
cttc-nr-demo tutorial

Release 3.0

Mar 15, 2024

CONTENTS

1	Acknowledgements	1
2	Acronyms	3
3	Overview	5
3.1	Program Overview	6
3.2	References	7
4	End-to-end observations	9
5	RAN lifecycle	11
5.1	EpcEnbApplication	11
5.2	NrGnbNetDevice	16
5.3	LteEnbRrc	16
5.4	LtePdcP	18
5.5	LteRlcUm	20
5.6	NrGnbMac	23
5.7	NrGnbPhy	25
5.8	NrUePhy	27
5.9	NrUeMac	29
5.10	LteRlcUm - UE Side	30
5.11	LtePdcP - UE Side	31
5.12	NrUeNetDevice	31
6	Conclusions	33

ACKNOWLEDGEMENTS

This tutorial has been originally created by Giovanni Grieco in the scope of GSoC 2023 project [IUNS-3 5G NR: Improving the Usability of ns-3's 5G NR Module](#) mentored by Tom Henderson, Katerina Koutlia, and Biljana Bojovic.

The main idea of this ns-3 GSoC project was to ease the learning curve for new users of the NR module. To achieve this goal, the first version of the NR tutorial created by Giovanni Grieco focuses on a `cttc-nr-demo` example and explains the internal functionality of the NR RAN by providing a detailed, layer-by-layer insights on the packet lifecycle as they traverse the RAN. The tutorial highlights all the points in the NR protocol stack where packets may be delayed or dropped, and how to log and trace such events. The first version of this tutorial was released with NR 2.6 in November 2023.

Table 1.1: Changelog

Version	Release Date	Authors
1.0	N/A	Giovanni Grieco < giovanni@grieco.dev > Tom Henderson < tomh@tomh.org > Katerina Koutlia < katerina.koutlia@cttc.es > Biljana Bojović < biljana.bojovic@cttc.es >

ACRONYMS

- 5th Generation (5G)
- Acknowledged Mode (AM)
- Acknowledgement (ACK)
- Bandwidth Part (BWP)
- Bearer Identifier (BID)
- Carrier Component (CC)
- Data/Control (D/C)
- Data Radio Bearer Identifier (DRBID)
- Downlink (DL)
- Downlink Control Indicator (DCI)
- Evolved Node-B (eNB)
- Evolved Packet Core (EPC)
- External Interface – User Plane (X2-U)
- Finite State Machine (FSM)
- General Packet Radio Service (GPRS)
- GPRS Tunneling Protocol – User data tunneling (GTP-U)
- Guaranteed Bit Rate (GBR)
- Head-of-Line (HOL)
- Hybrid Automatic Repeat Request (HARQ)
- Identifier (ID)
- International Mobile Subscriber Identity (IMSI)
- Internet Protocol version 4 (IPv4)
- Line of Sight (LoS)
- Listen Before Talk (LBT)
- Logical Channel Identifier (LCID)
- Long Term Evolution (LTE)
- Modulation and Coding Scheme (MCS)

- Maximum Transmission Unit (MTU)
- Medium Access Control (MAC)
- New Radio (NR)
- Network Simulator 3 (ns-3)
- Next-Generation Node-B (gNB)
- Non-Access Stratum (NAS)
- Non-Guaranteed Bit Rate (NGBR)
- Non-Standalone (NSA)
- Packet Data Convergence Protocol (PDCP)
- Packet Gateway (PGW)
- Power Spectrum Density (PSD)
- Protocol Data Unit (PDU)
- Quality of Service Class Indicator (QCI)
- Radio Access Network (RAN)
- Radio Link Control (RLC)
- Radio Network Temporary Identity (RNTI)
- Radio Resource Control (RRC)
- Random Access Control Channel (RACH)
- Receive (RX)
- Resource Block (RB)
- Round Robin (RR)
- Service Access Point (SAP)
- Service Data Unit (SDU)
- Serving Gateway (SGW)
- Signal-to-Interference-plus-Noise Ratio (SINR)
- System Frame Number / Subframe Number (SFN/SF)
- Time Division Duplex (TDD)
- Transmit (TX)
- Transmission Time Interval (TTI)
- Transparent Mode (TM)
- Transport Block (TB)
- Tunnel Endpoint Identifier (TEID)
- Unacknowledged Mode (UM)
- Uplink (UL)
- User Datagram Protocol (UDP)
- User Equipment (UE)

OVERVIEW

This is a tutorial focused on the data plane operation of the example program `cttc-nr-demo` found in the `examples/` directory of the `ns-3 nr` module. The objective of the tutorial is to provide a detailed, layer-by-layer walk through of a basic NR example, with a focus on the typical lifecycle of packets as they traverse the RAN. The tutorial points out all of the locations in the RAN model where packets may be delayed or dropped, and how to trace such events.

This document assumes that you have already installed `ns-3` with the `nr` module and you are familiar with how `ns-3` works. If this is not the case, please review the `nr` module [README](#), the [ns-3 Installation Guide](#) and its [Tutorial](#) as needed. The [Getting Started](#) page of the `nr` module should also be reviewed.

The companion to this tutorial is the detailed [manual](#) for the `nr` module, which goes into more detail about the design and testing of each of the components of the 5G NR module.

To check if you are ready to work through this tutorial, check first if you can run the following program:

```
$ ./ns3 run cttc-nr-demo
```

and that it outputs the following output

```
Flow 1 (1.0.0.2:49153 -> 7.0.0.2:1234) proto UDP
  Tx Packets: 6000
  Tx Bytes:   768000
  TxOffered:  10.240000 Mbps
  Rx Bytes:   767744
  Throughput: 10.236587 Mbps
  Mean delay: 0.271518 ms
  Mean jitter: 0.030032 ms
  Rx Packets: 5998
Flow 2 (1.0.0.2:49154 -> 7.0.0.3:1235) proto UDP
  Tx Packets: 6000
  Tx Bytes:   7680000
  TxOffered: 102.400000 Mbps
  Rx Bytes:   7667200
  Throughput: 102.229333 Mbps
  Mean delay: 0.900970 ms
  Mean jitter: 0.119907 ms
  Rx Packets: 5990

Mean flow throughput: 56.232960
Mean flow delay: 0.588507
```

The tutorial also makes extensive use of the `ns-3` logging framework. To check if logs are enabled in your `ns-3` libraries, try the following command and check if it outputs some additional verbose output:

```
$ NS_LOG="CttcNrDemo" ./ns3 run cttc-nr-demo
```

The command should print the following informational message on screen:

```
+0.000000000s -1 CttcNrDemo:main(): [INFO ] Creating 2 user terminals and 1 gNBs
```

It can be observed that the message is accompanied by some contextual information. From left to right, the message also tells us the simulation time at which the message has been produced, the node id, what object and function is producing such message, and the logging level. In some cases, the node id may be `-1` because here the code is independent by any node operating in the simulation. Later you will discover that such number will be positive in case the code being executed is dependent on which node is acting in the simulation, e.g., for a transmission or reception of some packets.

3.1 Program Overview

From what it can be deduced, the demo simulates two downlink *Flows*, each of them relying on a unicast and uni-directional communication. Such flows rely on the UDP to carry application data from an origin with IPv4 `1.0.0.2` to two recipients with IPv4 `7.0.0.2` and `7.0.0.3` for *Flow 1* and *Flow 2*, respectively.

The purpose of this example is to simulate a downlink scenario. Two data flows originate from a remote host, with specific characteristics. One flow emphasizes low-latency communications, while the other focuses on achieving a higher throughput. In the provided demo output, it is evident that the former exhibits significantly lower mean delay and jitter compared to the latter, whereas the opposite is true for the achieved throughput. In the code, the low-latency communication is referred to as `LowLat` to indicate its low-latency nature, while the one that achieves higher throughput is referred to as `Voice` to reflect the traditional traffic associated with high-quality voice communications.

For this communication, the source is an IPv4 address, specifically `1.0.0.2`, which is referred to as the “remote host.” The recipients of the data are two UEs.

