to predefined hopping/mirroring patterns. Same hopping pattern as in Figure 7.6.

Selection between hopping according to cell-specific hopping/mirroring patterns as discussed or hopping according to explicit information in the scheduling grant can be done dynamically. More specifically, for cell bandwidths less than 50 resource blocks, there is a single bit in the scheduling grant indicating if hopping should be according to the cell-specific hopping/mirroring patterns or should be according to information in the scheduling grant. In the latter case, the hop is always half of the hopping bandwidth. In the case of larger bandwidths (50 resource blocks and beyond), there are two bits in the scheduling grant. One of the combinations indicates that hopping should be according to the cell-specific hopping/mirroring patterns while the three remaining alternatives indicate hopping of 1/2, +1/4, and −1/4 of the hopping bandwidth. Hopping according to information in the scheduling grant for the case of a cell bandwidth corresponding to 50 resource blocks is illustrated in Figure 7.8. In the first subframe, the scheduling grant indicates a hop of one-half the hopping bandwidth. In the second subframe, the grant indicates a hop of one-fourth the hopping bandwidth (equivalent to a negative hop of three-fourths of the hopping bandwidth). Finally, in the third subframe, the grant indicates a negative hop of one-fourth the hopping bandwidth.

7.2 UPLINK REFERENCE SIGNALS

Similar to the downlink, reference signals are also transmitted on the LTE uplink. There are two types of reference signals defined for the LTE uplink:

- Uplink DM-RS are intended to be used by the base station for channel estimation for coherent demodulation of the uplink physical channels (PUSCH and PUCCH). A DM-RS is thus only transmitted together with PUSCH or PUCCH and is then spanning the same frequency range as the corresponding physical channel.

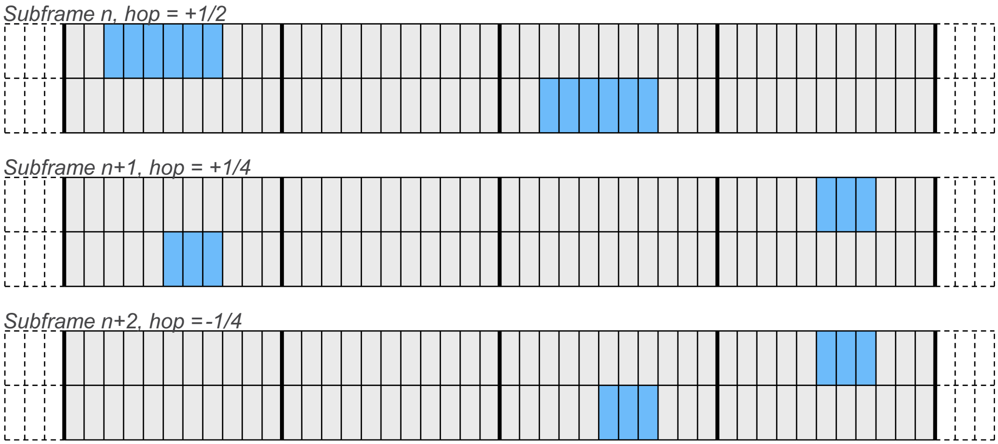

FIGURE 7.8

Frequency hopping according to explicit hopping information.

- Uplink SRS are intended to be used by the base station for channel-state estimation to support uplink channel-dependent scheduling and link adaptation. The SRS can also be used in other cases when uplink transmission is needed although there is no data to transmit. One example is when uplink transmission is needed for the network to be able to estimate the uplink receive timing as part of the *uplink-timing-alignment procedure*, see Section 7.6.

7.2.1 DEMODULATION REFERENCE SIGNALS

Uplink DM-RS are intended to be used for channel estimation for coherent demodulation of the PUSCH to which the UL-SCH transport channel is mapped, as well as for the PUCCH which carries different types of uplink L1/L2 control signaling. The basic principles for uplink DM-RS are the same for PUSCH and PUCCH transmission although there are some differences—for example, in terms of the exact set of OFDM symbols in which the reference signals are transmitted. The discussion in the following primarily focuses on PUSCH DM-RM. Some additional details on the PUCCH DM-RS structure is provided in Section 7.4 as part of the more general description of uplink L1/L2 control signaling.

7.2.1.1 Time–Frequency Structure

Due to the importance of low cubic metric and corresponding high power-amplifier efficiency for uplink transmissions, the principles for uplink reference-signal transmission are different compared to the downlink. In essence, transmitting reference signals frequency multiplexed with other uplink transmissions from the same device is not suitable for the uplink as that would negatively impact the device power-amplifier efficiency due to increased cubic metric. Instead, certain OFDM symbols within a subframe are used exclusively for DM-RS transmission—that is, the reference signals are *time multiplexed* with other uplink transmissions (PUSCH and PUCCH) from the same device. The structure of the reference signal itself then ensures a low cubic metric within these symbols as described in the following.

More specifically, in case of PUSCH transmission DM-RS is transmitted within the fourth symbol of each uplink slot[5] (Figure 7.9). Within each subframe, there are thus two reference-signal transmissions, one in each slot.

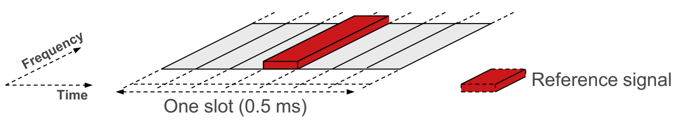

FIGURE 7.9

Transmission of uplink DM-RS within a slot in case of PUSCH transmission.

[5]The third symbol in the case of extended cyclic prefix.

In case of PUCCH transmission, the number of OFDM symbols used for DM-RS transmission in a slot, as well as the exact position of these symbols, differs between different *PUCCH formats* as further described in Section 7.4.

In general, there is no reason to estimate the channel outside the frequency band of the corresponding PUSCH/PUCCH transmission that is to be coherently demodulated. The frequency range spanned by an uplink DM-RS is therefore equal to the instantaneous frequency range spanned by the corresponding PUSCH/PUCCH transmission. This means that, for PUSCH transmission, it should be possible to generate reference signals of different bandwidths, corresponding to the possible bandwidths of a PUSCH transmission. More specifically, it should be possible to generate reference signals of a bandwidth corresponding to $12 \cdot N$ sub carriers, where N corresponds to the bandwidth of the PUSCH transmission measured in number of resource blocks.[6]

Regardless of the kind of uplink transmission (PUSCH or PUCCH), the basic structure of each reference-signal transmission is the same. As illustrated in Figure 7.10, an uplink DM-RS can be defined as a *frequency-domain reference-signal sequence* applied to consecutive inputs of an OFDM modulator—that is, to consecutive subcarriers. Referring to the preceding discussion, in case of PUSCH transmission, the frequency-domain reference-signal sequence should have a length $M = 12 \cdot N$ where N corresponds to the PUSCH bandwidth measured in number of resource blocks. In case of PUCCH transmission, the length of the reference-signal sequence should always be equal to 12.

7.2.1.2 Base Sequences

Uplink reference signals should preferably have the following properties:

- Small power variations in the frequency domain to allow for similar channel-estimation quality for all frequencies spanned by the reference signal. Note that this is equivalent to a well-focused time-domain auto-correlation of the transmitted reference signal.

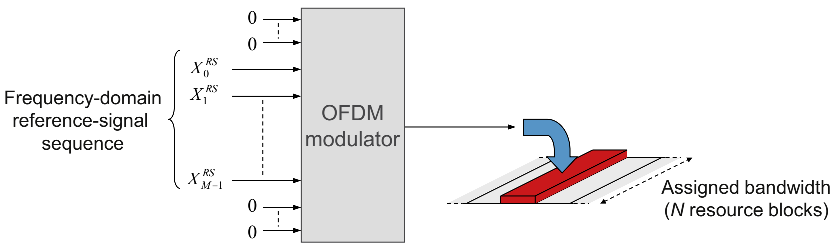

FIGURE 7.10

Generation of uplink DM-RS from a frequency-domain reference-signal sequence.

[6]Due to the imposed limitations on supported DFT sizes as described in Section 7.1.1 there will be some additional constraints on N.

- Limited power variations in the time domain, leading to low cubic metric of the transmitted signal.

Furthermore, a sufficient number of reference-signal sequences of a given length, corresponding to a certain reference-signal bandwidth, should be available in order to avoid an unreasonable planning effort when assigning reference-signal sequences to devices and cells.

So-called *Zadoff–Chu* sequences [34] have the property of constant power in both the frequency and time domains. The M_{ZC} elements of the q:th Zadoff–Chu sequence within the set of Zadoff–Chu sequences of (odd) length M_{ZC} can be expressed as:

$$Z_k^q = e^{-j\pi q \frac{k\cdot(k+1)}{M_{ZC}}} \quad 0 \leq k < M_{ZC} \tag{7.1}$$

From the point of view of small power variations in both the frequency and time domains, Zadoff–Chu sequences would thus be excellent as uplink reference-signal sequences. However, there are two reasons why Zadoff–Chu sequences are not suitable for direct use as uplink reference-signal sequences in LTE:

- The number of available Zadoff–Chu sequences of a certain length, corresponding to the number of possible values for the parameter q in Eq. (7.1), equals the number of integers that are relative prime to the sequence length M_{ZC}. To maximize the number of Zadoff–Chu sequences and thus, in the end, to maximize the number of available uplink reference signals, prime-length Zadoff–Chu sequences would therefore be preferred. At the same time, according to above, the length of the uplink reference-signal sequences should be a multiple of 12, which is not a prime number.
- For short sequence lengths, corresponding to narrow uplink transmission bandwidths, relatively few reference-signal sequences would be available even if they were based on prime-length Zadoff–Chu sequences.

Instead, for sequence lengths larger than or equal to 36, corresponding to transmission bandwidths larger than or equal to three resource blocks, basic reference-signal sequences, in the LTE specification referred to as *base sequences*, are defined as *cyclic extensions* of Zadoff–Chu sequences of length M_{ZC} (Figure 7.11), where M_{ZC} is the largest prime number smaller than or equal to the length of the reference-signal sequence. For example, the largest prime number less than or equal to 36 is 31, implying that reference-signal sequences of length 36 are defined as cyclic extensions of Zadoff–Chu sequences of length 31. The number of available sequences is then equal to 30—that is, one less than the length of the Zadoff–Chu sequence. For larger sequence lengths, more sequences are available. For example, for a sequence length equal to 72, there are 70 sequences available.[7]

For sequence lengths equal to 12 and 24, corresponding to transmission bandwidths of one and two resource blocks, respectively, special QPSK-based sequences have instead been

[7]The largest prime number smaller than or equal to 72 is 71. The number of sequences is then one less than the length of the Zadoff–Chu sequence, that is, 70.

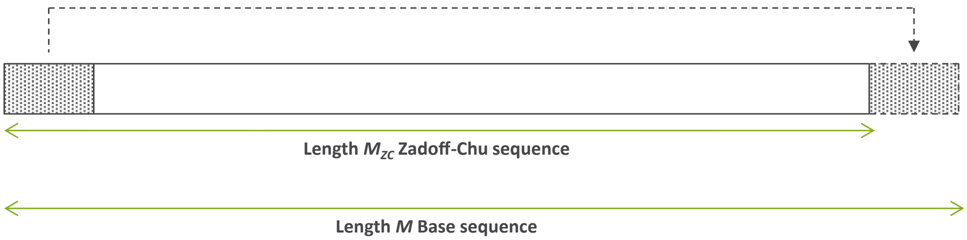

FIGURE 7.11

Length-M basic reference-signal sequence derived from cyclic extension of a length-M_{ZC} Zadoff–Chu sequence.

found from computer searches and are explicitly listed in the LTE specifications. For each of the two sequence lengths, 30 sequences have been defined.

Thus there are at least 30 sequences available for each sequence length. However, not all of these sequences are actually being used as base sequences:

- for sequence lengths less than 72, corresponding to reference-signal bandwidths less than six resource blocks, 30 sequences are being used;
- for sequence lengths equal to 72 and beyond, corresponding to reference-signal bandwidths of six resource blocks and beyond, 60 sequences are being used.

These sequences are divided into 30 *sequence groups* where each group consists of one base sequence for each sequence length less than 72 and two base sequences for each sequence length equal to 72 and beyond. A base sequence of a given length is thus fully specified by a *group index* ranging from 0 to 29 together with, in case of sequence lengths equal to 72 and beyond, a *sequence index* taking the values 0 and 1.

7.2.1.3 Phase-Rotation and Orthogonal Cover Codes

From the base sequences previously described, additional reference-signal sequences can be generated by applying different linear phase rotations in the frequency domain, as illustrated in Figure 7.12.

Applying a linear phase rotation in the frequency domain is equivalent to applying a cyclic shift in the time domain. Thus, although being *defined* as different frequency-domain phase rotations in line with Figure 7.12, the LTE specification actually refers to this as applying different *cyclic shifts*. Here the term "phase rotation" will be used. However, it should be borne in mind that what is here referred to as phase rotation is referred to as cyclic shift in the LTE specifications.

DM-RS derived from different base sequences typically have relatively low but still nonzero cross correlation. In contrast, reference signals defined from different phase rotations of the same base sequence are, at least in theory, completely orthogonal if the parameter α in

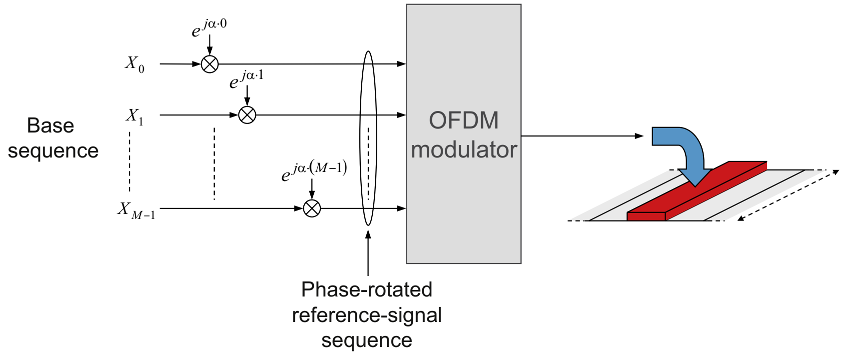

FIGURE 7.12

Generation of uplink reference-signal sequence from phase-rotated base sequence.

Figure 7.12 takes a value $m\pi/6$, where the integer m ranges from 0 to 11.[8] Up to 12 orthogonal reference signals can thus be derived from each base sequence by using different values of the parameter m.

However, to retain the orthogonality between these reference signals at the receiver side, the frequency response of the channel should be essentially constant over the span of one resource block—that is, over 12 subcarriers. Alternatively expressed, the main part of the channel time dispersion should not extend beyond the length of the cyclic shift mentioned in the preceding paragraphs. If that is not the case, a subset of the available values for α may be used—for example, only the values $\{0, 2\pi/6, 4\pi/6, \ldots, 10\pi/6\}$ or perhaps even fewer values. Limiting the set of possible values for α implies orthogonality over a smaller number of subcarriers and, as a consequence, less sensitivity to channel frequency selectivity. In other words, there is a trade-off between the number of orthogonal reference signals that can be generated from different phase rotations and the amount of channel frequency-selectivity that should be possible to cope with.

Another prerequisite for receiver-side orthogonality between reference signals defined from different phase rotations of the same base sequence is that the reference signals are transmitted well aligned in time. Thus the main uses of phase rotation include:

- to provide multiple simultaneous reference signals from the same device in case of uplink multi-layer transmission (uplink spatial multiplexing, see also Section 6.3.1);
- to provide the possibility for orthogonal reference signals between multiple devices being scheduled for PUSCH transmission on the same resource—that is, same set of resource blocks, within a cell (uplink MU-MIMO, see Section 6.3.2).

[8]The orthogonality is due to the fact that, for $\alpha = m\pi/6$, there will be an integer number of full-circle rotations over 12 subcarriers—that is, over one resource block.

Phase rotations may also be used to provide orthogonal reference signals between devices in neighbor cells, assuming tight synchronization between the cells. Finally, phase rotations are also used to separate reference signals of different devices in case of PUCCH transmission (see further Section 7.4).

In addition to the use of different phase rotations, orthogonal reference-signal transmissions can also be achieved by means of *Orthogonal Cover Codes* (OCC). As illustrated in Figure 7.13, two different length-two OCCs ([+1 +1] and [+1 −1], respectively) can be applied to the two PUSCH reference-signal transmissions within a subframe. This allows for overall DM-RS orthogonality over the subframe assuming that

- the channel does not change substantially over the subframe;
- the reference signals of the two slots are the same.[9]

Similar to phase rotations, receiver-side orthogonality between reference-signal transmissions based on different OCC requires that the transmissions are well-aligned in time at the receiver side. Thus the use of OCC is essentially the same as for phase rotations as described earlier:

- to provide multiple reference signals from the same device in case of uplink spatial multiplexing;
- to provide orthogonal reference signals between devices being scheduled on the same resource within a cell (uplink MU-MIMO);
- to allow for reference-signal orthogonality between uplink transmissions within neighbor cells in case of tight synchronization and time alignment between the cells.

