A concept for evaluating sub-THz communication for future 6G

Stefan Parkvall, Robert Baldemair, Erik Eriksson, Mattias Frenne,
Ajit Nimbalker, Claes Tidestav
Ericsson

Daewon Lee, Peng Lu, Gary Xiong, Yingyang Li, Sungho Moon, Jeff Zhuang
Intel Corporation

1 What is 6G?

Over the last three decades, cellular connectivity has undergone a remarkable evolution, going from voice-only systems to mobile data and IoT connectivity, serving not only smartphones but also facilitating new usage scenarios such as factory production, logistics, and smart cities with 5G NR from 3GPP representing state-of-the-art cellular standards. A similar evolution has taken place in other areas. For example, compute platforms have evolved from monolithic software run on mainframe computers via personal workstations into today's micro-service-based cloud computing. The evolution of compute power has enabled big data analytics and machine learning. Clearly, the evolution in these areas will continue.

In parallel, the research community has started to discuss 6G and wireless communication in 2030 and beyond. At that point, society has been shaped by 5G for 10 years, and new needs and services have appeared. It is expected that 6G systems will address multiple overlapping technical areas. Apart from limitless connectivity, that is, meeting all the connectivity needs of future applications, 6G will also further stress aspects such as trustworthy systems (protection and resiliency against intentional and unintentional disturbances), cognitive networks (where networks autonomously can perceive what is happening, take the appropriate actions, and explain to the operator why a certain action was taken), and network-compute fabric (the convergence of communication and compute into a single innovation platform).

A wide range of technologies, some which are yet to be defined, will together realize the 6G platform. Sub-THz communication, that is, communication in the 100–300 GHz frequency range, is one such technology. However, it is important to understand that sub-THz communication is only one of many technology components considered for 6G and that communication in lower frequency bands will continue to be important also in the 6G era.

Sub-THz communication can unleash vast amounts of spectrum for communication. However, due to the limited range of such sub-THz transmissions, operation in such frequency bands would be limited to very specific scenarios where extreme data rates and/or low latency in local areas are required. There are several potential use cases that can benefit from sub-THz communications such as untethered metaverse and over-the-air remote direct memory access (RDMA) connecting computers. How to deploy a system at sub-THz, including beamforming and mobility, is an interesting research challenge. On the hardware side, RF components at these frequency ranges are not yet mature for low-cost mass production, and hence, require further research.

To better understand the properties of sub-THz communication, Ericsson and Intel have jointly developed a concept to study the feasibility and performance of communication at these high frequencies, focusing on two important aspects: extensive use of beamforming and efficient processing to achieve very high data rates and/or super low latency. Implementations based on this concept can be used to evaluate sub-THz communication and various related technology components.

The remainder of this paper will describe the concept in detail and motivate the reasoning behind different design choices.

2 Transmission scheme and time-frequency structure

Generating RF energy at frequencies around 100 GHz and above is challenging and it is important to utilize a waveform requiring a low back-off in the power amplifier. DFTS-OFDM is a suitable waveform providing a low peak-to-average power ratio, and hence, allowing for a good power efficiency for the power amplifier under different process nodes.

One important aspect of DFTS-OFDM is the selection of the numerology, in particular the subcarrier spacing and the cyclic prefix length. To handle the impact from phase noise, which becomes more problematic at very high carrier frequencies, a large subcarrier spacing is beneficial. It also allows implementations covering a large bandwidth with a modest FFT size which affects the FFT processing latency relative to the much-shortened symbol period (521 ns). With this in mind, a subcarrier spacing of 1.92 MHz has been chosen, a value that is 128 times larger than the basic subcarrier spacing of 15 kHz in NR. The cyclic prefix is set to approximately 79 ns (with the exception for some symbols as discussed in Section 2.1, where it will be 120 ns). Although it may sound small, it is in fact sufficient to handle the delay spread in the targeted deployment scenarios with modest cell sizes and highly directive beamforming, as well as any timing misalignment between transmission points when the devices is simultaneously connected to multiple transmission points.

Table 2-1: DFTS-OFDM numerology ($T_{\bar{c}} = 1/(4096 \cdot 1.920 \cdot 10^6)$ s denotes the basic sampling time).

Subcarrier spacing	Useful symbol time, T_u	Cyclic prefix, T_{CP}
1920 kHz	$4096T_{\bar{c}} \approx 521$ ns	$624T_{\bar{c}} \approx 79$ ns[1]

2.1 Time-domain structure

Low latency is important aspect for 6G communication. To avoid waiting for the start of a slot when transmitting latency-critical information, transmissions are not bound to slot boundaries but may start and end at (almost) any DFTS-OFDM symbol. It is somewhat similar to the so-called ‘mini-slots’ in NR but takes it one step further by discarding the notion of a slot and relying on the DFTS-OFDM symbol as the basic time unit for transmission. Since the transmission can start and end at any DFT-OFDM symbol with ~600 ns in time duration, a minimal one-way latency of ~600 ns plus the processing time is possible. Reducing the processing time is essential and, together with the ultra-high data rates targeted, requires careful consideration.