To support such communications, a 5G Radio Access Network is configured, together with an LTE Core Network, referred to as the EPC. The entire architecture is defined as 5G NSA.

This demo is characterized by quasi-ideal conditions. For instance, the S1-U link, which interconnects the SGW with the gNB, has no delay. Furthermore, Direct Path Beamforming is used, which is an ideal algorithm based on the assumption that the transmitters always know the exact location of the receivers. Further information can be retrieved on the [manual](#), at Section 2.3.9. Shadowing is not considered, as buildings and any other kind of obstacles that could impair normal LoS conditions are absent. Finally, the channel model is updated only once, at the start of the simulation, given that the scenario is static, i.e., it does not change over time.

While both UEs are characterized by a Uniform Planar Array of 2x4 isotropic antenna elements, the gNB has the same array with a configuration of 4x8.

In terms of spectrum, 2 bands are created to support such communications. The first one operates at 28.0 GHz, while the second one at 28.2 GHz, both with a bandwidth of 50 MHz. In terms of numerology, i.e., the sub-carrier spacing, the former is 4, while the latter is 2. This simplifies spectrum allocation, given that each communication will operate on a dedicated BWP, on a single CC that occupies the entire band, resulting in the spectrum organized as below:

```
* The configured spectrum division is:
* -----Band1-----|-----Band2-----
* -----CC1-----|-----CC2-----
* -----BWP1-----|-----BWP2-----
```

Given that there is only one gNB, a total transmission power of 35 dBm, which is around 3.16 W, is spread among the two BWPs.

In terms of the BWP type and bearer, the former communication is configured to use a QCI with NGBR Low Latency, also known in the code as `NGBR_LOW_LAT_EMBB`, while the latter has a QCI with GBR and is named as `GBR_CONV_VOICE`. A list of other QCI types can be found at the [ns-3 doxygen page on QCI](#).

On the top of the stack, two UDP applications are used to simulate low-latency voice traffic and high quality one, respectively. In terms of the generated traffic, the former simulates a burst of 100 bytes each 100 us, while the latter generates packets of 1252 bytes every 100 us. The `FlowMonitorHelper` is used to gather data statistics about the traffic.

Finally, the EPC's PGW is then connected to a remote host with an ideal Point-to-Point channel: 100 Gbps of data rate with 2500 bytes of MTU and no delay.

3.2 References

[cttc-nr-demo] cttc-nr-demo program. Available at: <https://gitlab.com/cttc-lena/nr/-/blob/master/examples/cttc-nr-demo.cc>

END-TO-END OBSERVATIONS

We are mainly interested in observing the packet lifecycle as it moves through the RAN stack. We can make a few initial observations about the packet flow through logging.

Using the following logging-enabled command, it is possible to generate the log information from the `UdpClient` and `UdpServer` objects in the simulation, which simulate the two different types of communications previously mentioned, and redirect the output to two files, as follows:

```
$ NS_LOG="UdpClient=info|prefix_time|prefix_node|prefix_func" ./ns3 run 'cttc-nr-demo
↪' > log.client.out 2>&1
$ NS_LOG="UdpServer=info|prefix_time|prefix_node|prefix_func" ./ns3 run 'cttc-nr-demo
↪' > log.server.out 2>&1
```

Looking at the first couple of lines of the `log.client.out` file, one can see:

```
+0.400000000s 6 UdpClient:Send(): TraceDelay TX 100 bytes to 7.0.0.2 Uid: 8 Time: +0.
↪4s
+0.400000000s 6 UdpClient:Send(): TraceDelay TX 1252 bytes to 7.0.0.3 Uid: 9 Time: +0.
↪4s
```

These first two packets were sent at the same time to two different UEs, from node 6, which identifies the node that simulates the remote host. Next, observe the first packet arrivals on the UEs via the `log.server.out` file:

```
+0.400533031s 1 UdpServer:HandleRead(): TraceDelay: RX 100 bytes from 1.0.0.2
↪Sequence Number: 0 Uid: 8 TXtime: +4e+08ns RXtime: +4.00533e+08ns Delay: +533031ns
...
+0.402582140s 2 UdpServer:HandleRead(): TraceDelay: RX 1252 bytes from 1.0.0.2
↪Sequence Number: 0 Uid: 9 TXtime: +4e+08ns RXtime: +4.02582e+08ns Delay: +2.
↪58214e+06ns
```

The reception times (and packet delays) are quite different. One takes only 533 us to be delivered, while the other takes 2582 us to be delivered. In this tutorial, we will explain why this is so.

RAN LIFECYCLE

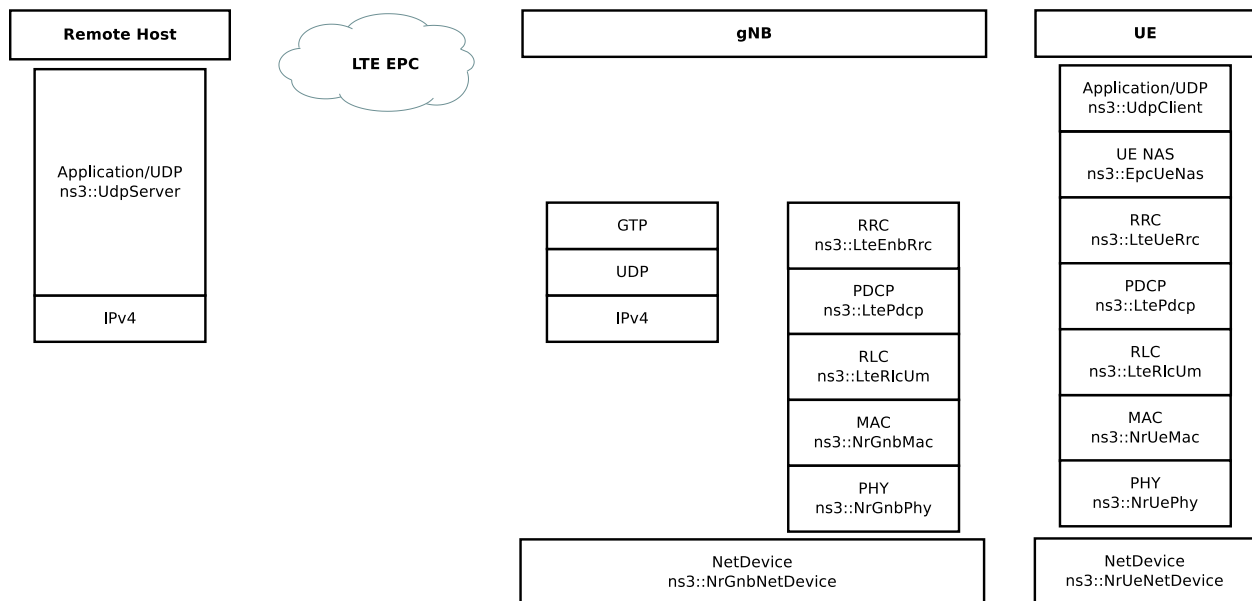


Fig. 5.1: Schematic topology

Fig. 5.1 depicts the objects that each packet will traverse through its lifecycle.

For reference, Fig. 5.2 and Fig. 5.3 are also provided as a map in order to better follow the packet being processed by each actor of the RAN, i.e., gNB and UE, respectively.

This tutorial will walk through each step of the way, starting with the entry point for these packets in the RAN– the `EpcEnbApplication`.