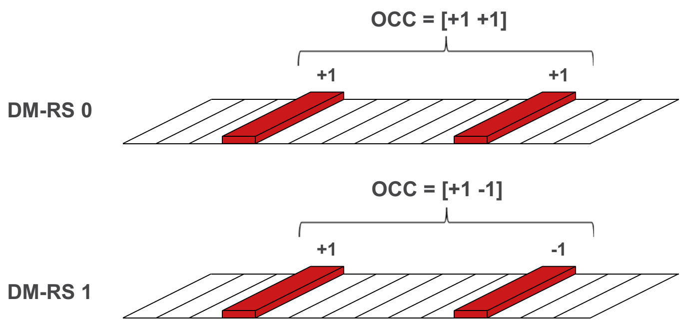

FIGURE 7.13

Generation of multiple DM-RS from orthogonal cover codes.

[9] Strictly speaking, the only thing required is that the correlation between the reference signals of DM-RS 0 and DM-RS 1 is the same for the two slots. If the reference signals are the same for the two slots, this is obviously the case.

It should be noted that, in contrast to phase rotations, orthogonality by means of OCC does not require that the same base sequence is used for the two DM-RS (DM-RS 0 and DM-RS 1 in Figure 7.13). Actually, the two reference signals do not even need to have the same bandwidth; having the same cross-correlation between the reference signals of DM-RS 0 and DM-RS 1 for the two slots is sufficient. Thus OCC can be used to achieve reference-signal orthogonality also for PUSCH transmissions of different bandwidths.

Similar to phase rotations, orthogonal codes can also be applied to DM-RS in case of PUCCH transmission, although in a somewhat different way compared to PUSCH due to the different time-domain structure of PUCCH DM-RS, see further Section 7.4.

7.2.1.4 Base-Sequence Assignment

According to the preceding discussion, each base sequence of a given length corresponds to a unique combination of a group index ranging from 0 to 29 and a sequence index taking the values 0 or 1. Base-sequence assignment—that is, determining which base sequence should be used by a specific device—is thus equivalent to assigning a corresponding group and sequence index.[10]

Prior to release 11, base-sequence assignment was *cell specific*—that is, for a given slot the group and sequence indices were the same for all devices having the same serving cell.

In the case of a *fixed* (*nonhopping*) group assignment, the sequence group to use for PUCCH transmission does not change between slots and was prior to release 11 directly given by the physical-layer cell identity. More specifically, the group index was equal to the cell identity modulo 30, where the cell identity may take values in the range 0 to 503 as described in Chapter 11. Thus, cell identities 0, 30, 60, …, 480 corresponded to sequence group 0, cell identities 1, 31, 61, …, 481 to sequence group 1, and so on.

In contrast, what sequence group to use for PUSCH transmission could be explicitly configured on a cell basis by adding an offset provided as part of the cell system information, to the PUCCH group index. The reason for providing the possibility for explicitly indicating what sequence group to use for PUSCH transmission in a cell was that it should be possible to use the same sequence group for PUSCH transmission in neighboring cells, despite the fact that such cells typically have different cell identities. In this case, the reference signals for PUSCH transmissions within the two cells would instead be distinguished by different phase rotations and/or OCC as discussed in Section 7.2.1.3, allowing for reference-signal orthogonality also *between* cells.[11]

In the case of *group hopping*, an additional cell-specific *group-hopping pattern* is added to the group index allowing for the group index of a cell to change on a slot basis. Prior to release 11, the group-hopping pattern was also derived from the cell identity and identical group-hopping patterns were used for PUSCH and PUCCH within a cell.

[10]For base-sequence lengths less than 72, the sequence index is always equal to zero.
[11]This assumes tight synchronization and time alignment between the cells.

In addition to the group index, for sequence lengths equal to or larger than 72, the reference-signal sequence also depends on the *sequence index*. The sequence index can either be fixed (nonhopping), in which case it always equals 0, or vary between slots (hopping) according to a *sequence-hopping pattern*. Similar to group hopping, prior to release 11 also the sequence-hopping pattern was cell-specific and given by the physical-layer cell identity.

In LTE release 11, the possibility for *device specific* base sequence assignment was introduced—that is, the group and sequence indices to use for PUSCH and PUCCH can, with release 11, be explicitly configured for a specific device regardless of the identity of the serving cell. The introduction of device-specific base-sequence assignment was done very similarly to the introduction of device-specific downlink DM-RS as described in Section 6.2.2—that is, by introducing the possibility for explicitly configuring a device with something that can be seen as a "virtual cell identity" which, if configured, replaces that actual physical-cell identity when deriving the group and sequence index. Similar to downlink DM-RS, if no virtual cell identity is configured the device should assume cell-specific base-sequence assignment according to the preceding discussion.

It should be pointed out that the device is not configured with a single "virtual-cell identity" to be used for deriving the base sequences for PUSCH and PUCCH. Rather, the device-specific configuration is done separately for PUSCH and PUCCH using two different "virtual cell identities."

The main reason for introducing the possibility for full device-specific assignment of uplink reference signals was to enhance the support for uplink *multi-point reception* (CoMP). In case of uplink multi-point reception, an uplink transmission may be received at a reception point not corresponding to the serving cell of the device but at a reception point corresponding to another cell. To enhance the reception quality it is beneficial to have orthogonal reference signals for different devices received at this reception point despite the fact that they, strictly speaking, have different serving cells. To allow for this, the reference-signal sequences should not be cell-specific but rather possible to assign on a device basis. Multi-point reception is discussed in somewhat more detail in Chapter 13 as part of a more general discussion on multi-point coordination and transmission.

7.2.1.5 Assignment of Phase Rotation and OCC

As discussed previously, the main use for phase rotations and OCC in case of PUSCH DM-RS is to provide a possibility for orthogonal reference signals for different layers in case of spatial multiplexing and for different devices scheduled on the same resource either within one cell (MU-MIMO) or in neighboring tightly synchronized cells.

In order not to limit the flexibility in terms of what devices can be jointly scheduled on the same resource, the assignment of phase rotation and OCC can be done dynamically. Thus, the exact phase rotation, given by the phase parameter m in Figure 7.12 and the OCC is provided jointly as a single parameter included as part of the uplink scheduling grant provided by the network. Each value of this parameter corresponds to a certain combination of phase rotation and OCC for each layer to be transmitted by the device. In case of spatial multiplexing, the different layers will then inherently be assigned different phase shifts and, possibly, different OCC. By providing different parameter values to different devices, the devices will be assigned different phase-shifts/OCC combinations allowing for orthogonal reference signals and thus providing enhanced MU-MIMO performance either within a cell or between cells.

7.2.2 SOUNDING REFERENCE SIGNALS

The DM-RS discussed in Section 7.2.1 are intended to be used by the base station for channel estimation to allow for coherent demodulation of uplink physical channels (PUSCH or PUCCH). A DM-RS is always transmitted together with and spanning the same frequency range as the corresponding physical channel.

In contrast, SRS are transmitted on the uplink to allow for the base station to estimate the uplink *channel state* at different frequencies. The channel-state estimates can then, for example, be used by the base-station scheduler to assign resource blocks of instantaneously good quality for uplink PUSCH transmission from the specific device (uplink channel-dependent scheduling). They can also be used to select different transmission parameters such as the instantaneous data rate and different parameters related to uplink multi-antenna transmission. The channel information obtained from the SRS can also be used for downlink transmission purposes exploiting channel reciprocity, for example downlink channel-dependent scheduling in TDD systems. As mentioned earlier, SRS transmission can also be used in other situations when uplink transmission is needed although there is no data to transmit, for example, for uplink timing estimation as part of the uplink-timing-alignment procedure. Thus, an SRS is not necessarily transmitted together with any physical channel and if transmitted together with, for example, PUSCH, the SRS may span a different, typically larger, frequency range.

There are two types of SRS transmission defined for the LTE uplink: *periodic* SRS transmission which has been available from the first release of LTE (release 8) and *aperiodic* SRS transmission introduced in LTE release 10.

7.2.2.1 Periodic SRS Transmission

Periodic SRS transmission from a device occurs at regular time intervals, from as often as once every 2 ms (every second subframe) to as infrequently as once every 160 ms (every 16th frame). When SRS is transmitted in a subframe, it occupies the last symbol of the subframe as illustrated in Figure 7.14. As an alternative, in the case of TDD operation, SRS can also be transmitted within the UpPTS.

In the frequency domain, SRS transmissions should span the frequency range of interest for the scheduler. This can be achieved in two ways:

- By means of a sufficiently wideband SRS transmission that allows for sounding of the entire frequency range of interest with a single SRS transmission as illustrated in the upper part of Figure 7.15.

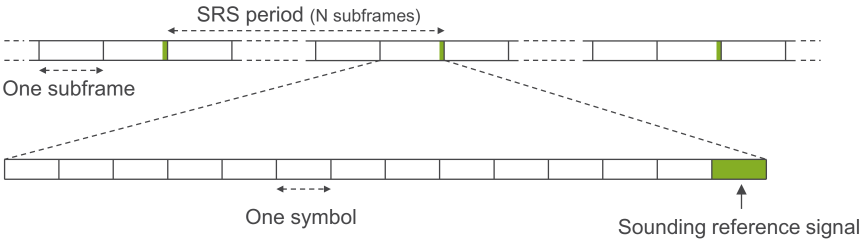

FIGURE 7.14

Transmission of SRS.

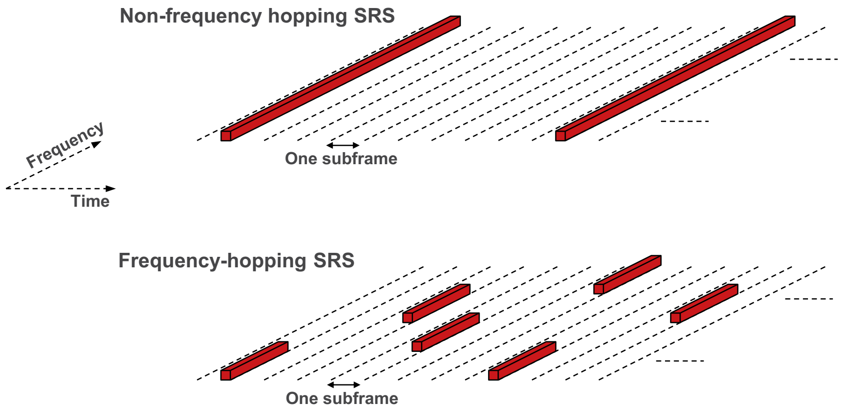

FIGURE 7.15

Nonfrequency-hopping (wideband) versus frequency-hopping SRS.

- By means of a more narrowband SRS transmission that is hopping in the frequency domain in such a way that a sequence of SRS transmissions jointly spans the frequency range of interest, as illustrated in the lower part of Figure 7.15.

The instantaneous SRS bandwidth is always a multiple of four resource blocks. Different bandwidths of the SRS transmission can simultaneously be available within a cell. A narrow SRS bandwidth, corresponding to four resource blocks, is always available, regardless of the uplink cell bandwidth. Up to three additional, more wideband SRS bandwidths may also be configured within a cell. A specific device within the cell is then explicitly configured to use one of the up to four SRS bandwidths available in the cell.

If a device is transmitting SRS in a certain subframe, the SRS transmission may very well overlap, in the frequency domain, with PUSCH transmissions from other devices within the cell. To avoid collision between SRS and PUSCH transmissions from different devices, devices should in general avoid PUSCH transmission in the OFDM symbols in which SRS transmission may occur. To achieve this, all devices within a cell are aware of the set of subframes within which SRS *may* be transmitted by *any* device within the cell. All devices should then avoid PUSCH transmission in the last OFDM symbol of those subframes. Information about the set of subframes in which SRS may be transmitted within a cell is provided as part of the cell system information.[12]

On a more detailed level, the structure for SRS is similar to that of uplink DM-RS described in Section 7.2.1. More specifically, a SRS is also defined as a frequency-domain reference-signal sequence derived in the same way as for DM-RS—that is, cyclic extension of prime-length Zadoff−Chu sequences for sequence lengths equal to 30 and above and special sequences for sequence lengths less than 30. However, in the case of SRS, the reference-signal sequence is mapped to *every second* subcarrier, creating a "comb"-like spectrum, as illustrated in Figure 7.16. Taking into account that the bandwidth of the SRS transmission is always a multiple of four resource blocks, the lengths of the reference-signal sequences for SRS are thus always multiples of 24.[13] The reference-signal sequence to use for SRS transmission within the cell is derived in the same way as the PUCCH DM-RS within the cell, assuming cell-specific reference-signal sequence assignment. Device-specific reference-signal sequences are not supported for SRS.

Starting from release 13, up to four different combs can be used instead of two as in previous releases, subject to higher-layer configuration. In case of four different combs, every fourth subcarrier is used instead of every second. The purpose of this is to increase the SRS multiplexing capacity to handle the increased number of antennas supported with the introduction of FD-MIMO, see Chapter 10.

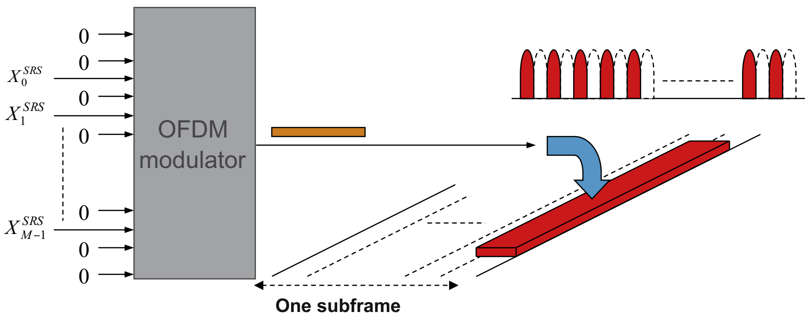

FIGURE 7.16

Generation of SRS from a frequency-domain reference-signal sequence.

[12]What is provided as part of the system information is a periodicity (2 to 160 ms) and a subframe offset, compare the following bullet list.

[13]Four resource blocks, each spanning 12 subcarriers but only every second subcarrier used for a certain SRS transmission.

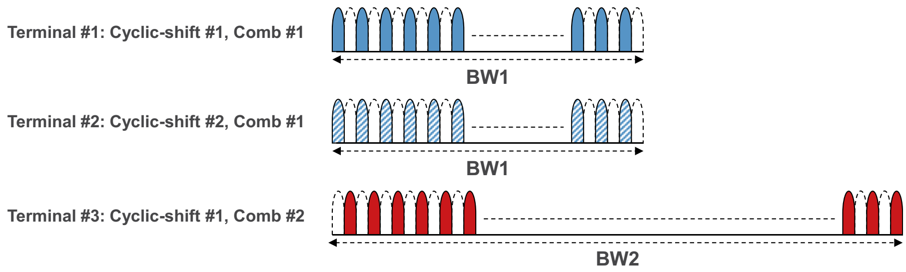

FIGURE 7.17

Multiplexing of SRS transmissions from different devices.

Similar to DM-RS, different phase rotations, also for SRS referred to as "cyclic shifts" in the LTE specifications, can be used to generate different SRS that are orthogonal to each other. By assigning different phase rotations to different devices, multiple SRS can thus be transmitted in parallel in the same subframe, as illustrated by devices #1 and #2 in the upper part of Figure 7.17. However, it is then required that the reference signals span the same frequency range.

Another way to allow for SRS to be simultaneously transmitted from different devices is to rely on the fact that each SRS only occupies every second (or every fourth) subcarrier. Thus, SRS transmissions from two devices can be *frequency multiplexed* by assigning them to different frequency shifts or "combs," as illustrated by device #3 in the lower part of Figure 7.17. In contrast to the multiplexing of SRS transmission by means of different "cyclic shifts," frequency multiplexing of SRS transmissions does not require the transmissions to cover identical frequency ranges.

To summarize, the following set of parameters defines the characteristics of an SRS transmission:

- SRS transmission time-domain period (from 2 to 160 ms) and subframe offset.
- SRS transmission bandwidth—the bandwidth covered by a single SRS transmission.
- Hopping bandwidth—the frequency band over which the SRS transmission is frequency hopping.
- Frequency-domain position—the starting point of the SRS transmission in the frequency domain.
- Transmission comb as illustrated in Figure 7.17.
- Phase rotation (or equivalently cyclic shift) of the reference-signal sequence.
- Number of combs (introduced in release 13).

A device that is to transmit SRS is configured with these parameters by means of higher layer (RRC) signaling.