Despite the symbol-oriented design, there is a need to define a longer time structure to be used for control signals, for example, when defining transmission of system information and

[1] For some symbols $T_{CP} \approx$120 ns.

random-access occasions. For this purpose, a 1 ms subframe is defined. A subframe is in turn divided into eight TDD periods, each consisting of 208 symbols (see Figure 2-1).

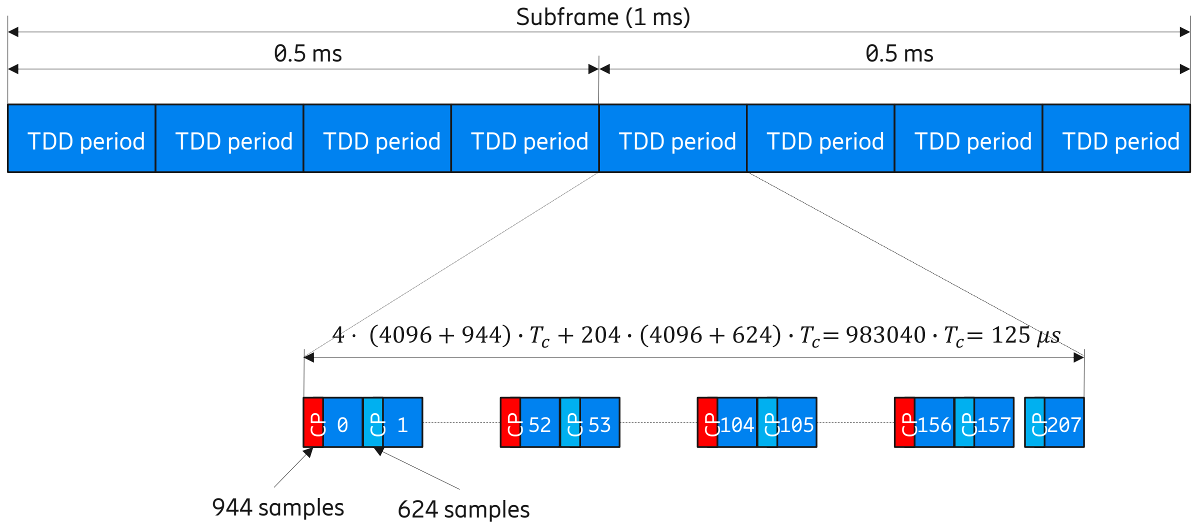

Figure 2-1: Time-domain structure.

To obtain an integer number of symbols per subframe, every 52nd symbol has a slightly longer cyclic prefix as seen in Figure 2-1. The reason for taking this approach, instead of a direct scaling of the NR structure which would result in a significantly longer cyclic prefix every 896th symbol, is to get a more uniform distribution of the samples used for the cyclic prefix.

2.2 Frequency-domain structure

In the frequency domain, the transmitted signal is organized in resource blocks of 12 subcarriers each. Up to 270 resource blocks, corresponding to a carrier bandwidth of almost 7 GHz, are supported. To simplify implementing DFTS-OFDM as a DFT-precoder followed by an OFDM transmitter, the size of the DFT-precoder is limited to multiples of 2, 3, and 5 as in NR.

2.3 Duplex scheme

Dynamic TDD is the basic duplexing scheme employed, in principle allowing for the uplink-downlink ratio to be dynamically changed. However, for simplicity, implementations may be based on fixed uplink-downlink ratios. Three examples hereof, representing downlink-heavy, balanced, and uplink-heavy traffic patterns, are given in Table 2-2 where each TDD period in Figure 2-1 is split into a downlink part and an uplink part as shown. Essentially the hybrid ARQ (HARQ) round-trip time will be in the units of TDD periods (that is, 125µs) which is adequate for initial evaluation purposes. Since the synchronization signal block (SSB) and physical random-access channel (PRACH) are located at the end of a TDD period as described in section 8, there are different number of symbols available for uplink data transmission depending on whether an SSB or a PRACH is part of that TDD period as illustrated in Figure 2-2.

Table 2-2: Example of uplink-downlink ratios.

Ratio	Downlink L_{DL}	Guard L_{GI}	Uplink		
			$L_{UL,1}$ (no PRACH, no SSB)	$L_{UL,2}$ (SSB)	$L_{UL,3}$ (PRACH)
3:1	140	12	56	42	32
1:1	90	12	106	92	82
1:3	48	12	148	134	124

Figure 2-2: Splitting a TDD period into downlink and uplink parts.

3 Transport-channel processing

This section provides a more detailed description of the downlink and uplink physical-layer functionality such as coding, modulation, multi-antenna precoding, resource-block mapping, and reference signal structure.

3.1 Overview

The overall transport-channel processing of a downlink or uplink transport block is illustrated in Figure 3- and is, on a high level, similar to NR. However, to allow for a parallelization-friendly implementation, which is necessary given the very high data rates targeted, several changes have been made compared to NR. In particular, channel coding and the mapping of coded bits to an integer number of DFTS-OFDM symbols is different from NR.