5.1 EpcEnbApplication

The `EpcEnbApplication` is installed on the gNB. It is responsible for receiving packets tunneled through the EPC model and sending them into the `NrGnbNetDevice`. Conceptually, this is just an application-level relay function. It is possible to run the `cttc-nr-demo` scenario by enabling the `EpcEnbApplication` log component, which can also be configured to print messages at INFO level and prefix the message by stating the time of the simulation, the node ID of interest, and the routine that prints that message. In this way, with the following command

```
$ NS_LOG="EpcEnbApplication=info|prefix_all" ./ns3 run 'cttc-nr-demo' > log.out 2>&1
```

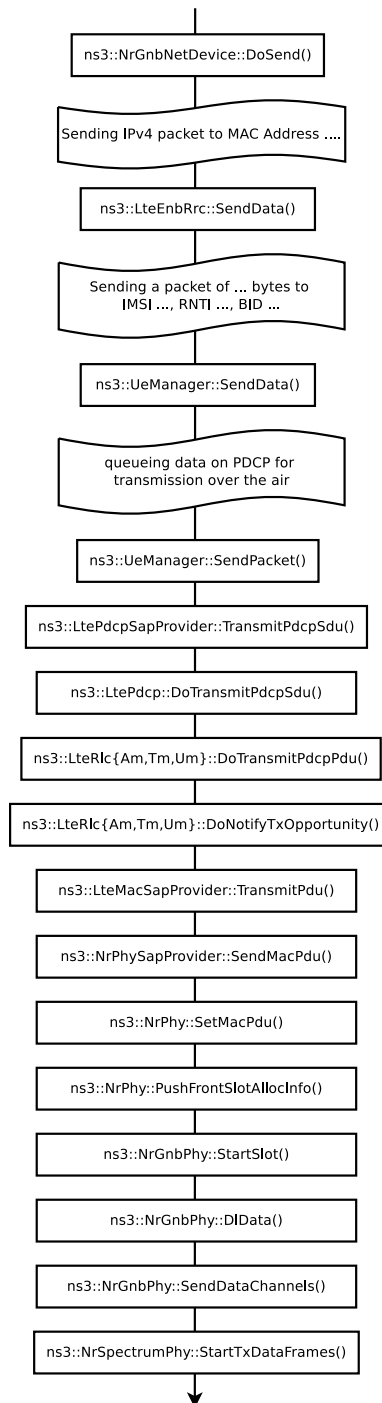


Fig. 5.2: gNB's downlink packet processing pipeline

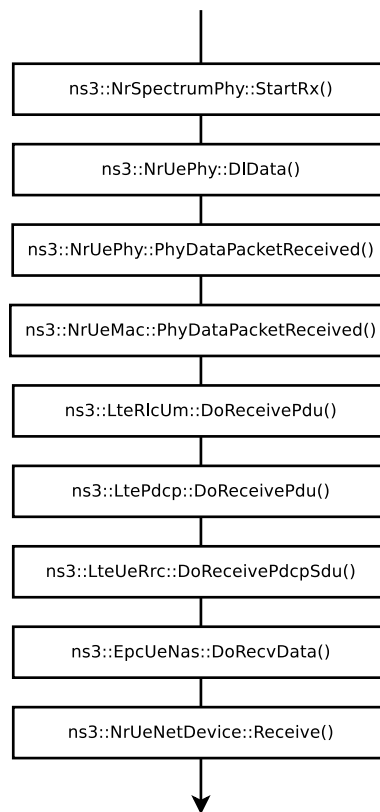


Fig. 5.3: UE's downlink packet processing pipeline

one can observe this relay function on the first packet, as follows:

```
+0.400000282s 0 EpcEnbApplication:RecvFromS1uSocket(): [INFO ] Received packet from
↳S1-U interface with GTP TEID: 2
+0.400000282s 0 EpcEnbApplication:SendToLteSocket(): [INFO ] Add EpsBearerTag with
↳RNTI 2 and bearer ID 2
+0.400000282s 0 EpcEnbApplication:SendToLteSocket(): [INFO ] Forward packet from eNB
↳'s S1-U to LTE stack via IPv4 socket.
```

The file `src/lte/model/epc-enb-application.cc` contains the source code.

The `EpcEnbApplication` adds the cell-specific UE ID (RNTI) and BID as tags through `EpsBearerTag`. This greatly simplifies packet processing in later sections. The application is able to find the right RNTI and BID given the TEID present on the GTP-U header of the received packet structure. This process can be found at `EpcEnbApplication::RecvFromS1uSocket()` source code:

```
GtpuHeader gtpu;
packet->RemoveHeader(gtpu);
uint32_t teid = gtpu.GetTeid();
std::map<uint32_t, EpsFlowId_t>::iterator it = m_teidRbidMap.find(teid);
if (it == m_teidRbidMap.end())
{
    NS_LOG_WARN("UE context at cell id " << m_cellId << " not found, discarding packet
↳");
    m_rxDropS1uSocketPktTrace(packet->Copy());
}
else
{
    m_rxS1uSocketPktTrace(packet->Copy());
    SendToLteSocket(packet, it->second.m_rnti, it->second.m_bid);
}
```

Notice that the hashmap `<uint32_t, EpsFlowId_t>` links together a TEID, handled by `uint32_t`, with RNTI and BID, grouped by `EpsFlowId_t` data structure.

The packet and its modifications can be visually represented as shown in Fig. 5.4, where removed portions of the packet are marked in red, whereas green ones are the added portions.

There is also a byte tag to trace the packet for flow statistics, which are bound to `Flow Monitor` and are relevant to the final results that are printed by the scenario.

The new modified packet is finally sent to `EpcEnbApplication::SendToLteSocket()`, which distinguishes the packet to be sent according to its L3 type, in this case IPv4.

As a final remark, the following traces can be used to track incoming and outgoing packets from this application:

- `RxFromEnb`: Receive data packets from LTE Enb Net Device
- `RxFromS1u`: Receive data packets from S1-U Net Device
- `RxFromS1uDrop`: Drop data packets from S1-U Net Device
- `TxToEnb`: Transmit data packets to LTE eNB Net Device
- `TxToS1u`: Transmit data packets to S1-U Net Device

Packet Tags



Byte Tags



Packet Contents




Fig. 5.4: Structure of the packet handled by `EpcEnbApplication`

5.1.1 Packet latency

Packets cannot incur latency in the `EnbEpcApplication`.

5.1.2 Packet drops

Packets can be dropped if there is no association between GTP-U TEID and gNB-UE link's RNTI and BID.

5.2 NrGnbNetDevice

After the `EpcEnbApplication` sends a packet, it is immediately received on the `NrGnbNetDevice::DoSend()` method. We can observe this in the logs:

```
$ NS_LOG="NrGnbNetDevice=info|prefix_all" ./ns3 run 'cttc-nr-demo' > log.out 2>&1
```

```
+0.400000282s 0 NrGnbNetDevice::DoSend(): [INFO ] Forward received packet to RRC Layer
```

The source code for this method is in the file `contrib/nr/model/nr-gnb-net-device.cc`. As an aside, notice how we are bouncing back and forth between the source code directories `src/lte/` (for the previous `EpcEnbApplication`) and `contrib/nr/` (for this object); this is true for the upper layers of the current `nr` module, which reuse pieces originally implemented for LTE.

The source code reveals that it is possible to trace the packet via `NrNetDevice/Tx`.

`NrGnbNetDevice` does not process the packet, but it just relays it to the RRC layer.

5.2.1 Packet latency

Packets cannot incur latency in the `NrGnbNetDevice`.

5.2.2 Packet drops

Packets cannot be dropped in the `NrGnbNetDevice`.