7.2.2.2 Aperiodic SRS Transmission

In contrast to periodic SRS, aperiodic SRS are *one-shot* transmissions triggered by signaling on PDCCH as part of the scheduling grant. The frequency-domain structure of an aperiodic SRS transmission is identical to that of periodic SRS. Also, in the same way as for periodic SRS transmission, aperiodic SRS are transmitted within the last symbol of a subframe. Furthermore, the time instants when aperiodic SRS may be transmitted are configured per device using higher-layer signaling.

The frequency-domain parameters for aperiodic SRS (bandwidth, odd or even "comb," and so on) are configured by higher-layer (RRC) signaling. However, no SRS transmission will actually be carried out until the device is explicitly triggered to do so by an explicit *SRS trigger* on PDCCH/EPDCCH. When such a trigger is received, a single SRS is transmitted in the next available aperiodic SRS instant configured for the device using the configured frequency-domain parameters. Additional SRS transmissions can then be carried out if additional triggers are received.

Three different parameter sets can be configured for aperiodic SRS, for example differing in the frequency position of the SRS transmission and/or the transmission comb. Information on what parameters to use when the SRS is actually transmitted is included in the PDCCH/EPDCCH L1/L2 control signaling information, which consists of two bits, three combinations of which indicate the specific SRS parameter set. The fourth combination simply dictates that no SRS should be transmitted.

7.3 UPLINK MULTI-ANTENNA TRANSMISSION

Downlink multi-antenna transmission was supported by the LTE specification from its first release (release 8). With LTE release 10, support for uplink multi-antenna transmission—that is, uplink transmission relying on multiple transmit antennas at the device side—was also introduced for LTE. Uplink multi-antenna transmission can be used to improve the uplink link performance in different ways:

- to improve the achievable data rates and spectral efficiency for uplink data transmission by allowing for antenna precoding supporting uplink beam-forming as well as spatial multiplexing with up to four layers for the uplink physical data channel PUSCH;
- to improve the uplink control-channel performance by allowing for transmit diversity for the PUCCH.

7.3.1 PRECODER-BASED MULTI-ANTENNA TRANSMISSION FOR PUSCH

As illustrated in Figure 7.18, the structure of the uplink antenna precoding is very similar to that of downlink antenna precoding (Section 6.3), including the presence of precoded DM-RS (one per layer) similar to downlink non-codebook-based precoding (Figure 6.18). Uplink

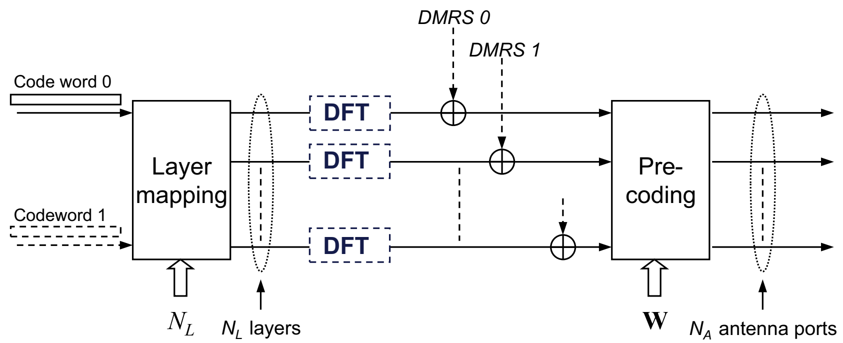

FIGURE 7.18

Precoder-based multi-antenna transmission for LTE uplink.

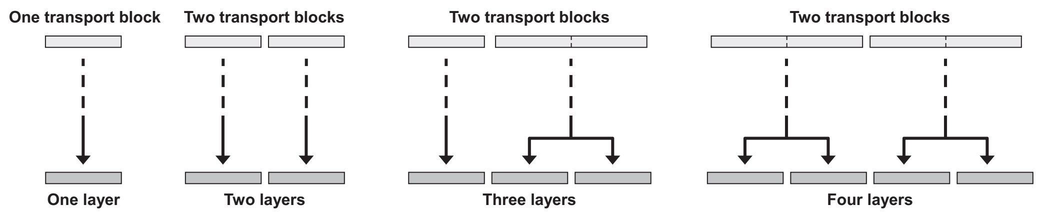

FIGURE 7.19

Uplink transport-channel-to-layer mapping (initial transmission).

antenna precoding supports transmission using up to four antenna ports, allowing for spatial multiplexing with up to four layers.

The principles for mapping of the modulation symbols to layers are also the same as for the downlink. For an initial transmission, there is one transport block in the case of a single layer and two transport blocks for more than one layer, as illustrated in Figure 7.19. Similar to the downlink, in the case of a hybrid-ARQ retransmission, a single transport block may also be transmitted on multiple layers in some cases.

As can be seen in Figure 7.18, the DFT precoding is actually taking place after layer mapping—that is, each layer is separately DFT precoded. To simplify the description this was not really visible in Figure 7.1 outlining the overall physical-layer transport-channel processing.

It can also be noted that, in contrast to Figure 6.19, the precoder in Figure 7.18 is not shaded. As discussed in Section 6.3.3, for downlink non-codebook-based precoding, the precoder part of the antenna mapping is not visible in the specification and the network can, in essence, apply an arbitrary precoding for the downlink transmission. Due to the use of

Table 7.1 Uplink Precoder Matrices for Two Antenna Ports

Transmission Rank	Codebook Index					
	0	1	2	3	4	5
1	$\frac{1}{\sqrt{2}}\begin{bmatrix}1\\1\end{bmatrix}$	$\frac{1}{\sqrt{2}}\begin{bmatrix}1\\-1\end{bmatrix}$	$\frac{1}{\sqrt{2}}\begin{bmatrix}1\\j\end{bmatrix}$	$\frac{1}{\sqrt{2}}\begin{bmatrix}1\\-j\end{bmatrix}$	$\frac{1}{\sqrt{2}}\begin{bmatrix}1\\0\end{bmatrix}$	$\frac{1}{\sqrt{2}}\begin{bmatrix}0\\1\end{bmatrix}$
2	$\frac{1}{\sqrt{2}}\begin{bmatrix}1 & 0\\0 & 1\end{bmatrix}$	–	–	–	–	–

precoded DM-RS, the device can recover the different layers without knowledge of exactly what precoding has been applied at the transmitted side.

The same is also true for the uplink—that is, the presence of precoded DM-RS would allow for the base station to demodulate the uplink multi-antenna transmission and recover the different layers without knowledge of the precoding taking place at the transmitter side. However, for LTE the uplink precoder matrix is selected by the network and conveyed to the device as part of the scheduling grant. The device should then follow the precoder matrix selected by the network. Thus, in the uplink, the precoder is visible in the specification and, in order to limit the downlink signaling, there is a limited set of precoder matrices specified for each transmission rank.

More specifically, for each combination of transmission rank N_L and number of antennas ports N_A, a set of precoder matrices of size $N_A \times N_L$ is defined, as illustrated in Tables 7.1 and 7.2 for two and four antenna ports, respectively. For full-rank transmission—that is, when the transmission rank or number of layers equals the number of transmit antennas—only a single precoder matrix is defined, namely the identity matrix of size $N_A \times N_A$ (not shown in the tables). Note that, for the case of four antenna ports, only a subset of the defined matrices is shown. In total there are 24 rank-1 matrices, 16 rank-2 matrices, and 12 rank-3 matrices defined for four antenna ports, in addition to the single rank-4 matrix.

As can be seen, all the precoder matrices in Table 7.1 contain one and only one nonzero element in each row, and this is generally true for all precoder matrices defined for the uplink. As a consequence, the signal transmitted on a certain antenna port (corresponding to a certain row of the precoder matrix) always depends on one and only one specific layer (corresponding to a specific column of the precoder matrix). Expressed alternatively, the precoder matrix maps the layers to the antenna ports *with at most one layer being mapped to each antenna port*. Due to this, the good cubic-metric properties of the transmitted signal are also preserved for each antenna port when antenna precoding is applied. The precoder matrices of Tables 7.1 and 7.2 are therefore also referred to as *cubic-metric-preserving precoder matrices*.

In order to select a suitable precoder, the network needs information about the uplink channel. Such information can, for example, be based on measurements on the uplink SRS

Table 7.2 Subset of Uplink Precoder Matrices for Four Antenna Ports and Different Transmission Ranks

	Codebook Index				
Transmission Rank	**0**	**1**	**2**	**3**	…
1	$\frac{1}{2}\begin{bmatrix}1\\1\\1\\-1\end{bmatrix}$	$\frac{1}{2}\begin{bmatrix}1\\1\\j\\j\end{bmatrix}$	$\frac{1}{2}\begin{bmatrix}1\\1\\-1\\1\end{bmatrix}$	$\frac{1}{2}\begin{bmatrix}1\\1\\-j\\-j\end{bmatrix}$	…
2	$\frac{1}{2}\begin{bmatrix}1&0\\1&0\\0&1\\0&-j\end{bmatrix}$	$\frac{1}{2}\begin{bmatrix}1&0\\1&0\\0&1\\0&j\end{bmatrix}$	$\frac{1}{2}\begin{bmatrix}1&0\\-j&0\\0&1\\0&1\end{bmatrix}$	$\frac{1}{2}\begin{bmatrix}1&0\\-j&0\\0&1\\0&-1\end{bmatrix}$	…
3	$\frac{1}{2}\begin{bmatrix}1&0&0\\1&0&0\\0&1&0\\0&0&1\end{bmatrix}$	$\frac{1}{2}\begin{bmatrix}1&0&0\\-1&0&0\\0&1&0\\0&0&1\end{bmatrix}$	$\frac{1}{2}\begin{bmatrix}1&0&0\\0&1&0\\1&0&0\\0&0&1\end{bmatrix}$	$\frac{1}{2}\begin{bmatrix}1&0&0\\0&1&0\\-1&0&0\\0&0&1\end{bmatrix}$	…
4	$\frac{1}{2}\begin{bmatrix}1&0&0&0\\0&1&0&0\\0&0&1&0\\0&0&0&1\end{bmatrix}$	–	–	–	…

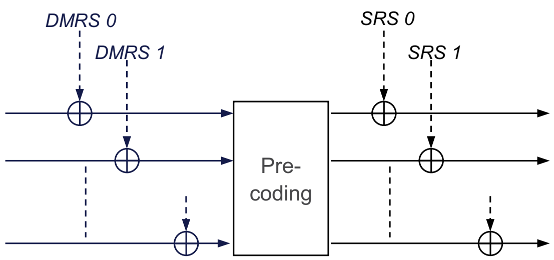

FIGURE 7.20

Illustration of SRS transmitted *after* uplink antenna precoding.

(Section 7.2.2). As indicated in Figure 7.20 SRS are transmitted non-precoded—that is, directly on the different antenna ports. The received SRS thus reflect the channel of each antenna port, not including any precoding. Based on the received SRS, the network can thus decide on a suitable uplink transmission rank and corresponding uplink precoder matrix, and provide information about the selected rank and precoder matrix as part of the scheduling grant.

The previous paragraph assumed the same number of antenna ports for PUSCH as for SRS. This is a relevant situation and the SRS is, in this case, used to aid the selection of the precoding matrix, as discussed in the preceding paragraphs. However, there are also situations when SRS and PUSCH use *different* numbers of antenna ports. One example is uplink transmission of two layers (two antenna ports), where the eNodeB would like to use SRS to probe the channel for potential four-layer transmission. In this case the SRS is transmitted on a *different* set of antenna ports than the PUSCH to aid the eNodeB in assessing the benefits, if any, of switching to four-layer transmission.

7.3.2 UPLINK MULTI-USER MIMO

As described in Section 6.3.5, downlink multi-user MIMO (MU-MIMO) implies downlink transmission to different devices using *the same time–frequency resource* and relying on the availability of multiple antennas, at least on the network side, to suppress interference between the transmissions. The term MU-MIMO originated from the resemblance to SU-MIMO (spatial multiplexing).

Uplink MU-MIMO is essentially the same thing but for the uplink transmission direction—that is, uplink MU-MIMO implies uplink transmissions from multiple devices using *the same uplink time–frequency resource* and relying on the availability of multiple receive antennas at the base station to separate the two or more transmissions. Thus, MU-MIMO is really just another term for uplink *space-division multiple access* (SDMA).

Actually, on the uplink, the relation between MU-MIMO and SU-MIMO (spatial multiplexing) is even closer. Uplink spatial multiplexing, for example with two antenna ports and two layers, implies that the device transmits two transport blocks with one transport block transmitted on each layer and thus on each antenna port,[14] as illustrated in the left part of Figure 7.21. As illustrated in the right part of the figure, MU-MIMO is essentially equivalent to separating the two antennas into two different devices and transmitting one transport block from each device. The base-station processing to separate the two

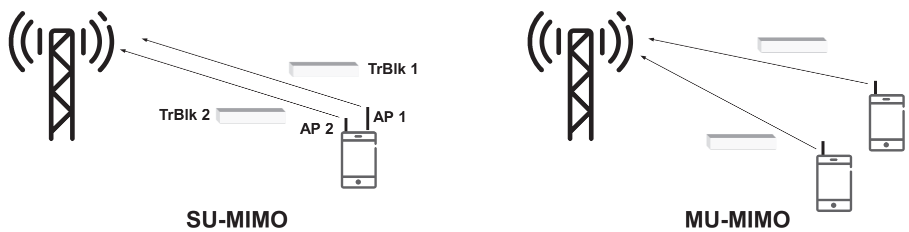

FIGURE 7.21

SU-MIMO and MU-MIMO.

[14]Note that the 2 × 2 precoder matrix is the identity matrix, see Table 7.1.

transmissions could essentially be identical to the processing used to separate the two layers in the case of spatial multiplexing. It should be noted that the separation of the two transmissions at the receiver side could be simplified, or at least the possible means to achieve this separation are extended, if the two devices are well separated in space, something which is not the case for two antennas attached to the same device. As an example, for sufficiently separated devices, classical beam-forming relying on correlated receiver antennas can be used to separate the uplink transmissions. Alternatively, uncorrelated receiver antennas can be used, and the separation means are then essentially the same as for SU-MIMO.

One important benefit of uplink MU-MIMO is that one, in many cases, can get similar gains in *system throughput* as SU-MIMO (spatial multiplexing) without the need for multiple transmit antennas at the device side, allowing for less complex device implementation. It should be noted though that spatial multiplexing could still provide substantial gains in terms of *user throughput* and peak data rates that can be provided from a single device. Furthermore, the potential system gains of uplink MU-MIMO rely on the fact that more than one device is actually available for transmission to in a subframe. The process of "pairing" devices that should share the time–frequency resources is also nontrivial and requires suitable radio-channel conditions.

Essentially, support for uplink MU-MIMO only requires the possibility to explicitly assign a specific orthogonal reference signal for the uplink transmission, thereby ensuring orthogonality between reference-signal transmissions from the different devices involved in the MU-MIMO transmission. As described in Section 7.2.1.5, this is supported by means of the dynamic assignment of DM-RS phase rotation and OCC as part of the uplink scheduling grant.

7.3.3 PUCCH TRANSMIT DIVERSITY

Precoder-based multi-layer transmission is only used for the uplink data transmission on PUSCH. However, in the case of a device with multiple transmit antennas, one obviously wants to use the full set of device antennas and corresponding device power amplifiers also for the L1/L2 control signaling on PUCCH in order to be able to utilize the full power resource and achieve maximum diversity. To achieve additional diversity, LTE release 10 also introduced the possibility for two-antenna *transmit diversity* for PUCCH. More specifically, the transmit diversity supported for PUCCH is referred to as *spatial orthogonal-resource transmit diversity* (SORTD).

The basic principle of SORTD is simply to transmit the uplink control signaling using different resources (time, frequency, and/or code) on the different antennas. In essence, the PUCCH transmissions from the two antennas will be identical to PUCCH transmissions from two different devices using different resources. Thus, SORTD creates additional diversity but achieves this by using twice as many PUCCH resources, compared to non-SORTD transmission.

For four physical antennas at the device, implementation-specific *antenna virtualization* is used. In essence, a transparent scheme is used to map the two-antenna-port signal to four physical antennas.

7.4 UPLINK L1/L2 CONTROL SIGNALING

Similar to the LTE downlink, there is also a need for uplink L1/L2 control signaling to support data transmission on downlink and uplink transport channels. Uplink L1/L2 control signaling consists of:

- hybrid-ARQ acknowledgments for received DL-SCH transport blocks;
- channel-state information (CSI) related to the downlink channel conditions, used to assist downlink scheduling; and
- scheduling requests, indicating that a device needs uplink resources for UL-SCH transmission.

There is no information indicating the UL-SCH transport format signaled on the uplink. As mentioned in Chapter 4, the eNodeB is in complete control of the uplink UL-SCH transmissions and the device always follows the scheduling grants received from the network, including the UL-SCH transport format specified in those grants. Thus, the network knows the transport format used for the UL-SCH transmission in advance and there is no need for any explicit transport-format signaling on the uplink.