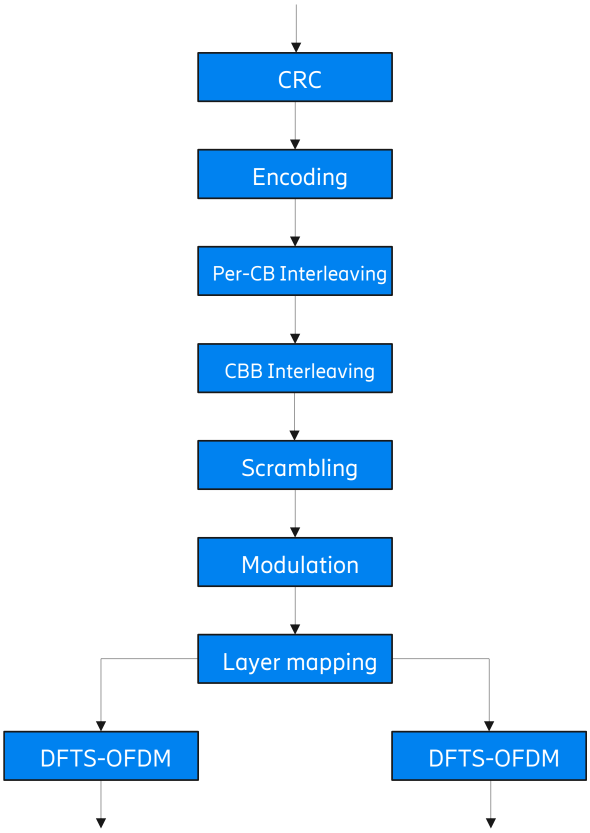

Figure 3-1: General transport-channel processing.

3.2 Channel coding and interleaving

A transport block is split into one or more code blocks (CBs) and the code blocks from the same transport block are known as a code-block bundle (CBB). To allow for efficient and parallel implementation, each CBB in a PDSCH/PUSCH, except for the lowest data rates, spans one OFDM symbol, that is, the CBB boundaries and DFTS-OFDM symbol boundaries are aligned. This allows the receiver to process each DFTS-OFDM symbol independently and in parallel, which is beneficial at the very high data rates targeted. For the lowest data rates, where a single code-block CBB per DFTS-OFDM symbol would result in a too high data rate assuming a reasonable minimum code block size, a CBB may span multiple symbols.

A per-code-block CRC is attached to each code block in a CBB. For the last CBB in a PDSCH/PUSCH transmission, a special CRC mask is used to enable early termination of the transmission, see further Section 5.1.

Channel coding for each CB uses the same LDPC base graph and circular buffer rate-matching as in NR. For the downlink, limited buffer rate-matching is used, while for the uplink, both limited and full buffer rate-matching are supported, with up to four redundancy versions.

After rate-matching, a per-code-block interleaver like in NR is applied to facilitate mapping of systematic bits to more reliable positions in QAM constellations for redundancy version 0. For CBBs containing multiple code blocks, per-CBB interleaving based on a rectangular interleaver is also used for modulation orders higher than π/2-BPSK to further improve performance.

3.3 Scrambling and Modulation

Scrambling is applied to the coded bits using a bit-level scrambling sequence. The scrambling generator is reinitialized for every CBB based on the radio-network temporary identifier (RNTI) and, in case of multiple transmission points (TRPs) as described in section 6.3, the leg identity. The scrambled bits are subsequently modulated using π/2-BPSK, QPSK, 16-QAM, 64-QAM, or 256-QAM.

3.4 Layer mapping

A fixed rank of two layers is used for transmission of PDSCH, PUSCH and PDCCH. Each of the two layers is directly mapped to one of the two polarizations, by using the 2x2 identity matrix as the MIMO precoder. This avoids the drawbacks of mixing signals of two layers through the same power amplifier that otherwise would increase the peak to average power ratio (PAPR). A fixed rank 2 and a fixed MIMO precoding matrix further simplify channel-state information (CSI) reporting since there is no need for reporting of precoding matrices or rank indicators, only reporting of CQI is necessary.

The use of a fixed rank 2 is further motivated by the use of beamforming with a narrow beamwidth, where the resulting propagation condition between network and mobile device is effectively line of sight (LOS). If dual-polarized antenna elements are used, the probability of rank 2 being preferred when using such a narrow beam is very high. Measurement data at 28 GHz, where it was seen that rank 1 transmissions are very rare, support this approach.

The mapping of modulated symbols is done first across layers, that is, symbols are mapped alternating to the two layers, then in time. Note that the π/2-BPSK modulation is applied per layer to obtain the low PAPR property of consecutive π/2-BPSK modulation symbols within each layer. This is achieved by using a rotated BPSK constellation for the modulation and performing the time-varying π/2 phase rotation per layer and symbol as shown in Figure 3-2.

Figure 3-2: π/2-BPSK modulation using the fixed two-layer transmission.