5.3 LteEnbRrc

The RRC layer of the gNB is handled by `LteEnbRrc` and its companion `UeManager`, both available in `lte/model/lte-enb-rrc.cc`. While packets incoming from the upper layers are processed by `LteEnbRrc::SendData()`, packets to be sent to UEs are handled by `UeManager::SendPacket()`. Indeed, if the simulation is started with the following command

```
$ NS_LOG="LteEnbRrc=info|prefix_all" ./ns3 run cttc-nr-demo &> out.log
```

it produces the following log messages:

```
+0.400000282s 0 LteEnbRrc:SendData(): [INFO ] Sending a packet of 128 bytes to IMSI 1,
↳ RNTI 2, BID 2
+0.400000282s 0 LteEnbRrc:SendData(): [INFO ] queueing data on PDCP for transmission_
↳ over the air
+0.400000282s 0 LteEnbRrc:SendPacket(): [INFO ] Send packet to PDCP layer
```

By taking a look at the source code, the RRC layer starts by extracting the RNTI, which is handled by the `EpsBearerTag` data structure. Thanks to the RNTI, it is possible to retrieve the corresponding `UeManager` instance, which keeps track of the gNB-UE link state through a FSM. The search is done by `LteEnbRrc`, which uses a hashmap that links each RNTI to a pointer of the `UeManager` of interest. Once found, the matching `SendData()` method is called with the packet's BID of reference.

```
EpsBearerTag tag;
bool found = packet->RemovePacketTag(tag);
NS_ASSERT_MSG(found, "no EpsBearerTag found in packet to be sent");
Ptr<UeManager> ueManager = GetUeManager(tag.GetRnti());

ueManager->SendData(tag.GetBid(), packet);
```

The `SendData()` method can be quickly understood with the following code excerpt:

```
switch (m_state)
{
case INITIAL_RANDOM_ACCESS:
case CONNECTION_SETUP:
    NS_LOG_WARN("not connected, discarding packet");
    m_packetDropTrace(p, bid);
    return;

case CONNECTED_NORMALLY:
case CONNECTION_RECONFIGURATION:
case CONNECTION_REESTABLISHMENT:
case HANDOVER_PREPARATION:
case HANDOVER_PATH_SWITCH: {
    NS_LOG_LOGIC("queueing data on PDCP for transmission over the air");
    SendPacket(bid, p);
}
break;
// ...
}
```

Let's ignore the states related to handover, given that in the NR module the X2-U interface is not implemented, yet.

If the UE is ready to receive the packet from the gNB viewpoint, i.e. it's in `CONNECTED_NORMALLY`, the `SendPacket()` method is called to continue with the transmission:

```
LtePdcpsapProvider::TransmitPdcpsduParameters params;
params.pdcpsdu = p;
params.rnti = m_rnti;
params.lcid = Bid2Lcid(bid);
uint8_t drbid = Bid2Drbid(bid);
// Transmit PDCP sdu only if DRB ID found in drbMap
std::map<uint8_t, Ptr<LteDataRadioBearerInfo>>::iterator it = m_drbMap.find(drbid);
if (it != m_drbMap.end())
{
    Ptr<LteDataRadioBearerInfo> bearerInfo = GetDataRadioBearerInfo(drbid);
    if (bearerInfo)
```

(continues on next page)

(continued from previous page)

```

{
    NS_LOG_INFO("Send packet to PDCP layer");
    LtePdcPcapProvider* pdcPcapProvider = bearerInfo->m_pdcPcap->
    GetLtePdcPcapProvider();
    pdcPcapProvider->TransmitPdcPcapSdu(params);
}
}

```

Over there, the packet is processed by creating a PDCP SDU structure with complementary information, known as `LtePdcPcapProvider::TransmitPdcPcapSduParameters`. These parameters are the RNTI and other two parameters which are dependent of the BID, such as LCID, which in this case is 4, and DRBID, which is 2. The latter is used to retrieve the `LteDataRadioBearerInfo` associated to a DRBID, if it is still in place. At last, the SDU is forwarded to the active `LtePdcPcapProvider` through `TransmitPdcPcapSdu()`.

With this procedure in mind, it is possible to understand how the packet is modified. If we run the same simulation by enabling the `LteEnbRrc` log component, we obtain the following message logs:

```

+0.400000282s 0 LteEnbRrc:SendData(): [INFO ] Received packet
+0.400000282s 0 LteEnbRrc:SendPacket(): [INFO ] Send packet to PDCP layer

```

As it can be noticed, packet tags are removed, as the RNTI and BID are transferred to the `LtePdcPcapProvider::TransmitPdcPcapSduParameters` instance. Byte tags and packet structure remain unaltered.

5.3.1 Packet latency

Packets cannot incur latency in `LteEnbRrc`.

5.3.2 Packet drops

There may be a chance that the UE is still asking for resources on the RACH, i.e., `INITIAL_RANDOM_ACCESS`, or there is still pending set up, i.e., `CONNECTION_SETUP`. In that case, the packet cannot be transmitted. To this end, it is placed in the `UeManager/Drop` trace source and discarded.

5.4 LtePdcPcap

The `LteEnbRrc` sends the packet down the stack by referencing a `LtePdcPcapProvider`. This is an abstract interface in order to implement any kind of PDCP layer. In this case, the simulation makes use of the implementation provided by `LtePdcPcap` available in `lte/model/lte-pdcPcap.cc`.

The `LtePdcPcap` log component can be enabled at INFO level to check out these log messages:

```

+0.400000282s 0 LtePdcPcap:DoTransmitPdcPcapSdu(): [INFO ] Received PDCP SDU
+0.400000282s 0 LtePdcPcap:DoTransmitPdcPcapSdu(): [INFO ] Transmit PDCP PDU

```

According to the `LtePdcPcapProvider` interface, the internal method used to receive an incoming packet from upper layers is called `DoTransmitPdcPcapSdu()`, which in this case is defined as the following:

```

Ptr<Packet> p = params.pdcPcapSdu;

// Sender timestamp
PdcPcapTag pdcPcapTag(Simulator::Now());

```

(continues on next page)

(continued from previous page)

```

LtePdcphHeader pdcphHeader;
pdcphHeader.SetSequenceNumber(m_txSequenceNumber);

m_txSequenceNumber++;
if (m_txSequenceNumber > m_maxPdcphSn)
{
    m_txSequenceNumber = 0;
}

pdcphHeader.SetDcBit(LtePdcphHeader::DATA_PDU);

NS_LOG_LOGIC("PDCP header: " << pdcphHeader);
p->AddHeader(pdcphHeader);
p->AddByteTag(pdcphTag, 1, pdcphHeader.GetSerializedSize());

m_txPdu(m_rnti, m_lcid, p->GetSize());

LteRlcSapProvider::TransmitPdcphPduParameters txParams;
txParams.rnti = m_rnti;
txParams.lcid = m_lcid;
txParams.pdcphPdu = p;

m_rlcSapProvider->TransmitPdcphPdu(txParams);

```

As it can be observed, a PDCP packet tag, identified by `PdcphTag`, and its header, represented by `LtePdcphHeader`, are created and attached to the SDU. The former stores the current simulation time and the latter is set with a specified sequence number and a D/C bit. While the sequence number is locally tracked through the `m_txSequenceNumber` property, the D/C bit is set to `DATA_PDU` to indicate that this is a data packet and not a control one. This change can be visually represented as [Fig. 5.5](#).

Finally, the packet is forwarded to the RLC layer through `LteRlcSapProvider::TransmitPdcphPdu()`, by providing some additional parameters to the PDCP PDU, such as the RNTI and the LCID, under the `LteRlcSapProvider::TransmitPdcphPduParameters` structure.


5.4.1 Packet latency

Packets cannot incur latency in the `LtePdcph`.

Packet Tags

N/A

Byte Tags



Packet Contents



Fig. 5.5: Structure of the PDCP PDU processed by `LtePdcP`. Green blocks are added at the PDCP layer, while blue ones represent unchanged structure.

5.4.2 Packet drops

Packets cannot be dropped in the `LtePdcP`.

5.5 LteRlcUm

Following the same design of `LtePdcP` template class and its implementation `LtePdcP`, the RLC layer is based on the same class structure. Here, the `LteRlcSapProvider` is a template class where multiple implementations of the RLC can be defined. On *ns-3*, the RLC layer is implemented by `LteRlcAm`, `LteRlcTm`, and `LteRlcUm`, according to how the PDCP PDU is transferred to the UE, if on Acknowledged Mode (AM), Transparent Mode (TM), or Unacknowledged Mode (UM), respectively. For this simulation, the `LteRlcUm` is adopted, which is implemented in `lte/model/lte-rlc-um.cc`.