Uplink L1/L2 control signaling needs to be transmitted on the uplink regardless of whether or not the device has any uplink transport-channel data to transmit and thus regardless of whether or not the device has been assigned any uplink resources for UL-SCH transmission. Hence, two different methods are supported for the transmission of the uplink L1/L2 control signaling, depending on whether or not the device has been assigned an uplink resource for UL-SCH transmission:

- *Nonsimultaneous transmission of UL-SCH and L1/L2 control.* If the device does not have a valid scheduling grant—that is, no resources have been assigned for the UL-SCH in the current subframe—a separate physical channel, the PUCCH, is used for transmission of uplink L1/L2 control signaling.
- *Simultaneous transmission of UL-SCH and L1/L2 control.* If the device has a valid scheduling grant—that is, resources have been assigned for the UL-SCH in the current subframe—the uplink L1/L2 control signaling is time multiplexed with the coded UL-SCH on to the PUSCH prior to DFT precoding and OFDM modulation. As the device has been assigned UL-SCH resources, there is no need to support transmission of the scheduling request in this case. Instead, scheduling information can be included in the MAC headers, as described in Chapter 13.

The reason to differentiate between the two cases is to minimize the cubic metric for the uplink power amplifier in order to maximize coverage. However, in situations when there is sufficient power available in the device, simultaneous transmission of PUSCH and PUCCH can be used with no impact on the coverage. The possibility for simultaneous PUSCH and PUCCH transmission was therefore introduced in release 10 as one part of several features[15] adding flexibility at the cost of a somewhat higher cubic metric. In situations where this cost is not acceptable, simultaneous PUSCH and PUCCH can always be avoided by using the basic mechanism introduced in the first version of LTE.

In the following section, the basic PUCCH structure and the principles for PUCCH control signaling are described, followed by control signaling on PUSCH.

7.4.1 BASIC PUCCH STRUCTURE

If the device has not been assigned an uplink resource for UL-SCH transmission, the L1/L2 control information (CSI reports, hybrid-ARQ acknowledgments, and scheduling requests) is transmitted on uplink resources (resource blocks) specifically assigned for uplink L1/L2 control on PUCCH. Transmission of control signaling on PUCCH is characterized by the PUCCH format used.

The first LTE releases, release 8 and release 9, defined two main PUCCH formats:[16]

- PUCCH format 1, carrying 0, 1, or 2 bits of information and used for hybrid-ARQ acknowledgments and scheduling requests;
- PUCCH format 2, carrying up to 11 bits of control information and used for reporting CSI.

There is one PUCCH per device. Given the relatively small payload size of PUCCH formats 1 and 2, the bandwidth of one resource block during one subframe is too large for the control signaling needs of a single device. Therefore, to efficiently exploit the resources set aside for control signaling, multiple devices can share the same resource-block pair.

With the introduction of carrier aggregation in release 10, where up to five component carriers can be aggregated, the number of hybrid-ARQ acknowledgments increased. To address this situation, an additional PUCCH format was introduced,

- PUCCH format 3, carrying up to 22 bits of control information.

Carrier aggregation was extended to handle up to 32 component carriers in release 13, calling for an even higher PUCCH capacity which was solved by

- PUCCH format 4, capable of a large number of hybrid-ARQ acknowledgments by using multiple resource-block pairs, and
- PUCCH format 5, capable of a payload between PUCCH formats 3 and 4.

---

[15]Other examples of such features are simultaneous transmission on multiple uplink component carriers and uplink multi-cluster transmission.

[16]There are actually three subtypes each of format 1 and 2, see further the detailed description.

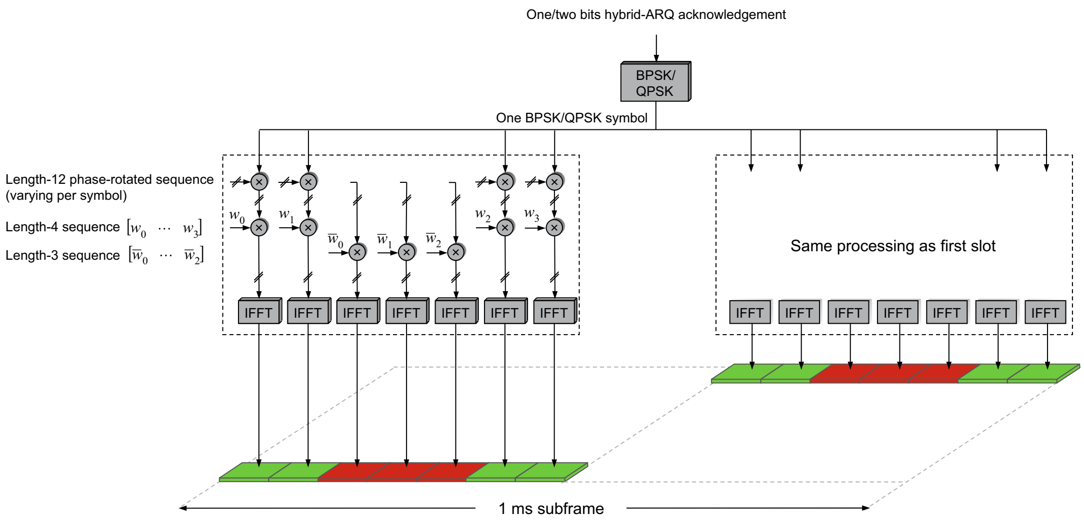

FIGURE 7.22

PUCCH format 1 (normal cyclic prefix).

The detailed structure of each of the different PUCCH formats is discussed in the following section, followed by an overview on how the formats are used.

7.4.1.1 PUCCH Format 1

PUCCH format 1,[17] for transmission of hybrid-ARQ acknowledgments and scheduling requests, is capable of carrying up to two bits of information. The same structure is used in the two slots of a subframe as illustrated in Figure 7.22. For transmission of a hybrid-ARQ acknowledgment, the one or two hybrid-ARQ acknowledgment bits are used to generate a BPSK or QPSK symbol, respectively. For a scheduling request, the same constellation point as for a negative acknowledgment is used. The modulation symbol is then used to generate the signal to be transmitted in each of the two PUCCH slots.

Different devices sharing the same resource-block pair in a subframe are separated by different orthogonal phase rotations of a length-12 frequency-domain sequence, where the sequence is identical to a length-12 reference-signal sequence. Furthermore, as described in conjunction with the reference signals in Section 7.2, a linear phase rotation in the frequency domain is equivalent to applying a cyclic shift in the time domain. Thus, although the term "phase rotation" is used here, the term cyclic shift is sometimes used with an implicit reference to the time domain. Similarly to the case of reference signals, there are up

[17]There are actually three variants in the LTE specifications, formats 1, 1a, and 1b, used for transmission of scheduling requests and one or two hybrid-ARQ acknowledgments, respectively. However, for simplicity, they are all referred to as format 1 herein.

to 12 different phase rotations specified, providing up to 12 different orthogonal sequences from each base sequence.[18] However, in the case of frequency-selective channels, not all 12 phase rotations can be used if orthogonality is to be retained. Typically, up to six rotations are considered usable in a cell from a radio-propagation perspective, although inter-cell interference may result in a smaller number being useful from an overall system perspective. Higher-layer signaling is used to configure the number of rotations that are used in a cell.

There are seven OFDM symbols per slot for a normal cyclic prefix (six in the case of an extended cyclic prefix). In each of those seven OFDM symbols, a length-12 sequence, obtained by phase rotation of the base sequence as described earlier, is transmitted. Three of the symbols are used as reference signals to enable channel estimation by the eNodeB and the remaining four[19] are modulated by the BPSK/QPSK symbols described earlier. In principle, the BPSK/QPSK modulation symbol could directly modulate the rotated length-12 sequence used to differentiate devices transmitting on the same time–frequency resource. However, this would result in unnecessarily low capacity on the PUCCH. Therefore, the BPSK/QPSK symbol is multiplied by a length-4 orthogonal cover sequence.[20] Multiple devices may transmit on the same time–frequency resource using the same phase-rotated sequence and be separated through different orthogonal covers. To be able to estimate the channels for the respective devices, the reference signals also employ an orthogonal cover sequence, with the only difference being the length of the sequence—three for the case of a normal cyclic prefix. Thus, since each base sequence can be used for up to $3 \cdot 12 = 36$ different devices (assuming all 12 rotations are available; typically at most six of them are used), there is a threefold improvement in the PUCCH capacity compared to the case of no cover sequence. The cover sequences are three Walsh sequences of length 4 for the data part and three DFT sequences of length 3 for the reference signals.

A PUCCH format 1 resource, used for either a hybrid-ARQ acknowledgment or a scheduling request, is represented by a single scalar resource index. From the index, the phase rotation and the orthogonal cover sequence are derived.

The use of a phase rotation of a base sequence together with orthogonal sequences as described earlier provides orthogonality between different devices in the same cell transmitting PUCCH on the same set of resource blocks. Hence, in the ideal case, there will be no intra-cell interference, which helps improve the performance. However, there will typically

[18]In releases 8 to 10, the base sequence is cell-specific, while release 11 adds the possibility of configuring a "virtual cell identity" from which the base sequence is derived. See further the discussion in conjunction with uplink reference signals in Section 7.2.
[19]The number of symbols used for reference signals and the acknowledgment is a trade-off between channel-estimation accuracy and energy in the information part; three symbols for reference symbols and four symbols for the acknowledgment have been found to be a good compromise.
[20]In the case of simultaneous SRS and PUCCH transmissions in the same subframe, a length-3 sequence is used, thereby making the last OFDM symbol in the subframe available for the SRS.

be inter-cell interference for the PUCCH as the different sequences used in neighboring cells are nonorthogonal. To randomize the inter-cell interference, the phase rotation of the sequence used in a cell varies on a symbol-by-symbol basis in a slot according to a hopping pattern derived from the physical-layer cell identity of the primary carrier. In release 11, the randomization can be configured to use the virtual cell identity instead of the physical one. On top of this, slot-level hopping is applied to the orthogonal cover and phase rotation to further randomize the interference. This is exemplified in Figure 7.23 assuming normal cyclic prefix and six of 12 rotations used for each cover sequence. To the phase rotation given by the cell-specific hopping a slot-specific offset is added. In cell A, a device is transmitting on PUCCH resource number 3, which in this example corresponds to using the (phase rotation, cover sequence) combination (6, 0) in the first slot and (11, 1) in the second slot of this particular subframe. PUCCH resource number 11, used by another device in cell A transmitting in the same subframe, corresponds to (11, 1) and (8, 2) in the first and second slots, respectively, of the subframe. In another cell the PUCCH resource numbers are mapped to different sets (rotation, cover sequence) in the slots. This helps to randomize the inter-cell interference.

For an extended cyclic prefix, the same structure as in Figure 7.22 is used with the difference being the number of reference symbols in each slot. In this case, the six OFDM symbols in each slot are divided such that the two middle symbols are used for reference signals and the remaining four symbols used for the information. Thus, the length of the orthogonal sequence used to spread the reference symbols is reduced from 3 to 2 and the multiplexing capacity is lower. However, the general principles described in the preceding paragraphs still apply.

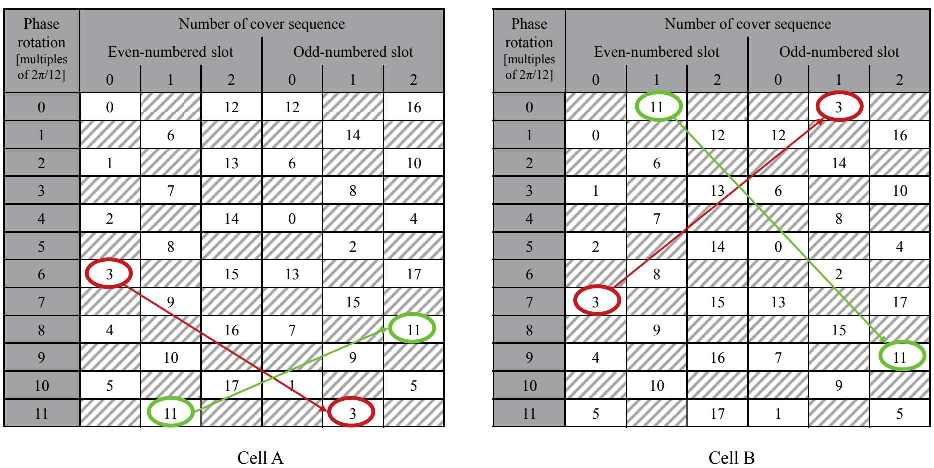

FIGURE 7.23

Example of phase rotation and cover hopping for two PUCCH resource indices in two different cells.

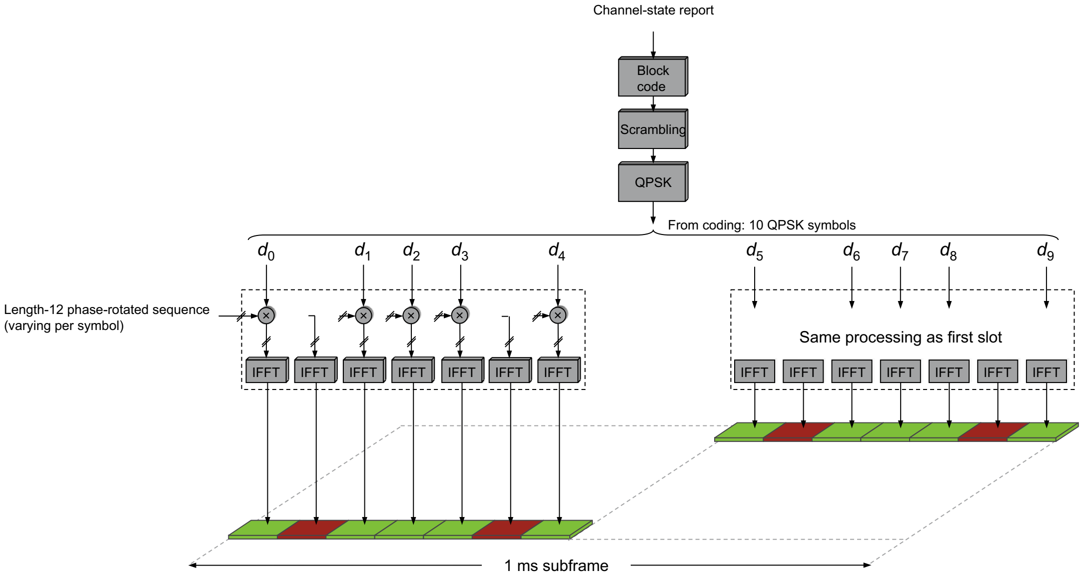

FIGURE 7.24

PUCCH format 2 (normal cyclic prefix).

7.4.1.2 PUCCH Format 2

PUCCH format 2, used for CSI reporting, is capable of handling up to 11 information bits per subframe.[21] Similarly to PUCCH format 1, multiple devices, using the same resource-block pair in a subframe, are separated through different orthogonal phase rotations of a length-12 sequence as illustrated for normal cyclic prefix in Figure 7.24. After block coding using a punctured Reed–Müller code and QPSK modulation, there are 10 QPSK symbols to transmit in the subframe: the first five symbols are transmitted in the first slot and the remaining five in the last slot.

Assuming a normal cyclic prefix, there are seven OFDM symbols per slot. Of the seven OFDM symbols in each slot, two[22] are used for reference-signal transmission to allow coherent demodulation at the eNodeB. In the remaining five, the respective QPSK symbol to be transmitted is multiplied by a phase-rotated length-12 base sequence and the result is transmitted in the corresponding OFDM symbol. For an extended cyclic prefix, where there

[21]There are actually three variants in the LTE specifications, formats 2, 2a, and 2b, where the last two formats are used for simultaneous transmission of CSI reports and hybrid-ARQ acknowledgments as discussed later in this section. However, for simplicity, they are all referred to as format 2 here.

[22]Similarly to format 1, the number of symbols used for reference signals and the coded channel-quality information is a trade-off between channel-estimation accuracy and energy in the information part. Two symbols for reference symbols and five symbols for the coded information part in each slot were found to be the best compromise.

are six OFDM symbols per slot, the same structure is used but with one reference-signal symbol per slot instead of two.

Basing the format 2 structure on phase rotations of the same base sequence as format 1 is beneficial as it allows the two formats to be transmitted in the same resource block. As phase-rotated sequences are orthogonal, one rotated sequence in the cell can be used either for one PUCCH instance using format 2 or three PUCCH instances using format 1. Thus, the "resource consumption" of one CSI report is equivalent to three hybrid-ARQ acknowledgments (assuming normal cyclic prefix). Note that no orthogonal cover sequences are used for format 2.