3.5 Reference signals

To achieve a low PAPR property of the reference signals, Zadoff-Chu-based reference signals are used, based on the reference-signal structure used for the NR DFTS-OFDM uplink. A reference signal is mapped to every fourth subcarrier, resulting in a reference signal comb-structure. Within each such comb, two cyclic shifts are defined. Consequently, an OFDM symbol used for demodulation reference signals (DM-RS) can handle up to eight different reference signal ports. Due to the much narrow beamwidth resulted from the use of a much larger array to deliver the required EIRP, often there can be just a single user served in each narrow beam with two layers. User capacity can be increased with spatial domain multiplexing using multiple beams. Eight antenna ports will enable four rank-2 users to be served simultaneously.

In the time domain, the DM-RS is “front-loaded”, that is, a scheduled PDSCH or PUSCH transmission always starts with a DM-RS symbol. For scenario with fast channel variation, additional DM-RS symbols can be configured every 8th, 16th, or 32nd OFDM symbol to allow for channel interpolation and improved channel-estimation performance. For random access, Msg-3 PUSCH uses a DM-RS every 8th symbol, in addition to the front-loaded symbol, to provide robust performance in unknown channel conditions.

In the frequency domain, the DM-RS is contained within the scheduled bandwidth of PDSCH and PUSCH, that is, no DM-RS is transmitted outside the scheduled bandwidth. Moreover, since a fixed rank-two transmission is used (see Section 3.4), two ports are always selected, and there are two possibilities to assigning DM-RS by using the same or different DM-RS combs. In any case, the different cyclic shifts are used to provide orthogonality between ports while keeping the PAPR at its minimum.

To compensate for phase noise, a π/2-BPSK modulated phase-tracking reference signal (PT-RS) is used for the PDSCH and PUSCH. Two, four, or eight groups of PT-RS, each group consisting of a length-4 Hadamard code, are inserted prior to the DFTS-precoding, where the number of groups depends on the scheduled bandwidth. The PT-RS is present in both layers, and the Hadamard code ensures that the PT-RS are orthogonal between different layers.

Figure 3-3 PT-RS design where group of four PT-RS samples (marked with the X) are mapped to the center of each of four intervals of data symbols, (also two or eight intervals is used, depending in the scheduled bandwidth).

For measurements supporting beam management, a CSI-RS is also defined with a similar design as the DM-RS. Eight CSI-RS ports are defined within an OFDM symbol, and a CSI-RS resource always uses two ports, one per polarization. The CSI-RS is aperiodically triggered and follows the bandwidth of the simultaneously scheduled PDSCH. A gap between the PDSCH and CSI-RS is used to accommodate different timing requirements for the PDSCH and CSI-RS.

4 Physical-layer control signaling

To support the transmission of the PDSCH and PUSCH there is a need for associated control signaling, often referred to as L1/L2 control signaling. In this section, the transmission of downlink control information (DCI) and uplink control information (UCI) is described.

4.1 Downlink

Downlink control information is used for scheduling assignments and scheduling grants. At a high level, the general principle is similar to NR with blind decoding at the UE, but several details are different.

After 24-bit CRC attachment using the same CRC polynomial and initialization procedure as in NR, the bits are encoded, see Figure 4-1. Encoding and rate-matching follows the NR downlink polar coding with maximum code size of 512 with two exceptions: Firstly, the distributed CRC interleaver is omitted since the early termination benefits may be quite limited in practice. Secondly, the maximum DCI payload size is increased to 256 bits including CRC bits (compared to the NR max DCI payload size of 164 bits).

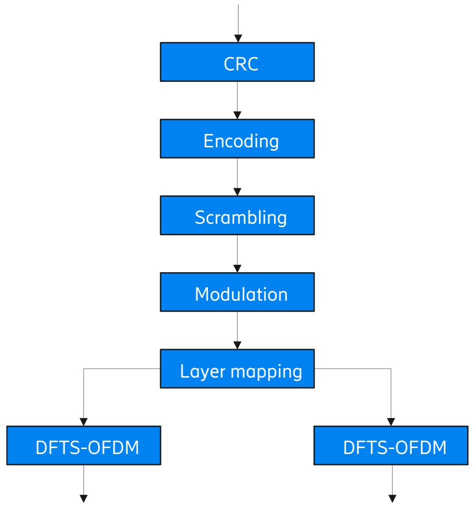

Figure 4-1: PDCCH processing chain.

The channel coding step is followed by scrambling, modulation, and layer mapping. The PDCCH is always transmitted with rank 2 for the same reason as the PDSCH and uses π/2-BPSK modulation to minimize the PAPR.

Each PDCCH is transmitted in a control resource set (CORESET) consisting of two or more symbols. The first symbol is used for the PDCCH DM-RS, while the following symbols carry the control information. The CORESETs are centered in the carrier bandwidth but can have different bandwidths.

Two different CORESETs are defined, CORESET0 and CORESET1, each with an associated search space used for blind decoding by the UE. Each search space has a single blind decode for a single DCI size.