If the `LteRlcUm` log component is enabled at INFO level, it produces the following log messages:

```
+0.400000282s 0 LteRlcUm:DoTransmitPdcPpdu(): [INFO ] Received RLC SDU
+0.400000282s 0 LteRlcUm:DoTransmitPdcPpdu(): [INFO ] New packet enqueued to the RLC
→Tx Buffer
```

As it can be deduced, the RLC SDU is received at `LteRlcUm::DoTransmitPdcPpdu()`, with its implementation that is:

```
if (m_txBufferSize + p->GetSize() <= m_maxTxBufferSize)
{
    if (m_enablePdcPdiscarding)
        // Ignored, as this flag is false by default

    /** Store PDCP PDU */
    LteRlcSduStatusTag tag;
    tag.SetStatus(LteRlcSduStatusTag::FULL_SDU);
    p->AddPacketTag(tag);

    NS_LOG_INFO("New packet enqueued to the RLC Tx Buffer");
```


The packet is then appended to a buffer called `m_txBuffer`, which is a `vector<TxPdu>`, where `TxPdu` is a structure containing (i) a pointer to the packet, and (ii) the time when the packet has been added to the buffer, as dictated by the `Simulator::Now()` call in `m_txBuffer.emplace_back()` instruction.

At the end of the procedure, `DoReportBufferStatus()` is called to alert the queue size at the MAC layer, handled by `LteMacSapProvider`. The buffer report focuses on reporting the queue size and its HOL delay, together with the RNTI and LCID of reference, as it can be noticed by the following excerpt of the said method:

```
Time holDelay(0);
uint32_t queueSize = 0;

if (!m_txBuffer.empty())
{
    holDelay = Simulator::Now() - m_txBuffer.front().m_waitingSince;

    queueSize =
        m_txBufferSize + 2 * m_txBuffer.size(); // Data in tx queue + estimated_
        ↪ headers size
}

LteMacSapProvider::ReportBufferStatusParameters r;
r.rnti = m_rnti;
r.lcid = m_lcid;
r.txQueueSize = queueSize;
r.txQueueHolDelay = holDelay.GetMilliSeconds();
r.retTxQueueSize = 0;
r.retTxQueueHolDelay = 0;
r.statusPduSize = 0;

NS_LOG_LOGIC("Send ReportBufferStatus = " << r.txQueueSize << ", " << r.
        ↪ txQueueHolDelay);
m_macSapProvider->ReportBufferStatus(r);
```

There is no mention of the actual `m_txBuffer`, which is kept at the RLC layer, until a transmission opportunity is found at the MAC layer, for which there is an upcall to `LteRlcUm::DoNotifyTxOpportunity()` done by the MAC scheduler. Indeed, the INFO log messages suggest that the buffer acquires two RLC SDUs before a transmission opportunity takes place:

```
+0.400000282s 0 LteRlcUm:DoTransmitPdcPdu(): [INFO ] Received RLC SDU
+0.400000282s 0 LteRlcUm:DoTransmitPdcPdu(): [INFO ] New packet enqueued to the RLC_
        ↪ Tx Buffer
+0.400002262s 0 LteRlcUm:DoTransmitPdcPdu(): [INFO ] Received RLC SDU
+0.400002262s 0 LteRlcUm:DoTransmitPdcPdu(): [INFO ] New packet enqueued to the RLC_
        ↪ Tx Buffer
+0.400062500s 0 LteRlcUm:DoNotifyTxOpportunity(): [INFO ] RLC layer is preparing data_
        ↪ for the following Tx opportunity of 36 bytes for RNTI=2, LCID=4, CCID=0, HARQ ID=15,
        ↪ MIMO Layer=0
```

The `LteRlcUm::DoNotifyTxOpportunity()` prepares the data to transmit in the following way:

```
if (txOpParams.bytes <= 2)
{
    // Stingy MAC: Header fix part is 2 bytes, we need more bytes for the data
    NS_LOG_INFO("TX opportunity too small - Only " << txOpParams.bytes << " bytes");
    return;
}

Ptr<Packet> packet = Create<Packet>();
```

(continues on next page)

(continued from previous page)

```
LteRlcHeader rlcHeader;
```

```
// Build Data field
uint32_t nextSegmentSize = txOpParams.bytes - 2;
// ...
```

First of all, the opportunity window size is checked to ensure that we have more than two bytes, given that the MAC header requires that size. Next, an empty packet is created, called `packet`. Such packet must be equivalent to the size of the transmission opportunity, for which the variable `nextSegmentSize` is used to understand how much data contained in `m_txBuffer` can be transferred. In this case, we can transfer up to 34 bytes, which is the result of the `txOpParams.bytes`, set at 36 bytes, minus 2 bytes due to the MAC header size.

```
Ptr<Packet> firstSegment = m_txBuffer.begin()->m_pdu->Copy();
Time firstSegmentTime = m_txBuffer.begin()->m_waitingSince;
// ...
while (firstSegment && (firstSegment->GetSize() > 0) && (nextSegmentSize > 0))
{
    if ((firstSegment->GetSize() > nextSegmentSize) ||
        // Segment larger than 2047 octets can only be mapped to the end of the
        ↪Data field
        (firstSegment->GetSize() > 2047))
    {
        // The packet 'firstSegment' must be fragmented to fit in our segment being
        ↪prepared for MAC TX
    }
    else if ((nextSegmentSize - firstSegment->GetSize() <= 2) || m_txBuffer.empty())
    {
        // If the packet fits the segment and there are no other packets to TX, just
        ↪add it without fragmenting it
    }
    else // (firstSegment->GetSize() < m_nextSegmentSize) && (m_txBuffer.size() > 0)
    {
        // If there are still other packets to TX, add the current packet
        ↪'firstSegment' and update nextSegmentSize
    }
}
```

From the head of `m_txBuffer`, the first RLC SDU is taken and it is called `firstSegment`, together with its waiting time. Until we have space for the transmission opportunity, tracked by `nextSegmentSize`, the while loop is executed.

Of course, if the `firstSegment` is too big to fit in the transmission opportunity, it is fragmented. The `LteRlcSduStatusTag` is then updated to check if the fragment is the `FIRST_SEGMENT`, `MIDDLE_SEGMENT`, or `LAST_SEGMENT`, according to how many parts the packet is fragmented. Fragments that do not fit the opportunity window size are inserted back to the head of `m_txBuffer`. If such fragment is not the last, the RLC Header `DATA_FIELD_FOLLOWS` bit is set to alert that this packet does not contain the entire PDCP packet. This will be useful to schedule new transmission opportunities in the future and ensure correct packet integration.

There are other cases that could happen, which are handled by the conditional clause inside the while loop. If there is only one packet left in the buffer and it fits the current segment, it is just placed to the segment. Furthermore, if there is still more space left on the segment and other packets are waiting in the `m_txBuffer`, the current packet `firstSegment` is placed in the segment and the left space, tracked by `nextSegmentSize`, is evaluated to add a new packet in the segment at the next while loop iteration.

This operation can be better visualized by observing [Fig. 5.6](#).

where the “Hdr” that precedes the MAC SDU fragment is the subheader according to [3GPP TS 38.321](#).

After the fragmentation procedure, the RLC header, handled by `LteRlcHeader` data structure, is set up. The sequence




Fig. 5.6: Overview of how the packets buffered by the RLC layer are combined to prepare a segment with a size equal to the transmission opportunity granted at the MAC layer

number is set and kept track via the local property `m_sequenceNumber`. Moreover, the framing information is set to declare if the last byte of the segment corresponds or not to the end of a RLC SDU. Finally, a `RlcTag` byte tag is added and linked to the RLC header, which saves the current simulation time.

It is possible to enable INFO log messages to understand when the packet is sent to the MAC:

```
+0.400062500s 0 LteRlcUm:DoNotifyTxOpportunity(): [INFO ] Forward RLC PDU to MAC Layer
```

Once the RLC PDU is ready, it is pushed to the MAC SAP through the `TransmitPdu()` interface method, together with the RNTI, LCID, CCID, HARQ ID, and MIMO layer.