The phase rotations to use in the different symbols for PUCCH format 2 are hopping in a similar way as for format 1, motivated by interference randomization. Resources for PUCCH format 2 can, similar to format 1, be represented by a scalar index, which can be seen as a "channel number."

7.4.1.3 PUCCH Format 3

For downlink carrier aggregation, see Chapter 12, multiple hybrid-ARQ acknowledgment bits need to be fed back in the case of simultaneous transmission on multiple component carriers. Although PUCCH format 1 with resource selection can be used to handle the case of two downlink component carriers, this is not sufficient as a general solution as carrier aggregation of up to five component carriers is part of release 10. PUCCH format 3 was therefore introduced in release 10 to enable the possibility of transmitting up to 22 bits of control information on PUCCH in an efficient way. A device capable of more than two downlink component carriers—that is, capable of more than four bits for hybrid-ARQ acknowledgments—needs to support PUCCH format 3. For such a device, PUCCH format 3 can also be used for less than four bits of feedback relating to simultaneous transmission on multiple component carriers if configured by higher-layer signaling not to use PUCCH format 1 with resource selection.

The basis for PUCCH format 3, illustrated in Figure 7.25, is DFT-precoded OFDM—that is, the same transmission scheme as used for UL-SCH. The acknowledgment bits, one or two per downlink component carrier depending on the transmission mode configured for that particular component carrier, are concatenated with a bit reserved for scheduling request into a sequence of bits where bits corresponding to unscheduled transport blocks are set to zero. Block coding is applied,[23] followed by scrambling to randomize inter-cell interference. The resulting 48 bits are QPSK-modulated and divided into two groups, one per slot, of 12 QPSK symbols each.

Assuming a normal cyclic prefix, there are seven OFDM symbols per slot. Similarly to PUCCH format 2, two OFDM symbols (one in the case of an extended cyclic prefix) in each slot are used for reference-signal transmission, leaving five symbols for data transmission. In each slot, the block of 12 DFT-precoded QPSK symbols is transmitted in the

[23] A (32,k) Reed−Müller code is used, but for twelve or more bits in TDD two Reed−Müller codes are used in combination.

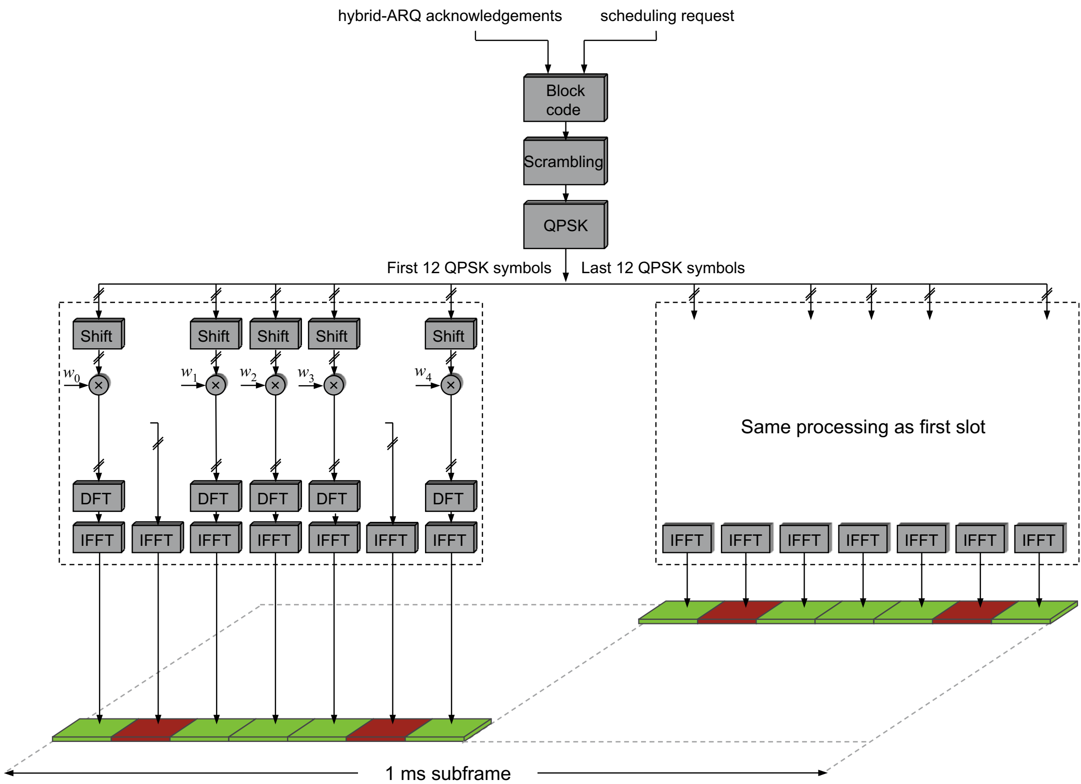

FIGURE 7.25

PUCCH format 3 (normal cyclic prefix).

five available DFTS-OFDM symbols. To further randomize the inter-cell interference, a cyclic shift of the 12 inputs to the DFT, varying between OFDM symbols in a cell-specific manner, is applied to the block of 12 QPSK symbols prior to DFT precoding (in releases 11 and later, the cyclic shift can be based on the virtual cell identity instead of the physical one).

To increase the multiplexing capacity, a length-5 orthogonal sequence is used with each of the five OFDM symbols carrying data in a slot being multiplied by one element of the sequence. Thus, up to five devices may share the same resource-block pair for PUCCH format 3. Different length-5 sequences are used in the two slots to improve the performance in high-Doppler scenarios. To facilitate channel estimation for the different transmissions sharing the same resource block, different reference-signal sequences are used.

The length-5 orthogonal cover sequences are obtained as five DFT sequences. There is also the possibility to use a length-4 Walsh sequence for the second slot in order to leave the last OFDM symbol unused for the case when sounding is configured in the subframe.

In the same manner as for the other two PUCCH formats, a resource can be represented by a single index from which the orthogonal sequence and the resource-block number can be derived.

Note that, due to the differences in the underlying structure of PUCCH format 3 compared to the two previous formats, resource blocks cannot be shared between format 3 and formats 1 and 2.

7.4.1.4 PUCCH Format 4

With the extension of carrier aggregation to handle up to 32 component carriers, the payload capacity of PUCCH format 3 is not sufficient to handle the resulting number of hybrid-ARQ acknowledgments. PUCCH formats 4 and 5 were introduced in release 13 to address this problem.

PUCCH format 4, illustrated in Figure 7.26, is to a large extent modeled after the PUSCH processing with a single DFT precoder covering multiple resource-block pairs. An 8-bit CRC is added to the payload, followed by tailbiting convolutional coding and rate matching to match the number of coded bits to the number of available resource elements. Scrambling,

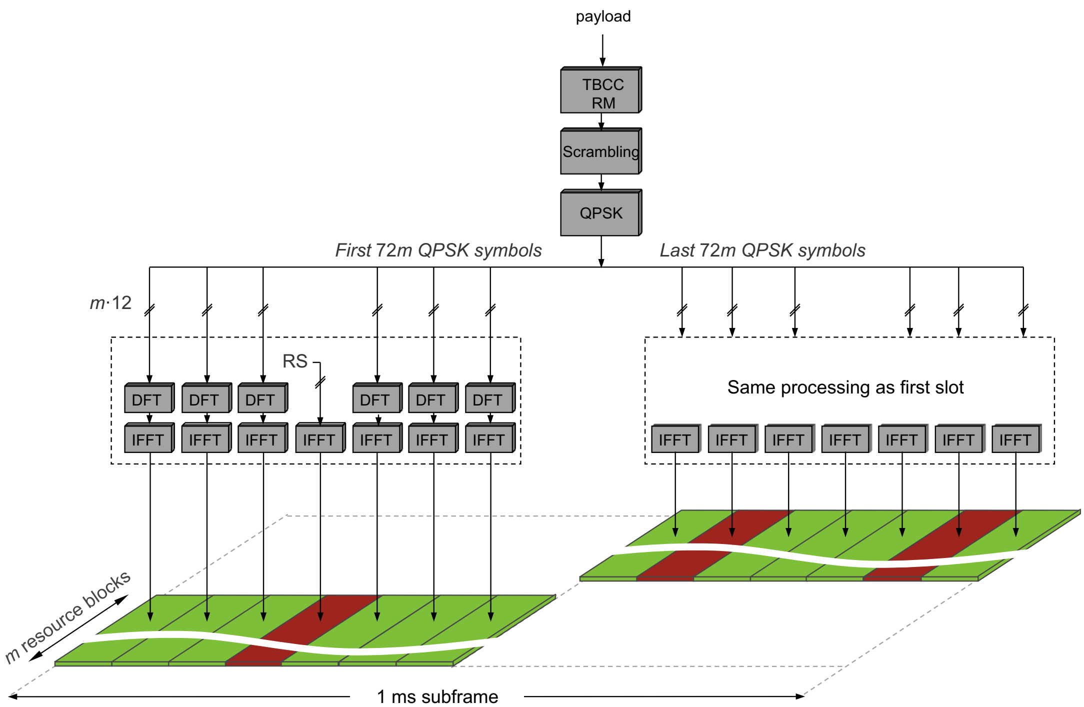

FIGURE 7.26

PUCCH format 4 (normal cyclic prefix).

QPSK modulation, DFT-precoding, and mapping to resource elements follow the same structure as the PUSCH—that is each DFT-spread OFDM symbol carries separate sets of coded bits. Multiple resource blocks in the frequency domain, 1, 2, 3, 4, 5, 6, or 8, can be used for PUCCH format 4, allowing for a very large payload. Inter-slot frequency hopping is used, similar to the other PUCCH formats and both normal and extended cyclic prefix is supported. There is also a possibility for a shortened format, leaving the last OFDM symbol in the subframe unused, for the case when sounding is configured in the subframe.

7.4.1.5 PUCCH Format 5

PUCCH format 5 is, similar to format 4, used to handle a large number of hybrid-ARQ feedback bits. However, as it uses a single resource-block pair, the payload capacity is smaller than PUCCH format 4. In addition, format 5 multiplexes two users onto a single resource-block pair as shown in Figure 7.27, making it suitable for efficient handling of payloads larger than format 3 but smaller than format 4.

The channel coding including CRC attachment and rate matching, and the QPSK modulation is identical to PUCCH format 4. However, the resource-block mapping and the usage of spreading is different. Each DFT-spread OFDM symbol carries six QPSK symbols. Prior to

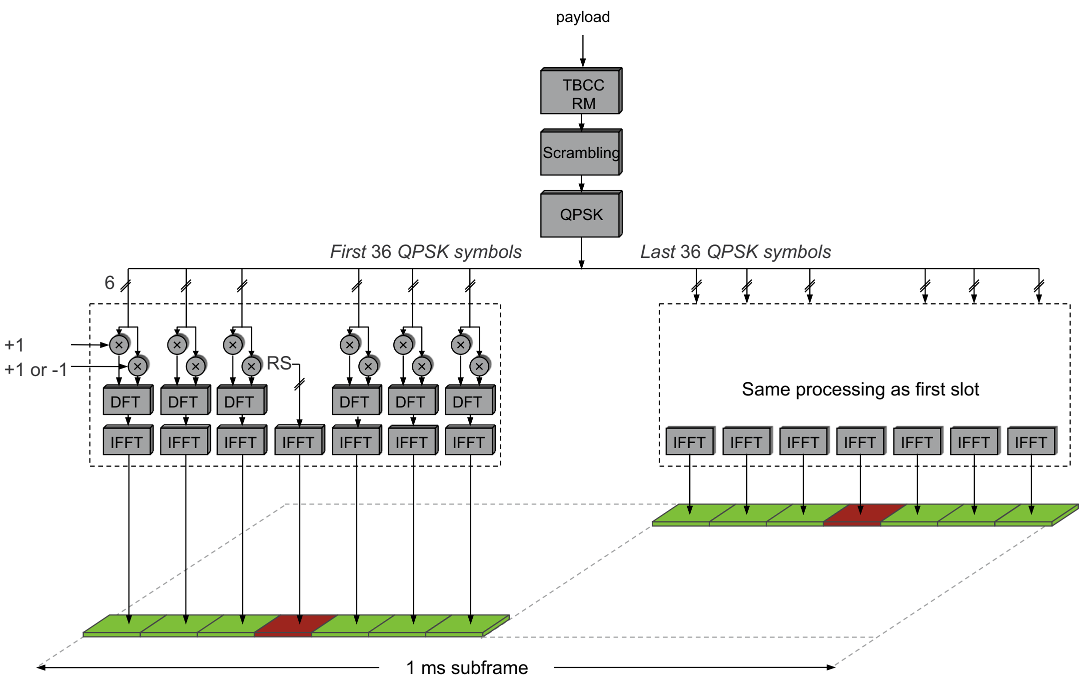

FIGURE 7.27

PUCCH format 5 (cyclic prefix).

DFT precoding, the six QPSK symbols are block-wise repeated where the second block is multiplied with +1 or −1 depending on which of the two orthogonal sequences are used. Hence, by using different orthogonal sequences, two users may share the same resource-block pair and transmit 144 coded bits each.[24]

The two users sharing the same resource-block pair use mutually orthogonal reference-signal sequences.

7.4.1.6 Resource-Block Mapping for PUCCH

The signals described for all of the PUCCH formats are, as already explained, transmitted on a (set of) resource-block pair. The resource-block pair to use is determined from the PUCCH resource index. Multiple resource-block pairs can be used to increase the control-signaling capacity in the cell; when one resource-block pair is full, the next PUCCH resource index is mapped to the next resource-block pair in sequence.

The resource-block pair(s) where a PUCCH is transmitted is located at the edges of the bandwidth allocated to the primary component carrier[25] as illustrated in Figure 7.28. To provide frequency diversity, frequency hopping on the slot boundary is used—that is, one "frequency resource" consists of 12 (or more in case of PUCCH format 4) subcarriers at the upper part of the spectrum within the first slot of a subframe and an equally sized resource at the lower part of the spectrum during the second slot of the subframe (or vice versa).

The reason for locating the PUCCH resources at the edges of the overall available spectrum is twofold:

- Together with the frequency hopping described previously, this maximizes the frequency diversity experienced by the control signaling.
- Assigning uplink resources for the PUCCH at other positions within the spectrum—that is, not at the edges—would have fragmented the uplink spectrum, making it impossible

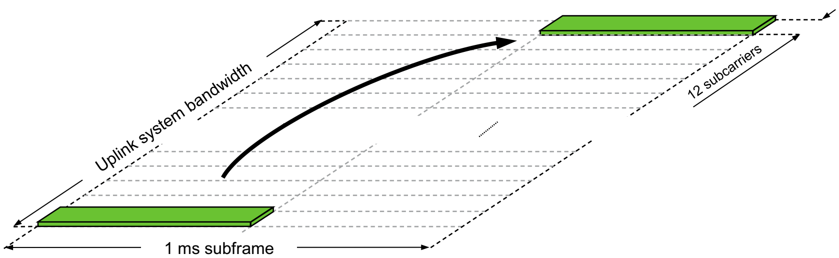

FIGURE 7.28

Uplink L1/L2 control signaling transmission on PUCCH.

[24]There are 12 OFDM symbols, each carrying 6 QPSK symbols for one user, resulting in $12 \cdot 6 \cdot 2 = 144$ bits for normal CP.

[25]Note that the primary component carrier in the uplink is specified on a per-device basis. Hence, different devices may view different carriers as their primary component carrier.

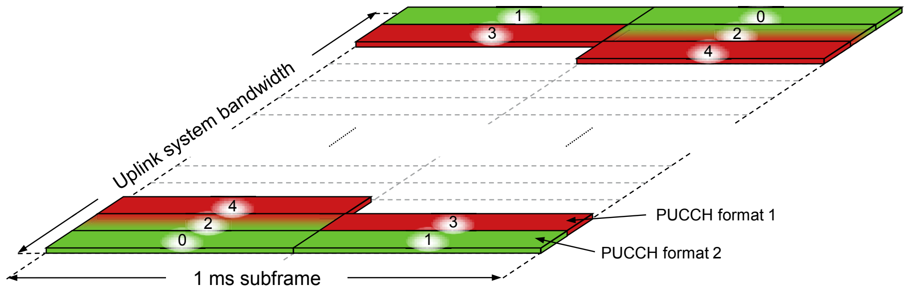

FIGURE 7.29

Allocation of resource blocks for PUCCH.

to assign very wide transmission bandwidths to a single device and still preserve the low-cubic-metric properties of the uplink transmission.