CORESET0 consists of three symbols (one DM-RS symbol and two payload symbols), spanning 30 RBs. CORESET0 is associated with one search space placed at the end of the downlink portion of each TDD period. Two DCI formats are monitored in this search space, format 0 used for scheduling, and format 2 used for initial access.

CORESET1 consists of two symbols (one DM-RS symbol, one payload symbol). Associated with CORESET1 is a search space starting in symbol 0 and symbol 2 of each TDD period. CORESET1 is associated with one search space. In this search space a single DCI format, format 1, is monitored.

DCI formats 0 and 1 are both for scheduling of data and control and contain resource allocation for both uplink and downlink in the same DCI as a joint scheduling grant. A single DCI for both downlink and uplink is to allow simplified uplink control signaling design as to be described in the next section. Of course, the reduced search space also helps UE to improve PDCCH decoding latency and thus buffer less PDSCH data while waiting for PDCCH to be decoded. Format 0 is significantly smaller and is also sent over the double-symbol CORESET0. This format has limited resource-allocation flexibility both in time and frequency domain and is mainly targeting fallback operation when CSI is unknown. Format 1, on the other hand, offers a very large flexibility in both frequency and time domain as well as much more detailed HARQ operation compared to format 0.

Format 2 is used for initial access where the DCI carries the random-access response with a grant for uplink resources as well as a TC-RNTI.

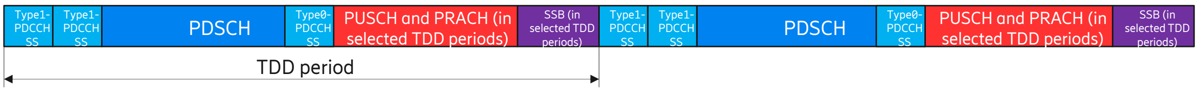

Figure 4-2: CORESETs, search spaces, and TDD periods.

4.2 Uplink

Uplink Control Information (UCI) is composed of HARQ-ACK feedback, aperiodic CQI reports for link adaptation, and reference-signal received power (RSRP) measurements based on SSB for beam management.

UCI is always transmitted on the PUSCH, either together with other data or alone, and no PUCCH is defined. Using the PUSCH for UCI in all scenarios instead of defining a PUCCH avoids complicated prioritization rules which would otherwise be needed and simplifies implementation.

The UCI is encoded using a simplified version of the NR channel coding, that is, repetition coding, simplex coding, Reed-Muller-based linear block coding, or Polar coding. The encoded UCI information is transmitted at the beginning of a PUSCH directly following the DM-RS symbol, see Figure 4-3. The number of symbols used for UCI transmission is indicated in the UL part of the joint downlink-uplink scheduling grant and ranges from one to seven symbols. UCI is transmitted using rank-2 and layer mapping follows PDCCH. Encoded UCI bits on each layer are modulated using π/2-BPSK (if PUSCH uses π/2-BPSK) or QPSK and transmitted using DFTS-OFDM.

Figure 4-3: UCI is transmitted on PUSCH directly after DM-RS.

Hybrid ARQ retransmissions are used, see Section 5.2 for details, and there is thus a need for HARQ acknowledgments (HARQ-ACKs) to be fed back to the base station. HARQ-ACK can be fed back for multiple scheduling assignments and also on a fine granular HARQ-process level which can lead to large UCI sizes. Depending on the available resources for UCI transmission HARQ-ACK feedback is dynamically bundled to fit into the available resources.

The CSI report in the UCI consists of a single wideband CQI, see Section 3.4 for the motivation. The rank-2 CQI is estimated based on the PDSCH DMRS, and if there is no

PDSCH scheduled, the PDCCH DMRS is used. For link adaptation, a 6-bit modulation-and-coding scheme (MCS) table is used to support a wide range of spectral efficiencies.

For beam management, the UE performs measurements on a set of aperiodic CSI-RS. The UE reports a 7-bit RSRP value, and the corresponding index of the CSI-RS resource in the set.

5 Scheduling and HARQ

Transmissions are dynamically scheduled, similar to NR, and a hybrid-ARQ mechanism is used to correct occasional transmission errors. In this section, some details around scheduling and hybrid-ARQ will be discussed.

5.1 Resource allocation

Single-carrier operation leaves limited room for frequency-domain adaptation and multiplexing. Hence, the focus of the system is on time- and spatial-domain adaptation. In frequency domain the start-length allocation scheme from NR is reused but in groups of 15 RBs.

The time-domain allocation is very flexible using DCI format 1, but even if the system is described as dynamic TDD some restrictions are put on the duplex flexibility to limit signaling overhead. A downlink symbol can start at any point, from symbol 0 and last any number of symbols within a TDD period. Uplink can start from symbol 28 up to the end of the TDD period.

For DCI Format 0 a set of seven tabulated frequency-domain resource allocations ranging from 15 to 270 PRBs and centered in the carrier are supported. This is done in order to keep the size of DCI format 0 small.