5.5.1 Packet latency

Packets do incur latency in `LteRlcUm`, and in the RLC layer in general, according to when the transmission opportunities take place and how much large is the buffer. A small buffer may improve latency on one hand, at the cost of risking the event of a full transmission buffer, which causes packet drops.

5.5.2 Packet drops

Packets can be dropped if the RLC buffer `m_txBuffer` is full, or the `m_enablePdcPdiscarding` is enabled, which analyzes timing budget and conditions of the RLC SDU and may drop it.

5.6 NrGnbMac

The `NrGnbMac`, available in `nr/model/nr-gnb-mac.cc`, implements the MAC SAP for NR communications.

If we enable the `NrGnbMac` log component at INFO level, it is possible to see the connection initialization with the UEs, both in terms of RACH communication and RNTI allocation:

```
+0.016125000s 0 [ CellId 2, bwpId 0] NrGnbMac:DoSlotDlIndication(): [INFO ]
↳Informing MAC scheduler of the RACH preamble for RAPID 1 in slot FrameNum: 1
↳SubFrameNum: 6 SlotNum: 4; Allocated RNTI: 1
+0.016125000s 0 [ CellId 2, bwpId 0] NrGnbMac:DoSlotDlIndication(): [INFO ]
↳Informing MAC scheduler of the RACH preamble for RAPID 0 in slot FrameNum: 1
↳SubFrameNum: 6 SlotNum: 4; Allocated RNTI: 2
+0.016125000s 0 [ CellId 2, bwpId 0] NrGnbMac:SendRar(): [INFO ] In slot FrameNum: 1
↳SubFrameNum: 6 SlotNum: 2 send to PHY the RAR message for RNTI 1 rapId 1
+0.016125000s 0 [ CellId 2, bwpId 0] NrGnbMac:SendRar(): [INFO ] In slot FrameNum: 1
↳SubFrameNum: 6 SlotNum: 2 send to PHY the RAR message for RNTI 2 rapId 0 (continues on next page)
```

(continued from previous page)

At first glance, it is possible to note that these logs present two key information, i.e., `CellId` and `bwpId`. These can be used as contextual information to better follow the traffic served by a certain cell on a BWP. To learn more how these IDs work, please refer to Section 2.2 of the [NR manual](#).

Once the RLC transmission buffer is updated, the `NrGnbMac::DoReportBufferStatus()` intermediates with the MAC scheduler SAP to report the buffer status:

```
+0.400000282s 0 [ CellId 2, bwpId 0] NrGnbMac:DoReportBufferStatus(): [INFO ]
↳Reporting RLC buffer status update to MAC Scheduler for RNTI=2, LCID=4,
↳Transmission Queue HOL Delay=0, Transmission Queue Size=132, Retransmission Queue
↳HOL delay=0, Retransmission Queue Size=0, PDU Size=0
+0.400002262s 0 [ CellId 3, bwpId 1] NrGnbMac:DoReportBufferStatus(): [INFO ]
↳Reporting RLC buffer status update to MAC Scheduler for RNTI=1, LCID=4,
↳Transmission Queue HOL Delay=0, Transmission Queue Size=1284, Retransmission Queue
↳HOL delay=0, Retransmission Queue Size=0, PDU Size=0
```

By following these messages it is possible to track the `Transmission Queue Size`, which is filled by incoming RLC PDUs and emptied by the scheduled transmissions. By default, the `NrHelper` considers the TDMA RR as MAC Scheduler for this example, handled by `NrMacSchedulerTdmaRR`. Schedulers are extensively covered in both the [NR manual](#), Sections 2.5.2 and 2.5.3, and the [doxygen documentation](#).

The MAC layer receives scheduled allocations with the call to `NrGnbMac::DoSchedConfigIndication()`. All the scheduler parameters are packed in a `NrMacSchedSapUser::SchedConfigIndParameters` structure. Among these parameters, the HARQ Process ID is critical to fully trace (i) when an allocation is scheduled, (ii) when a frame is transmitted, and (iii) when the UE ACKs the correct reception of the frame. These are critical information that helps also evaluate the packet latency in this layer.

For instance, let's take a look at these INFO log messages, focusing the attention on finding correspondences with the HARQ Process ID:

```
+0.400062500s 0 [ CellId 2, bwpId 0] NrGnbMac:DoSchedConfigIndication(): [INFO ] New
↳scheduled data TX in DL for HARQ Process ID: 15, Var. TTI from symbol 1 to 13. 1
↳TBs of sizes [ 39 ] with MCS [ 0 ]
+0.400062500s 0 [ CellId 2, bwpId 0] NrGnbMac:DoSchedConfigIndication(): [INFO ]
↳Notifying RLC of TX opportunity for HARQ Process ID 15 LC ID 4 stream 0 size 36
↳bytes
+0.400062500s 0 [ CellId 2, bwpId 0] NrGnbMac:DoTransmitPdu(): [INFO ] Sending MAC
↳PDU to PHY Layer
...
+0.400374995s 0 [ CellId 2, bwpId 0] NrGnbMac:DoDlHarqFeedback(): [INFO ] HARQ-ACK
↳UE RNTI 2 HARQ Process ID 15 Stream ID 0
```

From the first message we can observe that a new transmission of just one stream of 39 bytes is scheduled for transmission. Consequently, from the second log message the MAC notifies this opportunity to the RLC layer. After that, the MAC PDU is prepared. After a while, the ACK from the UE arrives. This means that the frame took 312 us to be sent and acknowledged.

Here the RLC PDU is made and forwarded to the MAC layer by calling its `TransmitPdu()` method, which in this case is implemented by `NrGnbMac::DoTransmitPdu()`. At the end of packet encapsulation, the PDU is forwarded to the PHY SAP, along with its `SfnSf` and the starting symbol index as indicated in the DCI, shown in the log messages as `Var. TTI from symbol 1 to 13`. To learn more how a variable TTI works, please refer to Section 2.5.1 of the [NR manual](#).

Furthermore, observe how the TB size is larger for BWP with ID 1:

```
+0.400250000s 0 [ CellId 3, bwpId 1] NrGnbMac:DoSchedConfigIndication(): [INFO ] New
↳scheduled data TX in DL for HARQ Process ID: 15, Var. TTI from symbol 1 to 13. 1
↳TBs of sizes [ 171 ] with MCS [ 0 ]
...
+0.401499997s 0 [ CellId 3, bwpId 1] NrGnbMac:DoDlHarqFeedback(): [INFO ] HARQ-ACK
↳UE RNTI 1 HARQ Process ID 15 Stream ID 0
```

but the latency has increased to 1,249 us.

5.6.1 Packet latency

Packets do incur latency in `NrGnbMac`, as they can be left in the buffer until a successful HARQ ACK is received from the UE. Furthermore, there is a difference in latency, depending on the BWP configuration.

5.6.2 Packet drops

Packets cannot be dropped in `NrGnbMac`.

5.7 NrGnbPhy

`NrGnbPhy` tailors the `NrPhy` according to how the gNB should behave at its PHY layer. Its implementation can be found in `model/nr-gnb-phy.cc`, `model/nr-phy.cc`, and `model/nr-phy-mac-common.cc`. The PHY layer is extensively covered in Section 2.3 of the [NR manual](#). Furthermore, there is more technical description on how it works at the [doxygen of `NrGnbPhy`](#) and [doxygen of `NrPhy`](#) classes.