The resource-block mapping is in principle done such that PUCCH format 2 (CSI reports) is transmitted closest to the edges of the uplink cell bandwidth with PUCCH format 1 (hybrid-ARQ acknowledgments, scheduling requests) next as illustrated in Figure 7.29. The locations of PUCCH formats 3, 4, and 5 are configurable and can, for example, be located between formats 1 and 2. A semi-static parameter, provided as part of the system information, controls on which resource-block pair the mapping of PUCCH format 1 starts. Furthermore, the semi-statically configured scheduling requests are located at the outermost parts of the format 1 resources, leaving dynamic acknowledgments closest to the data. As the amount of resources necessary for hybrid-ARQ acknowledgments varies dynamically, this maximizes the amount of contiguous spectrum available for PUSCH.

In many scenarios, the configuration of the PUCCH resources can be done such that the three PUCCH formats are transmitted on separate sets of resource blocks. However, for the smallest cell bandwidths, this would result in too high an overhead. Therefore, it is possible to mix PUCCH formats 1 and 2 in one of the resource-block pairs—for example, in Figure 7.29 this is the case for the resource-block pair denoted "2." Although this mixture is primarily motivated by the smaller cell bandwidths, it can equally well be used for the larger cell bandwidths. In the resource-block pair where PUCCH formats 1 and 2 are mixed, the set of possible phase rotations are split between the two formats. Furthermore, some of the phase rotations are reserved as "guard"; hence the efficiency of such a mixed resource-block pair is slightly lower than a resource-block pair carrying only one of the first two PUCCH formats.

7.4.2 UPLINK CONTROL SIGNALING ON PUCCH

Having described the three PUCCH formats, the details on how these different formats are used to convey uplink control information can be discussed. As already mentioned, uplink control signaling on PUCCH can in principle be any combination of hybrid-ARQ

Table 7.3 Usage of Different PUCCH Formats for Different Pieces of Information (the Superscripts 10, 11, and 13 Denote the First Release Supporting this Combination)

	PUCCH Format						
	Format 1						
Information		**Selection**	**Bundling**	**Format 2**	**Format 3**	**Format 4**	**Format 5**
ACK	•	•[10]			•[10]	•[13]	•[13]
SR	•						
SR + ACK	•		•[10]		•[10]	•[13]	•[13]
CSI				•		•[13]	•[13]
CSI + ACK				•	•[11]	•[13]	•[13]
CSI + SR						•[13]	•[13]
CSI + SR + ACK					•[11]	•[13]	•[13]

acknowledgments (ACK), CSI, and scheduling requests (SR). Depending on whether these pieces of information are transmitted alone or in combination, different PUCCH formats and mechanisms are used as summarized in Table 7.3. In principle, simultaneous transmission of multiple control signaling messages from a single device could use multiple PUCCHs. However, this would increase the cubic metric and a single PUCCH structure supporting simultaneous transmission of multiple feedback signals is used instead.

7.4.2.1 Hybrid-ARQ Acknowledgments

Hybrid-ARQ acknowledgments are used to acknowledge receipt of one (or two in the case of spatial multiplexing) transport blocks on the DL-SCH. PUCCH format 1 is used in absence of carrier aggregation but can also support carrier aggregation of up to two downlink carriers—that is, up to four acknowledgment bits—as discussed in the following. PUCCH formats 3, 4, or 5 are used for more than four acknowledgments bits.

The hybrid-ARQ acknowledgment is only transmitted when the device correctly received control signaling related to DL-SCH transmission intended for this device on an PDCCH or EPDCCH. If no valid DL-SCH-related control signaling is detected, then nothing is transmitted on the PUCCH (that is, DTX). Apart from not unnecessarily occupying PUCCH resources that can be used for other purposes, this allows the eNodeB to perform three-state detection, ACK, NAK, or DTX, on the PUCCH received when using PUCCH format 1. Three-state detection is useful as NAK and DTX may need to be treated differently. In the case of NAK, retransmission of additional parity bits is useful for incremental redundancy, while for DTX the device has most likely missed the initial transmission of systematic bits and a better alternative than transmitting additional parity bits is to retransmit the systematic bits.

Transmission of one or two hybrid-ARQ acknowledgment bits uses PUCCH format 1. As mentioned in Section 7.4.1.1, a PUCCH resource can be represented by an index. How to determine this index depends on the type of information and whether the PDCCH or the EPDCCH was used to schedule the downlink data transmission.

For PDCCH-scheduled downlink transmissions, the resource index to use for a hybrid-ARQ acknowledgment is given as a function of the first CCE in the PDCCH used to schedule the downlink transmission to the device. In this way, there is no need to explicitly include information about the PUCCH resources in the downlink scheduling assignment, which of course reduces overhead. Furthermore, as described in Chapter 8, hybrid-ARQ acknowledgments are transmitted a fixed time after the reception of a DL-SCH transport block and when to expect a hybrid ARQ on the PUCCH is therefore known to the eNodeB.

For EPDCCH-scheduled transmissions, the index of the first ECCE in the EPDCCH cannot be used alone. Since the ECCE numbering is configured per device and therefore is device-specific, two different devices with control signaling on different resource blocks may have the same number of the first ECCE in the EPDCCH. Therefore, the ACK/NAK resource offset (ARO) being part of the EPDCCH information (see Section 6.4.6) is used in addition to the index of the first ECCE to determine the PUCCH resource. In this way, PUCCH collisions between two devices scheduled with EPDCCH can be avoided.

In addition to dynamic scheduling by using the (E)PDCCH, there is also, as described in Chapter 9, the possibility to semi-persistently schedule a device according to a specific pattern. In this case there is no PDCCH or EPDCCH to derive the PUCCH resource index from. Instead, the configuration of the semi-persistent scheduling pattern includes information on the PUCCH index to use for the hybrid-ARQ acknowledgment. In either of these cases, a device is using PUCCH resources only when it has been scheduled in the downlink. Thus, the amount of PUCCH resources required for hybrid-ARQ acknowledgments does not necessarily increase with an increasing number of devices in the cell, but, for dynamic scheduling, is rather related to the number of CCEs in the downlink control signaling.

The description in the preceding paragraphs addressed the case of downlink carrier aggregation not being used. Extensions to handle carrier aggregation are covered in Chapter 12.

7.4.2.2 Scheduling Request

Scheduling requests are used to request uplink resources for data transmission. Obviously, a scheduling request should only be transmitted when the device is requesting resources, otherwise the device should be silent to save battery resources and not create unnecessary interference.

Unlike the hybrid-ARQ acknowledgment, whose occurrence is known to the eNodeB from the downlink scheduling decision, the need for uplink resources for a certain device is in principle unpredictable by the eNodeB. One way to handle this would be to have a contention-based mechanism for requesting uplink resources. The random-access mechanism is based on this principle and can, to some extent, also be used for scheduling requests, as discussed in Chapter 9. Contention-based mechanisms typically work well for low intensities, but for higher scheduling-request intensities, the collision rate between different devices simultaneously requesting resources becomes too large. Therefore, LTE provides a contention-free scheduling-request mechanism on the PUCCH, where each device in the cell is given a reserved resource on which it can transmit a request for uplink resources. Unlike hybrid-ARQ acknowledgments, no explicit information bit is transmitted by the scheduling request; instead the information is conveyed by the presence (or absence) of energy on the corresponding PUCCH. However, the scheduling request, although used for a completely different purpose, shares the same PUCCH format as the hybrid-ARQ acknowledgment, namely PUCCH format 1.

The contention-free scheduling-request resource is represented by a PUCCH format 1 resource index as described earlier, occurring at every *n*th subframe. The more frequently these time instants occur, the lower the scheduling-request delay at the cost of higher PUCCH resource consumption. As the eNodeB configures all the devices in the cell, when and on which resources a device can request resources is known to the eNodeB. A single scheduling request resource is also sufficient for the case of carrier aggregation, as it only represents a request for uplink resources, which is independent of whether carrier aggregation is used or not.

7.4.2.3 Hybrid-ARQ Acknowledgments and Scheduling Request

The discussion in the two previous sections concerned transmission of *either* a hybrid-ARQ acknowledgment *or* a scheduling request. However, there are situations when the device needs to transmit *both* of them.

If PUCCH format 1 is used for the acknowledgments, simultaneous transmission of the acknowledgments and scheduling request is handled by transmitting the hybrid-ARQ acknowledgment on the scheduling-request resource (see Figure 7.30). This is possible as the same PUCCH structure is used for both of them and the scheduling request carries no explicit information. By comparing the amount of energy detected on the acknowledgment resource and the scheduling-request resource for a specific device, the eNodeB can determine whether or not the device is requesting uplink data resources. Once the PUCCH resource used for transmission of the acknowledgment is detected, the hybrid-ARQ acknowledgment can be decoded. Other, more advanced methods jointly decoding hybrid-ARQ and scheduling request can also be envisioned.

Channel selection, which is a way to transmit up to four acknowledgments on PUCCH in absence of a simultaneous scheduling request, cannot be used for joint transmission of

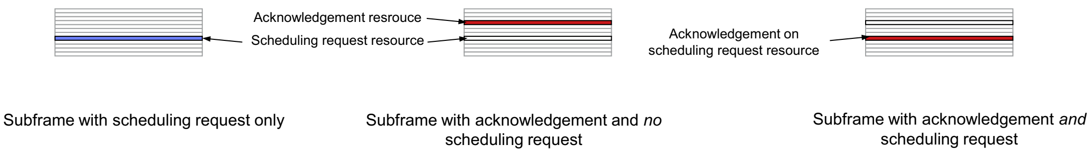

FIGURE 7.30

Multiplexing of scheduling request and hybrid-ARQ acknowledgment from a single device.

acknowledgments and scheduling request. Instead, up to four acknowledgment bits are bundled (combined) into two bits that are transmitted as described in the previous paragraph. Bundling implies that two or more acknowledgment bits are combined into a smaller number of bits. In essence, one acknowledgment bit represents the decoding outcome of multiple transport blocks and all these transport blocks need to be retransmitted as soon as one of them is incorrectly received.

PUCCH formats 3 to 5 support joint coding of acknowledgments and scheduling request in conjunction with carrier aggregation as described in Chapter 12.

7.4.2.4 Channel-State Information

CSI reports, the contents of which are discussed in Chapter 10, are used to provide the eNodeB with an estimate of the downlink radio-channel properties as seen from the device to aid channel-dependent scheduling. A CSI report consists of multiple bits transmitted in one subframe. There are *two* types of CSI reports, namely

- *periodic* reports, occurring at regular time instants;
- *aperiodic* reports, triggered by downlink control signaling on the PDCCH (or EPDCCH).

Aperiodic reports can only be transmitted on PUSCH as described later in Section 7.4.3, while periodic reports can be transmitted on PUCCH using PUCCH format 2.

A PUCCH format 2 resource is represented by an index and higher-layer signaling is used to configure each device with a resource to transmit its CSI report on, as well as when those reports should be transmitted. Hence, the eNodeB has full knowledge of when and on which resources each of the devices will transmit CSI on PUCCH.

7.4.2.5 Hybrid-ARQ Acknowledgments and CSI

Transmission of data in the downlink implies transmission of hybrid-ARQ acknowledgments in the uplink. At the same time, since data is transmitted in the downlink, up-to-date CSI is beneficial to optimize the downlink transmissions. Hence, simultaneous transmission of hybrid-ARQ acknowledgments and CSI is supported by LTE.

The handling of simultaneous transmission of acknowledgments and a CSI report depends on the number of acknowledgment bits as well as higher-layer configuration. There is also the possibility to configure the device to drop the CSI report and only transmit the acknowledgments.

The basic way of supporting transmission of one or two acknowledgments simultaneously with CSI, part of releases 8 and later, is based on PUCCH format 2, although the detailed solution differs between the two.

For a normal cyclic prefix, each slot in PUCCH format 2 has two OFDM symbols used for reference signals. When transmitting a hybrid-ARQ acknowledgment at the same time as CSI report, the second reference signal in each slot is modulated by the acknowledgment, as illustrated in Figure 7.31(a). Either BPSK or QPSK is used, depending on whether one or two acknowledgment bits are to be fed back. The fact that the acknowledgment is superimposed

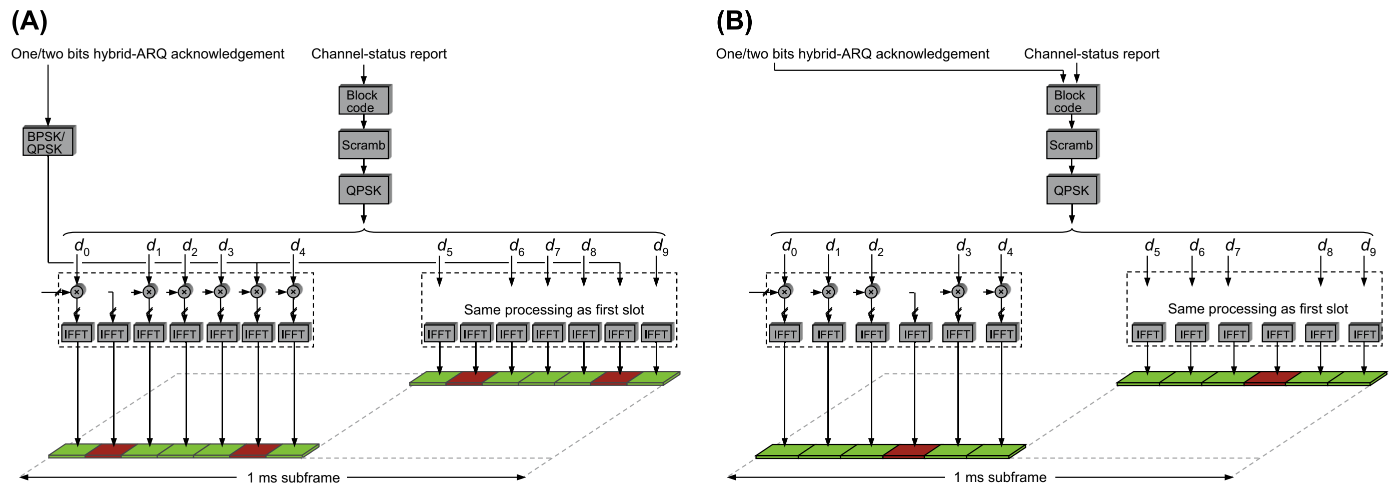

FIGURE 7.31

Simultaneous transmission of CSI and hybrid-ARQ acknowledgments: (A) normal cyclic prefix and (B) extended cyclic prefix.

on the reference signal needs to be accounted for at the eNodeB. One possibility is to decode the acknowledgment bit(s) modulated on to the second reference symbol using the first reference symbol for channel estimation. Once the acknowledgment bit(s) have been decoded, the modulation imposed on the second reference symbol can be removed and channel estimation and decoding of the CSI report can be handled in the same way as in the absence of simultaneous hybrid-ARQ acknowledgment. This two-step approach works well for low to medium Doppler frequencies; for higher Doppler frequencies the acknowledgment and CSI reports are preferably decoded jointly.

For an extended cyclic prefix, there is only a single reference symbol per slot. Hence, it is not possible to overlay the hybrid-ARQ acknowledgment on the reference symbol. Instead, the acknowledgment bit(s) are jointly coded with the CSI report prior to transmission using PUCCH format 2, as illustrated in Figure 7.31(b).

The time instances for which to expect CSI reports and hybrid-ARQ acknowledgments are known to the eNodeB, which therefore knows whether to expect a hybrid-ARQ acknowledgment along with the CSI report or not. If the (E)PDCCH assignment is missed by the device, then only the CSI report will be transmitted as the device is not aware that it has been scheduled. In the absence of a simultaneous CSI report, the eNodeB can employ DTX detection to discriminate between a missed assignment and a failed decoding of downlink data. However, one consequence of the structures described is that DTX detection is cumbersome, if not impossible. This implies that incremental redundancy needs to be operated with some care if the eNodeB has scheduled data such that the acknowledgment occurs at the same time as a CSI report. As the device may have missed the original transmission attempt in the downlink, it may be preferable for the eNodeB to select the redundancy version of the retransmission such that systematic bits are also included in the retransmission.

One possibility to circumvent this is to configure the device to drop the CSI report in the case of simultaneous transmission of a hybrid-ARQ acknowledgment. In this case, the eNodeB can detect DTX as the acknowledgment can be transmitted using PUCCH format 1, as described earlier. There will be no CSI report sent in this case, which needs to be taken into account in the scheduling process.

PUCCH formats 3 to 5 support joint coding of acknowledgments and scheduling request, see Chapter 10, and there is no need to drop the CSI report in this case (unless the device is configured to do so).

7.4.2.6 Scheduling Request and CSI

The eNodeB is in control of when a device may transmit a scheduling request and when the device should report the channel state. Hence, simultaneous transmission of scheduling requests and channel-state information can be avoided by proper configuration. If this is not done the device drops the CSI report and transmits the scheduling request only. Missing a CSI report is not detrimental and only incurs some degradation in the scheduling and rate-adaptation accuracy, whereas the scheduling request is critical for uplink transmissions.