In addition to being able to signal where a transmission ends it is possible to terminate a transmission earlier than allocated. This provides additional flexibility, for example when a decision to transmit latency-critical information must be taken without knowing the exact amount of data to transmit. In this case the scheduler can allocate a sufficiently large allocation and the transmitter signals the termination of the transmission by encoding the codeblocks in the last transport block with a scrambling of the CRC as mentioned in Section 3.2.

5.2 HARQ operation

A single DCI may schedule a large number of independent transport blocks. HARQ builds on detecting errors and providing more information to enable decoding. Since each transport block is independently encoded it makes sense to retransmit only the ones that fail. This, however, is a trade-off between signaling overhead and the resources needed for unnecessary retransmission.

HARQ operation is done in groups of four transport blocks defined as a HARQ-process bundle (HPB). To limit signaling a set of possible starting points can be signaled and the following HPBs are assumed to be in increasing order. To indicate if a given HPB is used for new data or retransmissions a bitmap is added to the DCI for the allocated HPBs. In order to support retransmission of some HPBs without transmitting any new data an explicit retransmission indicator is used, in contrast to the toggling new-data indicator (NDI) bit in

NR/LTE. The HPBs with retransmissions are mapped first, in time, to allow termination of a transmission after completing the retransmissions.

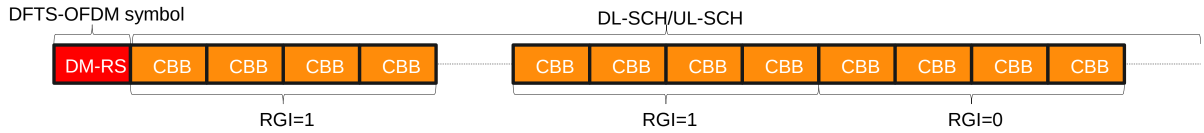

Figure 5-1 Retransmission group indicator (RGI) and the mapping to CBBs with retransmissions first. RGI set to 1 indicates that the HARQ process bundle is used for retransmissions of a group of CBBs.

To facilitate retransmissions also for the case when a DCI may have been lost, for example if no uplink was scheduled and HARQ feedback reception failed, an additional option for retransmission is provided by a single toggling NDI valid for all CBBs assigned by the DCI.

5.3 HARQ feedback

In order to support HARQ operation for downlink transmissions, feedback about the decoding results is needed from the UE. Feedback granularity has been aligned with the retransmission signaling so that one bit is provided per HPB. However, to support also more uplink-limited use cases, additional bundling of HARQ feedback can be used. This is done dynamically by computing the bundling factor needed to make all requested feedback fit within the provided UCI resource. Feedback related to different DCIs and different retransmission states are bundled independently.

6 Mobility and beam management

Beamforming is an essential aspect when operating in the sub-THz range to achieve the desired array gain. At these very high frequencies, analog beamforming is a likely choice for an implementation. With analog beamforming, beam-management procedures are needed to select the transmit and receive beam both at the UE and base station. Fundamentally, PDCCH, PDSCH and PUSCH are all transmitted and received over the same beam. Since a much larger number of beams are expected when operating at sub-THz frequencies compared to the case of NR, beam alignment between the base station and the UE is an important aspect to be evaluated in order to reduce the paring time and improve the precision. Even though the concepts described below are very similar to those in 5G-NR (for example, hierarchical beam refinement around SSB-level beams), the designed framework still allows the evaluation of many implementation-specific ideas on both base station and UE sides, to improve pairing latency, reliability, and recovery speed.

6.1 Beam management with and without beam indication

The concept supports beam management with and without beam indication. In beam management without beam indication, the UE and the base station update their beams independently and apply a new best beam without informing the other node. For beam management with beam indication, the UE and base station simultaneously change beams at a predetermined point in time.

In both cases, the beam adjustment is based on measurements on a set of aperiodic CSI-RS resources. These CSI-RS resources are transmitted in non-overlapping time resources.

The base station triggers the CSI-RS transmission in DCI, which provides the format of the upcoming CSI-RS transmission. The format describes how many times the CSI-RS resource is repeated in the same Tx beam, and how many such repeated bundles of CSI-RS resources are transmitted.

Beam management without beam indication is performed independently for the base station and the UE. To determine its beam, the base station transmits one CSI-RS resource in each beam it wants to probe and selects the beam based on measurement reports from the UE. The selected beam is used for subsequent PDCCH and PDSCH transmissions, as well as for subsequent PUSCH receptions. To let the UE determine its beam, the base station transmits all the CSI-RS resources in the same beam, which allows the UE to probe different Rx beams, and select the beam that results in the largest RSRP. That Rx beam is used for subsequent PDCCH and PDSCH receptions, as well as PUSCH transmissions.

To perform beam management with beam indication, the base station includes a reference to an SSB in the DCI that triggers the aperiodic CSI-RS measurement. The UE uses the Rx beam corresponding to the referenced SSB when it performs the measurement on the CSI-RS. The UE measures the RSRP for the CSI-RS resources and reports the CSI-RS with the largest RSRP. A certain time after the measurement report has been sent, the UE changes its Rx beam.