This simulation is configured by enabling both `NrGnbPhy` and `NrPhy` log components:

```
$ NS_LOG="NrGnbPhy=info|prefix_all:NrPhy=info|prefix_all" ./ns3 run cttc-nr-demo &>
↳out.log
```

At the start of the simulation, the PHY layer is configured. As we are in a TDD context, its pattern is set. Furthermore, according to the bandwidth available, the number of RBs per BWP is set. Finally, the channel access is requested and obtained. Given the simplicity of this simulation, the channel is granted for a long time, but take into account that it is released as soon as the gNB does not use it for data transmission. All this information can be obtained in the following log messages:

```
+0.000000000s -1 [ CellId 0, bwpId 65535] NrGnbPhy:SetTddPattern(): [INFO ] Set
↳pattern : F|F|F|F|F|F|F|F|F|F|
+0.000000000s -1 [ CellId 0, bwpId 0] NrPhy:DoUpdateRbNum(): [INFO ] Updated RbNum
↳to 16
+0.000000000s -1 [ CellId 0, bwpId 1] NrPhy:DoUpdateRbNum(): [INFO ] Updated RbNum
↳to 66
+0.000000000s -1 [ CellId 2, bwpId 0] NrPhy:DoUpdateRbNum(): [INFO ] Updated RbNum
↳to 6
+0.000000000s -1 [ CellId 2, bwpId 0] NrPhy:DoUpdateRbNum(): [INFO ] Updated RbNum
↳to 6
+0.000000000s 0 [ CellId 2, bwpId 0] NrGnbPhy:StartSlot(): [INFO ] Channel not
↳granted, request the channel
+0.000000000s 0 [ CellId 2, bwpId 0] NrGnbPhy:ChannelAccessGranted(): [INFO ]
↳Channel access granted for +9.22337e+18ns, which corresponds to 147573952589675
↳slot in which each slot is +62500ns. We lost +62500ns
```

(continues on next page)

(continued from previous page)

```
...
+0.000062500s 0 [ CellId 2, bwpId 0] NrGnbPhy:EndSlot(): [INFO ] Release the channel.
↳because we did not have any data to maintain the grant
```

The `F` in the TDD pattern means that for each TDD slot it will be possible to handle any type of frame, from DL to UL, which can be of `CTRL` or `DATA` type. More information can be found in the [NR doxygen regarding the `LteNrTddSlotType`](#).

The gNB's PHY layer receives the MAC PDU through the `NrPhy::SetMacPdu()` method. First of all, the numerology is checked to understand if the PHY is working at that numerology at the moment. If such condition is true, the PDU is appended to a `ns3::PacketBurst` object, which is an abstraction to a list of packets. Such list is mapped to contextual information, which is composed by the stream ID and starting symbol.

Indeed, it is possible to notice this INFO log message, which signals the MAC PDU entrance in this layer, along with its properties, such as `SfnSf` (used to keep track of the frame, subframe, and slot number) and starting symbol:

```
+0.400125000s 0 [ CellId 2, bwpId 0] NrPhy:SetMacPdu(): [INFO ] Adding a packet for
↳the Packet Burst of FrameNum: 40 SubFrameNum: 0 SlotNum: 4 at sym 1
```

As the timeslot order is handled by the scheduler, the transmission is led by an event loop, which calls `NrGnbPhy::StartSlot()` and queries the packet burst through `NrPhy::PushFrontSlotAllocInfo()`. Packets are finally transmitted over the air via `DlData()`, which then interacts with the spectrum through `SendDataChannels()`. Upon the same logic, packets are received from `UlData()` and `PhyDataPacketReceived()`.

From this point onwards, the `NrGnbPhy` interacts with `NrSpectrumPhy::StartTxDataFrames()`, which acts as an interface between the gNB PHY layer and the channel. `NrSpectrumPhy` acts as a state machine to know what the PHY layer (at the BWP of interest) is currently doing, from transferring/receiving data or control information or it is in idle state. Once a packet burst is given with its related set of control messages and duration of transmission, the structure `NrSpectrumSignalParameterDataFrame` is created. This object contains the aforementioned information, plus the cell identifier (`GetCellId()`) and the transmission PSD. Such information is then forwarded to the channel.

The log messages are quite verbose but in a constant pattern until there is data to transmit. Indeed, it is possible to notice these log messages once data arrive at the PHY layer:

```
+0.400187500s 0 [ CellId 2, bwpId 0] NrGnbPhy:RetrieveDciFromAllocation(): [INFO ]
↳Send DCI to RNTI 2 from sym 1 to 13
+0.400187500s 0 [ CellId 2, bwpId 0] NrGnbPhy:FillTheEvent(): [INFO ] Scheduled
↳allocation RNTI=0|DL|SYM=0|NSYM=1|TYPE=2|BWP=0|HARQP=0|RBG=[0;15] at +0ns
+0.400187500s 0 [ CellId 2, bwpId 0] NrGnbPhy:FillTheEvent(): [INFO ] Scheduled
↳allocation RNTI=2|DL|SYM=1|NSYM=12|McsStream0=0|TBsStream0=39|NdiStream0=1|RvS
tream0=0|TYPE=1|BWP=0|HARQP=15|RBG=[0;15] at +4464ns
+0.400187500s 0 [ CellId 2, bwpId 0] NrGnbPhy:FillTheEvent(): [INFO ] Scheduled
↳allocation RNTI=0|UL|SYM=13|NSYM=1|TYPE=2|BWP=0|HARQP=0|RBG=[0;15] at +58032ns
+0.400191964s 0 [ CellId 2, bwpId 0] NrGnbPhy:DlData(): [INFO ] ENB TXing DL DATA
↳frame FrameNum: 40 SubFrameNum: 0 SlotNum: 3 symbols 1-12 start +4.00192e+08ns end
↳+4.00246e+08ns
+0.400250000s 0 [ CellId 3, bwpId 1] NrGnbPhy:EndSlot(): [INFO ] Release the channel
↳because we did not have any data to maintain the grant
```

Here the `NrGnbPhy::FillTheEvent()` scheduled the Variable TTI in the simulator as events. The first scheduled information sends one symbol of control information in DL to the UE; the second one sends is the data TB, whereas the third one is to receive in UL the ACK from the UE.

5.7.1 Packet latency

At this layer each packet will get latency based on different features that characterize NR. For instance, the processing times affect packet latency and it is parametrized by `L1L2CtrlLatency` in timeslots. This parameter is hardcoded to 2, which indicates that the allocation requires two timeslots before seeing the packet going in the air. Furthermore, latency may be impacted by NR processing delays. Specifically, setting `NODelay` and consequently `K0` to a value greater than 0 will result in additional latency. For further information, please refer to Section 2.5.6 of the [NR manual](#).

It can be noticed that the packet gained 0.19 ms of latency by reaching the time scheduled for transmission and the total transmission of the data of interest requested 54 us.

5.7.2 Packet drops

Packets cannot be dropped in the `NrGnbPhy`, but they may be dropped in `NrSpectrumPhy` upon their reception if they are found to be corrupted. Whether the packet is corrupted is evaluated in `NrSpectrumPhy` by calling the function `NrErrorModel::GetTbDecodificationStats()` which takes into account the HARQ history (used for the HARQ model being configured for this simulation, i.e., Chase Combining or Incremental Redundancy.)

5.8 NrUePhy

The other end of the channel, which in this case is the UE, is implemented by `NrUePhy`, available at `model/nr/nr-ue-phy.cc`. The log components can be enabled and correlated to that of the gNB to evaluate what happens at the PHY layer:

```
$ NS_LOG="NrGnbPhy:NrUePhy" ./ns3 run cttc-nr-demo &> out.log
```

The following log excerpt can be analyzed:

```
+0.000000000s 1 [ CellId 2, bwpId 0] NrUePhy:StartEventLoop(): [INFO ] PHY starting.
↳Configuration:
    TxPower: 2 dBm
    NoiseFigure: 5
    TbDecodeLatency: 100 us
    Numerology: 4
    SymbolsPerSlot: 14
    Pattern: F|F|F|F|F|F|F|F|F|F|
Attached to physical channel:
    Channel bandwidth: 18000000 Hz
    Channel central freq: 2.8e+10 Hz
    Num. RB: 6
...
+0.400191964s 1 [ CellId 2, bwpId 0] NrUePhy:DlData(): [INFO ] UE2 HARQ ID 15 stream
↳0 RXing DL DATA frame for symbols 1-12 num of rbg assigned: 16. RX will take place
↳for +53568ns
+0.400245531s 1 [ CellId 2, bwpId 0] NrUePhy:GenerateDlCqiReport(): [INFO ] Stream 0
↳WB CQI 15 avrg MCS 28 avrg SINR (dB) 66.8459
+0.400245531s 1 [ CellId 2, bwpId 0] NrUePhy:NotifyDlHarqFeedback(): [INFO ] HARQ
↳Feedback for ID 15 Stream 0
...
+0.400767857s 2 [ CellId 3, bwpId 1] NrUePhy:DlData(): [INFO ] UE1 HARQ ID 19 stream
↳0 RXing DL DATA frame for symbols 1-12 num of rbg assigned: 66. RX will take place
↳for +214284ns
...
```