7.4.2.7 Hybrid-ARQ Acknowledgments, CSI and Scheduling Request

For devices not supporting or configured to use PUCCH format 3, 4, or 5, simultaneous transmission of acknowledgments, CSI, and scheduling request are handled similarly to the description in the previous section; the CSI report is dropped and the acknowledgments and scheduling request are multiplexed as previously described in Section 7.4.2.3. However, devices that are using PUCCH format 3 or higher for multiple acknowledgments support simultaneous transmission of all three pieces of information. Since there is a bit reserved for scheduling requests in this case, the transmission structure is no different from the case of simultaneous transmission of acknowledgments and CSI reports described in Section 7.4.2.5.

7.4.3 UPLINK L1/L2 CONTROL SIGNALING ON PUSCH

If the device is transmitting data on PUSCH—that is, has a valid scheduling grant in the subframe—control signaling is time multiplexed[26] with data on the PUSCH instead of using the PUCCH (in release 10 and later simultaneous PUSCH and PUCCH can be used, avoiding the need for control signaling on PUSCH for most cases at the cost of a somewhat worse cubic metric). Only hybrid-ARQ acknowledgments and CSI reports are transmitted on the PUSCH. There is no need to request a scheduling grant when the device is already scheduled; instead, in-band buffer-status reports are sent as part of the MAC headers, as described in Chapter 9.

Time multiplexing of CSI reports and hybrid-ARQ acknowledgments is illustrated in Figure 7.32. However, although they both use time multiplexing there are some differences in the details for the two types of uplink L1/L2 control signaling motivated by their different properties.

The hybrid-ARQ acknowledgment is important for proper operation of the downlink. For one and two acknowledgments, robust QPSK modulation is used, regardless of the modulation scheme used for the data, while for a larger number of bits the same modulation scheme as for the data is used. Channel coding for more than two bits is done in the same way as for the PUCCH and bundling is applied if the number of bits exceeds a limit—that is, two transport blocks on the same component carrier share a single bit instead of having independent bits. Furthermore, the hybrid-ARQ acknowledgment is transmitted near to the reference symbols as the channel estimates are of better quality close to the reference symbols. This is especially important at high Doppler frequencies, where the channel may vary during a slot. Unlike the data part, the hybrid-ARQ acknowledgment cannot rely on retransmissions and strong channel coding to handle these variations.

In principle, the eNodeB knows when to expect a hybrid-ARQ acknowledgment from the device and can therefore perform the appropriate demultiplexing of the acknowledgment and the data parts. However, there is a certain probability that the device has missed the scheduling assignment on the downlink control channels (PDCCH or EPDCCH), in which case the

[26]In the case of spatial multiplexing, the CQI/PMI is time multiplexed with one of the code words, implying that it is spatially multiplexed with the other code word.

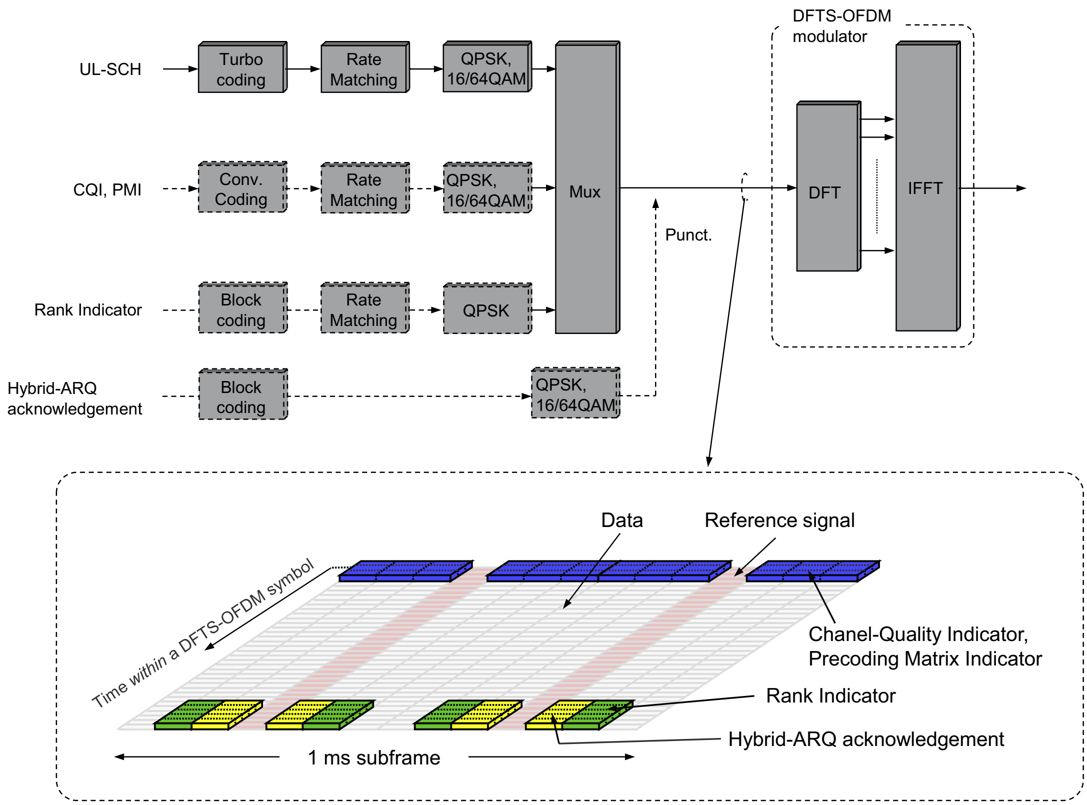

FIGURE 7.32

Multiplexing of control and data onto PUSCH.

eNodeB will expect a hybrid-ARQ acknowledgment while the device will not transmit one. If the rate-matching pattern was to depend on whether an acknowledgment is transmitted or not, all the coded bits transmitted in the data part could be affected by a missed assignment, which is likely to cause the UL-SCH decoding to fail. To avoid this error, the hybrid-ARQ acknowledgments are therefore punctured into the coded UL-SCH bit stream. Thus, the non-punctured bits are not affected by the presence/absence of hybrid-ARQ acknowledgments and the problem of a mismatch between the rate matching in the device and the eNodeB is avoided.

The contents of the CSI reports are described in Chapter 9; at this stage it suffices to note that a CSI report consists of *channel-quality indicator* (CQI), *precoding matrix indicator* (PMI), and *rank indicator* (RI). The CQI and PMI are time multiplexed with the coded data bits from PUSCH and transmitted using the same modulation as the data part. CSI reports are mainly useful for low-to-medium Doppler frequencies for which the radio channel is relatively constant, hence the need for special mapping is less pronounced. The RI, however, is mapped differently than the CQI and PMI; as illustrated in Figure 7.32, the RI is located near to the reference symbols using a similar mapping as the hybrid-ARQ acknowledgments. The more robust mapping of the RI is motivated by the fact that the RI is required in order to correctly interpret the CQI/PMI. The CQI/PMI, on the other hand, is simply mapped across the full subframe duration. Modulation-wise, the RI uses QPSK.

For uplink spatial multiplexing, in which case two transport blocks are transmitted simultaneously on the PUSCH, the CQI and PMI are multiplexed with the coded transport block using the highest modulation-and-coding scheme (MCS), followed by applying the previously described multiplexing scheme per layer (Figure 7.33). The intention behind this approach is to transmit the CQI and PMI on the (one or two) layers with the best quality.[27]

The hybrid-ARQ acknowledgments and the rank indicator are replicated across all transmission layers and multiplexed with the coded data in each layer in the same way as the single layer case described in the preceding paragraphs. The bits may, though, have been scrambled differently on the different layers. In essence, as the same information is transmitted on multiple layers with different scrambling, this provides diversity.

The basis for CSI reporting on the PUSCH is *aperiodic reports*, where the eNodeB requests a report from the device by setting the CSI request bit in the scheduling grant, as mentioned in Chapter 6. UL-SCH rate matching takes the presence of the CSI reports into account; by using a higher code rate a suitable number of resource elements is made available for transmission of the CSI report. Since the reports are explicitly requested by the eNodeB, their presence is known and the appropriate rate de-matching can be done at the receiver. If one of the configured transmission instances for a periodic report coincides with the

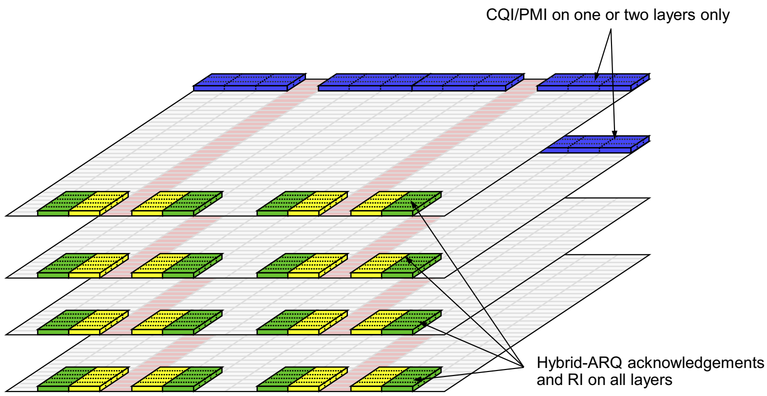

FIGURE 7.33

Multiplexing of CQI/PMI, RI and hybrid-ARQ acknowledgments in case of uplink spatial multiplexing.

[27]Assuming the MCS follows the channel quality, this holds for one, two, and four layers but not necessarily for three layers.

device being scheduled on the PUSCH, the periodic report is "rerouted" and transmitted on the PUSCH resources. Also, in this case there is no risk of mismatch in rate matching; the transmission instants for periodic reports are configured by robust RRC signaling and the eNodeB knows in which subframes such reports will be transmitted.

The channel coding of the CSI reports depends on the report size. For the smaller sizes such as a periodic report that otherwise would have been transmitted on the PUCCH, the same block coding as used for the PUCCH reports is used. For the larger reports, a tail-biting convolutional code is used for CQI/PMI, whereas the RI uses a (3, 2) block code for a single component carrier.

Unlike the data part, which relies on rate adaptation to handle different radio conditions, this cannot be used for the L1/L2 control-signaling part. Power control could, in principle, be used as an alternative, but this would imply rapid power variations in the time domain, which negatively impact the RF properties. Therefore, the transmission power is kept constant over the subframe and the amount of resource elements allocated to L1/L2 control signaling—that is, the code rate of the control signaling—is varied according to the scheduling decision for the data of the data part. High data rates are typically scheduled when the radio conditions are advantageous and hence a smaller amount of resource needs to be used by the L1/L2 control signaling compared to the case of poor radio conditions. To account for different hybrid-ARQ operating points, an offset between the code rate for the control-signaling part and the MCS used for the data part can be configured via higher-layer signaling.

7.5 UPLINK POWER CONTROL

Uplink power control for LTE is the set of algorithms and tools by which the transmit power for different uplink physical channels and signals are controlled to ensure that they, to the extent possible, are received with the appropriate power. This means that the transmission should be received with sufficient power to allow for proper demodulation of the corresponding information. At the same time, the transmit power should not be unnecessarily high as that would cause unnecessary interference to other transmissions in the same or other cells. The transmit power will thus depend on the channel properties, including the channel attenuation and the noise and interference level at the receiver side. Furthermore, in the case of UL-SCH transmission on PUSCH, if the received power is too low one can either increase the transmit power or reduce the data rate by use of rate control. Thus, for PUSCH transmission there is an intimate relation between power control and the link adaptation (rate control).

How to set the transmit power for random access is discussed in Chapter 11. Here we mainly discuss the power-control mechanism for the PUCCH and PUSCH physical channels. We also briefly discuss the power setting for SRS. Uplink DM-RS are always transmitted together and time-multiplexed with PUSCH or PUCCH. The DM-RS are then transmitted with the same power as the corresponding physical channel. This is also true in the case of uplink spatial multiplexing if the reference-signal power is defined as the total power of all DM-RS transmitted by the device. Expressed differently, the power of a single DM-RS is equal to the corresponding *per-layer* PUSCH power.

Fundamentally, LTE uplink power control is a combination of an *open-loop* mechanism, implying that the device transmit power depends on estimates of the downlink path loss, and a *closed-loop* mechanism, implying that the network can, in addition, directly adjust the device transmit power by means of explicit *power-control commands* transmitted on the downlink. In practice, these power-control commands are determined based on prior network measurements of the received uplink power, thus the term "*closed loop*."

7.5.1 UPLINK POWER CONTROL: SOME BASIC RULES

Before going into the details of the power-control algorithms for PUSCH and PUCCH, some basic rules for the power assignment to different physical channels will be discussed. These rules mainly deal with the presence of different transmit-power limitations and how these limitations impact the transmit-power setting for different physical channels. This is especially of interest in the case of the simultaneous transmission of multiple physical channels from the same device, a situation that may occur for LTE releases 10 and beyond:

- Release 10 introduced the possibility for carrier aggregation, implying that multiple PUSCH may be transmitted in parallel on different component carriers.
- Release 10 also introduced the possibility for simultaneous PUSCH/PUCCH transmission on the same or different component carriers.

In principle, each physical channel is separately and independently power controlled. However, in the case of multiple physical channels to be transmitted in parallel from the same device, the total power to be transmitted for all physical channels may, in some cases, exceed the maximum device output power P_{TMAX} corresponding to the device power class. As will be seen below, the basic strategy is then to first ensure that transmission of any L1/L2 control signaling is assigned the power assumed to be needed for reliable transmission. The remaining available power is then assigned to the remaining physical channels.

For each uplink component carrier configured for a device there is also an associated and explicitly configured *maximum per-carrier transmit power* $P_{\mathrm{CMAX},c}$, which may be different for different component carriers (indicated by the index c). Furthermore, although it obviously does not make sense for $P_{\mathrm{CMAX},c}$ to exceed the maximum device output power P_{TMAX}, the sum of $P_{\mathrm{CMAX},c}$ for all configured component carriers may very well, and typically will, exceed P_{TMAX}. The reason is that, in many cases, the device will not be scheduled for uplink transmission on all its configured component carriers and the device should also in that case be able to transmit with its maximum output power.

As will be seen in the next sections, the power control of each physical channel explicitly ensures that the total transmit power for a given component carrier does not exceed $P_{\mathrm{CMAX},c}$ for that carrier. However, the separate power-control algorithms do not ensure that the total transmit power for all component carriers to be transmitted by the device does not exceed the maximum device output power P_{TMAX}. Rather, this is ensured by a subsequent *power scaling* applied to the physical channels to be transmitted. This power scaling is carried out in such a way that any L1/L2 control signaling has higher priority, compared to data (UL-SCH) transmission.

If PUCCH is to be transmitted in the subframe, it is first assigned the power determined by its corresponding power-control algorithm, before any power is assigned to any PUSCH to be transmitted in parallel PUCCH. This ensures that L1/L2 control signaling on PUCCH is assigned the power assumed to be needed for reliable transmission before any power is assigned for data transmission.

If PUCCH is not transmitted in the subframe but L1/L2 control signaling is multiplexed on to PUSCH, the PUSCH carrying the L1/L2 control signaling is first assigned the power determined by its corresponding power-control algorithm, before any power is assigned to any other PUSCH to be transmitted in parallel. Once again, this ensures that L1/L2 control signaling is assigned the power assumed to be needed before any power is assigned for other PUSCH transmissions only carrying UL-SCH. Note that, in the case of transmission of multiple PUSCH in parallel (carrier aggregation), at most one PUSCH may include L1/L2 control signaling. Also, there cannot be PUCCH transmission and L1/L2 control signaling multiplexed on to PUSCH in the same subframe. Thus, there will never be any conflict between the said rules.

If the remaining available transmit power is not sufficient to fulfill the power requirements of any remaining PUSCH to be transmitted, the powers of these remaining physical channels, which only carry UL-SCH, are scaled so that the total power for all physical channels to be transmitted does not exceed the maximum device output power.

Overall, the PUSCH power scaling, including the priority for PUSCH with L1/L2 control signaling, can thus be expressed as:

$$\sum_c w_c \cdot P_{PUSCH,c} \leq P_{TMAX} - P_{PUCCH} \tag{7.2}$$

where $P_{PUSCH,c}$ is the transmit power for PUSCH on carrier c as determined by the power-control algorithm (before power scaling but including the per-carrier limitation $P_{CMAX,c}$), P_{PUCCH} is the transmit power for PUCCH (which is zero if there is no PUCCH transmission in the subframe), and w_c is the power-scaling factor for PUSCH on carrier c ($w_c \leq 1$). For any PUSCH carrying L1/L2 control signaling the scaling factor w_c should be set to 1. For the remaining PUSCH, some scaling factors may be set to zero by decision of the device, in practice implying that the PUSCH, as well as the corresponding UL-SCH mapped to the PUSCH, are not transmitted. For the remaining PUSCH the scaling factors w_c are set to the same value less than or equal to 1 to ensure that the above inequality is fulfilled. Thus, all PUSCH that are actually transmitted are power scaled by the same factor.