The measurements on aperiodic CSI-RS are suitable to probe a limited set of beams, typically but not necessarily adjacent to the currently used beam. To find beams from other base stations, measurements on SSBs can be used. The UE measures RSRP and reports a number of SSB indices and PCIs over MAC. The base station can use this report to trigger aperiodic measurements and subsequent beam updates.

The beam management procedures are designed to handle the case where one TRP transmits data in several beams but are equally applicable to the case where the beams are transmitted from different TRPs.

6.2 Beam failure detection and recovery

If the beam management algorithms fail, the base station may lose track of the UE; it may not know which beam to use to reach the UE. The beam recovery procedure is designed to handle this situation. Beam recovery allows the UE to detect that the base station is unable to reach the UE (beam-failure detection) and it also allows the UE to recover the connection with a base station (beam-failure recovery, or beam recovery).

Beam recovery is performed using contention-based random access to any base station. The UE sends a PRACH preamble, and in the subsequent Msg-3 transmission, it includes its C-RNTI. The base station uses the information in Msg-3 to recover the connection with the UE.

Beam-failure detection is based on an estimate of the PDCCH block-error probability. For every PDCCH the UE detects, the UE estimates the block-error probability based on the PDCCH DM-RS. A received PDCCH where the block-error probability is estimated to be smaller than 10% is considered a good PDCCH reception. At every recovery opportunity, that is, every RACH occasion, the UE checks how many such a good PDCCH reception occurred during the last 10 ms. If no such good PDCCH receptions occurred during the last 10 ms, the UE starts the beam recovery procedure. Obviously, to avoid triggering the beam recovery procedure the base station must at least transmit one PDCCH to the UE in the 10 ms that precedes a random-access occasion.

6.3 Multi-leg communication

Communication with multiple TRPs at the same time, using different antenna panels, is supported, which can be critical to mitigate objects blocking the communication between the UE and the base station. A UE can maintain up to two legs at the same time, that is, can simultaneously communicate with up to two transmission points. Once a second leg is established, the communication is self-contained: each leg has its own PDCCH, which schedules PDSCH or PUSCH. No cross-leg scheduling is supported. Power control and TA adjustment (section 77) are performed independently per leg.

The multi-leg operation is managed using extensions to the procedures described in section 6.1. A DCI that triggers an aperiodic CSI-RS measurement with an SSB reference may include a leg ID. If this leg ID is unknown to the UE, the UE activates a new leg and associates it with the beam it reported. On the other hand, if the leg ID is known to the UE, the UE updates the beam of that leg. The leg remains activate until it is explicitly deactivated by the network, using a MAC control element.

The network can use the SSB reports to identify which TRPs can be used in multi-connectivity. There are two main constraints, namely 1) the pathloss to the TRPs must be small enough, and 2) the communication must be performed using different UE panels. The UE includes information in the SSB report that makes it possible for the network to determine which TRPs to involve in the multi-connectivity.

7 Power control and timing advance

In contrast to NR, where a combination of open-loop power control based on the pathloss estimated from a DL reference signal and closed-loop power control is supported, only closed-loop power control is used in the trial concept in order to simplify the overall scheme. The network sends power-control commands in the DCI, which are used to control the uplink transmit power in the UE. In every DCI, the PUSCH transmit power is adjusted in steps of {-1, 0, +1, +3} dB. The UE also adjusts its transmit power based on the used MCS and the number of scheduled PRBs, following the same principles as in NR.

To aid scheduling decisions, power headroom (PHR) reports are transmitted to the network in a similar way as in NR. The PHR is based on a virtual PUSCH allocation. Transmission of a report is triggered based on a timer, or after an aperiodic CSI-RS report with SSB reference.

When beam management with beam indication is performed, that is, when the UE is triggered to perform a measurement on CSI-RS with an SSB reference, the UE resets its accumulated TPC commands and adjusts its transmit power using the pathloss estimated from the CSI-RS with the largest measured RSRP. Note that this adjustment is only performed once; for all subsequent transmissions, the UE just relies on the closed-loop power control commands and resource allocation sent in DCI.

For random access, open-loop power control is used. The PRACH transmit power is set based on the pathloss estimated from the selected SSB and the PRACH SNR target. In case the UE does not detect Msg-2 in response to the PRACH transmission, the PRACH transmission is repeated with an increased Tx power.

Timing advance can be adjusted using a MAC control element. In addition, after the UE performs an RSRP measurement on an aperiodic CSI-RS with an SSB reference, the UE will autonomously adjust its transmission timing under the assumption that the downlink transmissions are synchronized, and that the observed difference in time-of-arrival of two downlink reference signals is only caused by a difference in propagation delay. This reduces the need for sending timing-advance commands in the downlink.

8 Initial access

Initial access uses a similar approach as NR, that is, SSBs transmitted by the network used for downlink synchronization, and random access in the uplink for uplink synchronization.