After this overview of some general rules for the power setting of different devices, especially for the case of multiple physical channels transmitted in parallel from the same device, the power control carried out separately for each physical channel will be described in more detail.

7.5.2 POWER CONTROL FOR PUCCH

For PUCCH, the appropriate received power is simply the power needed to achieve a desired—that is, a sufficiently low—error rate in the decoding of the L1/L2 control information transmitted on the PUCCH. However, it is then important to bear the following in mind:

- In general, decoding performance is not determined by the *received signal strength* but rather by the *received signal-to-interference-plus-noise ratio* (SINR). What is an appropriate received power thus depends on the interference level at the receiver side, an interference level that may differ between different deployments and which may also vary in time as, for example, the load of the network varies.
- As previously described, there are different PUCCH formats which are used to carry different types of uplink L1/L2 control information (hybrid-ARQ acknowledgments, scheduling requests, CSI, or combinations thereof). The different PUCCH formats thus carry different numbers of information bits per subframe and the information they carry may also have different error-rate requirements. The required received SINR may therefore differ between the different PUCCH formats, something that needs to be taken into account when setting the PUCCH transmit power in a given subframe.

Overall, power control for PUCCH can be described by the following expression:

$$P_{\mathrm{PUCCH}} = \min\{P_{\mathrm{CMAX},c}, P_{0,\mathrm{PUCCH}} + \mathrm{PL}_{\mathrm{DL}} + \Delta_{\mathrm{Format}} + \delta\} \tag{7.3}$$

where P_{PUCCH} is the PUCCH transmit power to use in a given subframe and $\mathrm{PL}_{\mathrm{DL}}$ is the downlink path loss as estimated by the device. The "min $\{P_{\mathrm{CMAX},c},\ldots\}$" term ensures that the PUCCH transmit power as determined by the power control will not exceed the per-carrier maximum power $P_{\mathrm{CMAX},c}$.

The parameter $P_{0,\mathrm{PUCCH}}$ in expression (7.3) is a cell-specific parameter that is broadcast as part of the cell system information. Considering only the part $P_{0,\mathrm{PUCCH}} + \mathrm{PL}_{\mathrm{DL}}$ in the PUCCH power-control expression and assuming that the (estimated) downlink path loss accurately reflects the true uplink path loss, it is obvious that $P_{0,\mathrm{PUCCH}}$ can be seen as the *desired* or *target* received power. As discussed earlier, the required received power will depend on the uplink noise/interference level. From this point of view, the value of $P_{0,\mathrm{PUCCH}}$ should take the interference level into account and thus vary in time as the interference level varies. However, in practice it is not feasible to have $P_{0,\mathrm{PUCCH}}$ varying with the instantaneous interference level. One simple reason is that the device does not read the system information continuously and thus the device would anyway not have access to a fully up-to-date $P_{0,\mathrm{PUCCH}}$ value. Another reason is that the uplink path-loss estimates derived from downlink measurements will anyway not be fully accurate, for example due to differences between the instantaneous downlink and uplink path loss and measurement inaccuracies.

Thus, in practice, $P_{0,\mathrm{PUCCH}}$ may reflect the average interference level, or perhaps only the relatively constant noise level. More rapid interference variations can then be taken care of by closed-loop power control, see below.

For the transmit power to reflect the typically different SINR requirements for different PUCCH formats, the PUCCH power-control expression includes the term Δ_{Format}, which adds a format-dependent power offset to the transmit power. The power offsets are defined such that a baseline PUCCH format, more exactly the format corresponding to the transmission of a single hybrid-ARQ acknowledgment (format 1 with BPSK modulation, as described in Section 7.4.1.1), has an offset equal to 0 dB, while the offsets for the remaining formats can be explicitly configured by the network. For example, PUCCH format 1 with QPSK modulation, carrying two simultaneous acknowledgments and used in the case of downlink spatial multiplexing, should have a power offset of roughly 3 dB, reflecting the fact that twice as much power is needed to communicate two acknowledgments instead of just a single acknowledgment.

Finally, it is possible for the network to directly adjust the PUCCH transmit power by providing the device with explicit power-control commands that adjust the term δ in the power-control expression (7.3). These power-control commands are *accumulative*—that is, each received power-control command increases or decreases the term δ by a certain amount. The power-control commands for PUCCH can be provided to the device by two different means:

- As mentioned in Section 6.4, a power-control command is included in each downlink scheduling assignment—that is, the device receives a power-control command every time it is explicitly scheduled on the downlink. One reason for uplink PUCCH transmissions is the transmission of hybrid-ARQ acknowledgments as a response to downlink DL-SCH transmissions. Such downlink transmissions are typically associated with downlink scheduling assignments on PDCCH and the corresponding power-control commands could thus be used to adjust the PUCCH transmit power prior to the transmission of the hybrid-ARQ acknowledgments.
- Power-control commands can also be provided on a special PDCCH that simultaneously provides power-control commands to multiple devices (PDCCH using DCI format 3/3A; see Section 6.4.8). In practice, such power-control commands are then typically transmitted on a regular basis and can be used to adjust the PUCCH transmit power, for example prior to (periodic) uplink CSI reports. They can also be used in the case of semi-persistent scheduling (see Chapter 9), in which case there may be uplink transmission of both PUSCH (UL-SCH) and PUCCH (L1/L2 control) without any explicit scheduling assignments/grants.

The power-control command carried within the uplink scheduling grant consists of two bits, corresponding to the four different update steps -1, 0, $+1$, or $+3$ dB. The same is true for the power-control command carried on the special PDCCH assigned for power control when this is configured to DCI format 3A. On the other hand, when the PDCCH is configured to use DCI format 3, each power-control command consists of a single bit, corresponding to the update steps −1 and +1 dB. In the latter case, twice as many devices can be power controlled by a single PDCCH. One reason for including the possibility for 0 dB (no change of power) as one power-control step is that a power-control command is included in *every* downlink scheduling assignment and it is desirable not to have to update the PUCCH transmit power for each assignment.

7.5.3 POWER CONTROL FOR PUSCH

Power control for PUSCH transmission can be described by the following expression:

$$P_{\mathrm{PUSCH},c} = \min\left\{P_{\mathrm{CMAX},c} - P_{\mathrm{PUCCH}}, P_{0,PUSCH} + \alpha\cdot \mathrm{PL_{DL}} + 10\cdot\log_{10}(M) + \Delta_{\mathrm{MCS}} + \delta\right\} \tag{7.4}$$

where M indicates the instantaneous PUSCH bandwidth measured in number of resource blocks and the term Δ_{MCS} is similar to the term Δ_{Format} in the expression for PUCCH power control—that is, it reflects the fact that different SINR is required for different modulation schemes and coding rates used for the PUSCH transmission.

Equation (7.4) is similar to the power-control expression for PUCCH transmission, with some key differences:

- The use of "$P_{\mathrm{CMAX},c} - P_{\mathrm{PUCCH}}$" reflects the fact that the transmit power available for PUSCH on a carrier is the maximum allowed per-carrier transmit power *after power has been assigned to any PUCCH transmission* on that carrier. This ensures priority of L1/L2 signaling on PUCCH over data transmission on PUSCH in the power assignment, as described in Section 7.5.1.
- The term $10\cdot\log_{10}(M)$ reflects the fact that what is fundamentally controlled by the parameter $P_{0,\mathrm{PUSCH}}$ is the power *per resource block*. For a larger resource assignment, a correspondingly higher received power and thus a correspondingly higher transmit power is needed.[28]
- The parameter α, which can take a value smaller than or equal to 1, allows for so-called *partial path-loss compensation*, as described in the following.

In general, the parameters $P_{0,\mathrm{PUSCH}}$, α, and Δ_{MCS} can be different for the different component carriers configured for a device.

In the case of PUSCH transmission, the explicit power-control commands controlling the term δ are included in the uplink scheduling grants, rather than in the downlink scheduling assignments. This makes sense as PUSCH transmissions are preceded by an uplink scheduling grant except for the case of semi-persistent scheduling. Similar to the power-control commands for PUCCH in the downlink scheduling assignment, the power-control

[28]One could also have included a corresponding term in the equation for PUCCH power control. However, as the PUCCH bandwidth always corresponds to one resource block, the term would always equal zero.

commands for PUSCH are multi-level. Furthermore, also in the same way as for PUCCH power control, explicit power-control commands for PUSCH can be provided on the special PDCCH that simultaneously provides power-control commands to multiple devices. These power-control commands can, for example, be used for the case of PUSCH transmission using semi-persistent scheduling.

Assuming α equal to 1, also referred to as *full path-loss compensation*, the PUSCH power-control expression becomes very similar to the corresponding expression for PUCCH. Thus, the network can select a MCS and the power-control mechanism, including the term Δ_{MCS}, will ensure that the received SINR will match the SINR required for that MCS, *assuming that the device transmit power does not reach its maximum value*.

In the case of PUSCH transmission, it is also possible to "turn off" the Δ_{MCS} function by setting all Δ_{MCS} values to zero. In that case, the PUSCH received power will be matched to a certain MCS given by the selected value of $P_{0,PUSCH}$.

With the parameter α less than 1, the PUSCH power control operates with so-called *partial path-loss compensation*—that is, an increased path loss is not fully compensated for by a corresponding increase in the uplink transmit power. In that case, the received power, and thus the received SINR per resource block, will vary with the path loss and, consequently, the scheduled MCS should vary accordingly. Clearly, in the case of fractional path-loss compensation, the Δ_{MCS} function should be disabled. Otherwise, the device transmit power would be further reduced when the MCS is reduced to match the partial path-loss compensation.

Figure 7.34 illustrates the differences between full path-loss compensation ($\alpha = 1$) and partial path-loss compensation ($\alpha < 1$). As can be seen, with partial path-loss compensation, the device transmit power increases more slowly than the increase in path loss (left in the figure) and, consequently, the received power, and thus also the received SINR, is reduced as

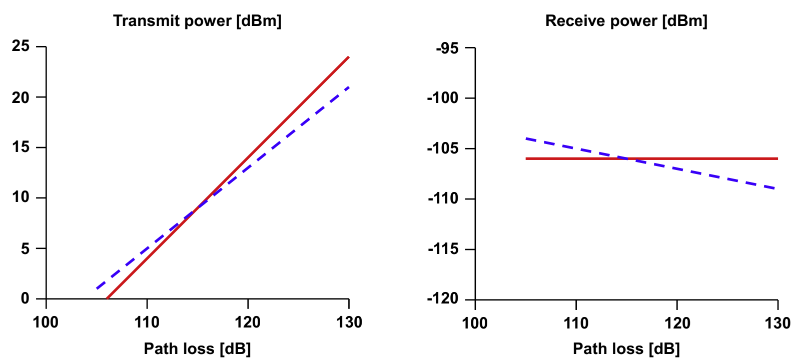

FIGURE 7.34

Full versus partial path-loss compensation. Solid curve: full compensation ($\alpha = 1$). Dashed curve: partial compensation ($\alpha = 0.8$).

the path loss increases (right in the figure). To compensate for this, the MCS—that is, the PUSCH data rate—should be reduced as the path loss increases.

The potential benefit of partial path-loss compensation is a relatively lower transmit power for devices closer to the cell border, implying less interference to other cells. At the same time, this also leads to a reduced data rate for these devices. It should also be noted that a similar effect can be achieved with full path-loss compensation by having the scheduled MCS depend on the estimated downlink path loss, which can be derived from the power headroom report, and rely on Δ_{MCS} to reduce the relative device transmit power for devices with higher path loss. However, an even better approach would then be to not only base the MCS selection on the path loss to the current cell, but also on the path loss to the neighboring interfered cells.

7.5.4 POWER CONTROL FOR SRS

The SRS transmit power basically follows that of the PUSCH, compensating for the exact bandwidth of the SRS transmission and with an additional power offset. Thus, the power control for SRS transmission can be described according to the equation:

$$P_{SRS} = \min\{P_{CMAX,c}, P_{0,PUSCH} + \alpha \cdot PL_{DL} + 10 \cdot \log_{10}(M_{SRS}) + \delta + P_{SRS}\} \quad (7.5)$$

where the parameters $P_{0,PUSCH}$, α, and δ are the same as for PUSCH power control, as discussed in Section 7.5.3. Furthermore, M_{SRS} is the bandwidth, expressed as number of resource blocks, of the SRS transmission and P_{SRS} is a configurable offset.

7.6 UPLINK TIMING ALIGNMENT

The LTE uplink allows for uplink intra-cell orthogonality, implying that uplink transmissions received from different devices within a cell do not cause interference to each other. A requirement for this *uplink orthogonality* to hold is that the signals transmitted from different devices within the same subframe but within different frequency resources (different resource blocks) arrive approximately time aligned at the base station. More specifically, any timing misalignment between received signals should fall within the cyclic prefix. To ensure such receiver-side time alignment, LTE includes a mechanism for *transmit-timing advance*.

In essence, timing advance is a negative offset, at the device, between the start of a received downlink subframe and a transmitted uplink subframe. By controlling the offset appropriately for each device, the network can control the timing of the signals received at the base station from the devices. Devices far from the base station encounter a larger propagation delay and therefore need to start their uplink transmissions somewhat in advance, compared to devices closer to the base station, as illustrated in Figure 7.35. In this specific example, the first device is located close to the base station and experiences a small propagation delay, $T_{P,1}$. Thus, for this device, a small value of the timing advance offset $T_{A,1}$ is sufficient to compensate for the propagation delay and to ensure the correct timing at the base

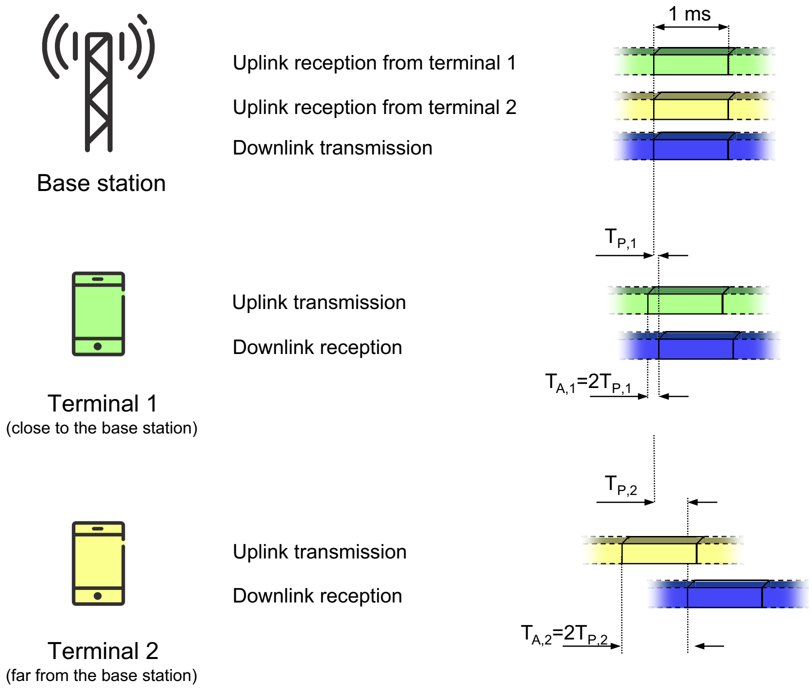

FIGURE 7.35

Uplink timing advance.

station. However, a larger value of the timing advance is required for the second device, which is located at a larger distance from the base station and thus experiences a larger propagation delay.

The timing-advance value for each device is determined by the network based on measurements on the respective uplink transmissions. Hence, as long as a device carries out uplink data transmission, this can be used by the receiving base station to estimate the uplink receive timing and thus be a source for the timing-advance commands. SRS can be used as a regular signal to measure upon, but in principle the base station can use any signal transmitted from the devices.

Based on the uplink measurements, the network determines the required timing correction for each device. If the timing of a specific device needs correction, the network issues a timing-advance command for this specific device, instructing it to retard or advance its timing relative to the current uplink timing. The user-specific timing-advance command is transmitted as a MAC control element (see Chapter 4 for a description of MAC control elements) on the DL-SCH. The maximum value possible for timing advance is 0.67 ms, corresponding to a device-to-base-station distance of slightly more than 100 km. This is also the value assumed when determining the processing time for decoding, as discussed in Section 8.1.

Typically, timing-advance commands to a device are transmitted relatively infrequently—for example, one or a few times per second.

If the device has not received a timing-advance command during a (configurable) period, the device assumes it has lost the uplink synchronization. In this case, the device must re-establish uplink timing using the random-access procedure prior to any PUSCH or PUCCH transmission in the uplink.