For the various array sizes considered for evaluation, up to 256 SSB transmissions using 256 beams within an SSB period of 40 ms are supported. An SSB consists of 14 consecutive symbols, split into 4 symbols PSS, 4 symbols SSS, and 6 symbols PBCH. The PSS and SSS are used to obtain downlink time synchronization and are based on m-sequences and complementary Golay sequence pairs, respectively. The SSS also serves as a demodulation reference signal when receiving the PBCH. The PBCH carries the master information block (MIB) with a content similar to NR but with no physical-layer-created payload. All components of the SSB are transmitted with 2 port transmit diversity and using DFTS-OFDM.

Upon detection of an SSB, the UE transmits a random-access preamble in the associated random-access opportunity. Four SSBs are mapped to the same PRACH opportunity. If the base station antenna setup supports simultaneously multiple, individually steerable, reception panels/beams, up to four PRACH opportunities can be monitored in parallel. The four random-access opportunities are allocated to different frequencies in the same symbol and form one random-access monitoring occasion, see Figure 8-1. In total, 16 random-access monitoring occasions are needed to map 256 SSBs to random-access opportunities.

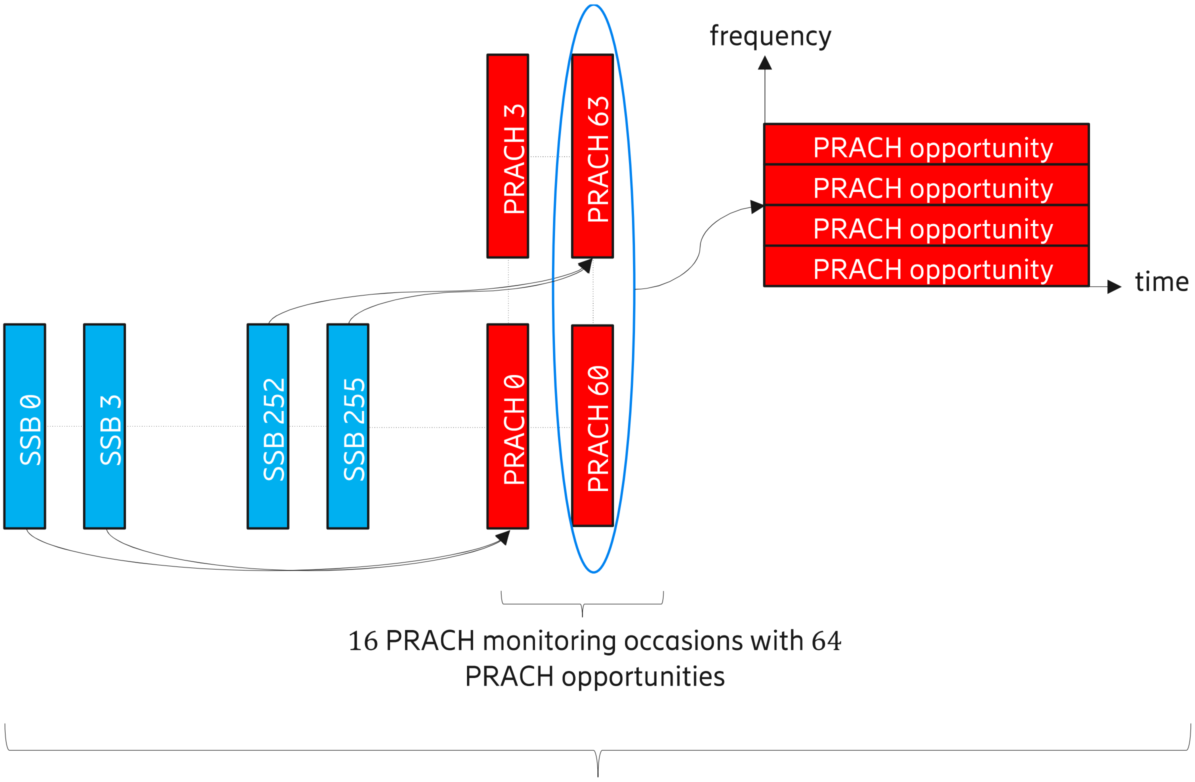

Figure 8-1: Mapping between SSB and random-access opportunities.

After transmitting a random-access preamble, the UE uses a similar random-access procedure as in NR, starting with reception of Msg-2. Noteworthy differences to NR include that Msg-2 is directly transmitted as DCI on the PDCCH (since it has a small payload) using DCI format 2, and that after receiving the timing-advance command in Msg-2 the UL timing can still be ambiguous in symbol timing (due to beamformed PRACH reception an UL symbol-timing ambiguity remains after PRACH). This timing uncertainty can be resolved by measuring the arrival timing of Msg-3 and sending a second timing advance command in Msg-4.

9 Conclusion

A concept for sub-THz communication has been described. Operating at such a high frequency range is challenging. New building practices compared to existing frequency bands are needed and commercial RF components are not yet mature. Baseband implementation capable of processing the data rates expected in the sub-THz range, many tens of Gbit/s, is another challenge. Consequently, the concept outlined in this document is different compared to NR in several aspects, some of which may be applicable for future radio access also in lower frequency bands.