(continues on next page)

(continued from previous page)

```
+0.400982140s 2 [ CellId 3, bwpId 1] NrUePhy:GenerateDlCqiReport(): [INFO ] Stream 0_
↳WB CQI 15 avrg MCS 28 avrg SINR (dB) 65.9225
+0.400982140s 2 [ CellId 3, bwpId 1] NrUePhy:NotifyDlHarqFeedback(): [INFO ] HARQ_
↳Feedback for ID 15 Stream 0
```

At the start of the simulation, an instance of `NrUePhy` is initialized for each UE and BWP ID. So in total we get 4 messages like this. It is clear how PHY is configured with the listed parameters, which can be useful to track and understand how traffic behaves at further points in the log file, by just looking at the message prefix `[CellId 2, bwpId 0]`.

It can be also observed that data is RX by different UE PHYs, with different BWP ID. While the first one is receiving the traffic for the low latency application, the second one is receiving traffic for the high-quality voice. It is clear the difference in the number of RBs assigned and delay needed for the data transfer.

After a frame is received completely, i.e., the packet burst is copied from the transmitter PHY to the receiver one, the corresponding CQI report is computed. For this reason `NrUePhy::GenerateDlCqiReport()` prints a log message that reports useful statistics for the last frame received, especially the average MCS and SINR. The computed SINR can be also tracked through the `DlDataSinr` trace.

Finally, the HARQ feedback can be tracked through `NrUePhy::NotifyDlHarqFeedback()` to fully understand its lifecycle after frame reception.

Even though it is possible to track the lifecycle of the packet in `NrUePhy`, the `NrSpectrumPhy` contains all the actual logic that allows the evaluation of the TB SINR and TBLER and finally decides if the TB gets corrupted during the reception. Such decision can be seen in `model/nr-spectrum-phy.cc` at `NrSpectrumPhy::EndRxData()` method. The TBLER is evaluated by the `NrLteMiErrorModel`. If the TBLER is less than the realization of a `UniformRandomVariable` between 0 and 1, the TB is marked as corrupted. This implies that a NACK is sent as HARQ status.

TB reception can be observed by enabling the corresponding log component and filter by `EndRxData`:

```
$ NS_LOG="NrSpectrumPhy=info|prefix_all" ./ns3 run cttc-nr-demo | grep EndRxData >_
↳output.log
```

obtaining log messages such as this:

```
+0.400245531s 1 NrSpectrumPhy:EndRxData(): [INFO ] Finishing RX, sinrAvg=4.83716e+06_
↳sinrMin=3.60456e+06 SinrAvg (dB) 66.8459
```

5.8.1 Packet latency

The same logic of `NrGnbPhy` applies, as the frame requires the given duration for its transfer from the gNB to the UE.

5.8.2 Packet drops

Upon the same logic of `NrGnbPhy` and the corresponding `NrSpectrumPhy`, the TBs dropped by `NrSpectrumPhy` on UE side cause a HARQ feedback at this layer in order to repeat the transmission.

5.9 NrUeMac

The NrUeMac is implemented in `model/nr-ue-mac.cc`. The NrUeMac receives notification on the availability of new MAC frames from PHY layer through its `DoReceivePhyPdu()` callback. Indeed, if we enable the component log like the following

```
$ NS_LOG="NrUeMac=info|prefix_all" ./ns3 run cttc-nr-demo
```

we get the following information:

```
+0.016316963s 1 [ CellId 2, bwpId 0, rnti 0] NrUeMac:DoReceiveControlMessage():
↳[INFO ] Received RAR in slot FrameNum: 1 SubFrameNum: 6 SlotNum: 5
+0.016316963s 2 [ CellId 2, bwpId 0, rnti 0] NrUeMac:DoReceiveControlMessage():
↳[INFO ] Received RAR in slot FrameNum: 1 SubFrameNum: 6 SlotNum: 5
+0.400345531s 1 [ CellId 2, bwpId 0, rnti 2] NrUeMac:DoReceivePhyPdu(): [INFO ]
↳Received PHY PDU from LCID 4 of size 81 bytes.
...
+0.400671427s 1 [ CellId 2, bwpId 0, rnti 2] NrUeMac:DoReceivePhyPdu(): [INFO ]
↳Received PHY PDU from LCID 4 of size 185 bytes.
...
+0.401082140s 2 [ CellId 3, bwpId 1, rnti 1] NrUeMac:DoReceivePhyPdu(): [INFO ]
↳Received PHY PDU from LCID 4 of size 345 bytes.
...
+0.402582140s 2 [ CellId 3, bwpId 1, rnti 1] NrUeMac:DoReceivePhyPdu(): [INFO ]
↳Received PHY PDU from LCID 4 of size 12065 bytes.
```

It is evident the amount of data being received by both BWPs. Finally, the MAC SDU is prepared by just removing the MAC header from its PDU and forwarded to the LTE RLC layer.

5.9.1 Packet latency

We can cross-check what has been transmitted by NrGnbMac and what has been received here. An example by enabling both log components is provided hereby

```
$ NS_LOG="NrGnbMac=info|prefix_all:NrUeMac=info|prefix_all" ./ns3 run cttc-nr-demo
```

If we filter the output only to track the first 3 TBs sent by each BWP, we obtain the following log messages:

```
+0.400062500s 0 [ CellId 2, bwpId 0] NrGnbMac:DoSchedConfigIndication(): [INFO ]
↳Notifying RLC of TX opportunity for HARQ Process ID 15 LC ID 4 stream 0 size 39
↳bytes
+0.400125000s 0 [ CellId 2, bwpId 0] NrGnbMac:DoSchedConfigIndication(): [INFO ]
↳Notifying RLC of TX opportunity for HARQ Process ID 14 LC ID 4 stream 0 size 39
↳bytes
+0.400187500s 0 [ CellId 2, bwpId 0] NrGnbMac:DoSchedConfigIndication(): [INFO ]
↳Notifying RLC of TX opportunity for HARQ Process ID 13 LC ID 4 stream 0 size 39
↳bytes
+0.400250000s 0 [ CellId 3, bwpId 1] NrGnbMac:DoSchedConfigIndication(): [INFO ]
↳Notifying RLC of TX opportunity for HARQ Process ID 15 LC ID 4 stream 0 size 171
↳bytes
+0.400345531s 1 [ CellId 2, bwpId 0, rnti 2] NrUeMac:DoReceivePhyPdu(): [INFO ]
↳Received PHY PDU from LCID 4 of size 39 bytes.
+0.400408031s 1 [ CellId 2, bwpId 0, rnti 2] NrUeMac:DoReceivePhyPdu(): [INFO ]
↳Received PHY PDU from LCID 4 of size 39 bytes.
+0.400470531s 1 [ CellId 2, bwpId 0, rnti 2] NrUeMac:DoReceivePhyPdu(): [INFO ]
↳Received PHY PDU from LCID 4 of size 39 bytes.
```

(continues on next page)

(continued from previous page)

```
+0.400500000s 0 [ CellId 3, bwpId 1] NrGnbMac:DoSchedConfigIndication(): [INFO ]_
↳Notifying RLC of TX opportunity for HARQ Process ID 14 LC ID 4 stream 0 size 171_
↳bytes
+0.400750000s 0 [ CellId 3, bwpId 1] NrGnbMac:DoSchedConfigIndication(): [INFO ]_
↳Notifying RLC of TX opportunity for HARQ Process ID 13 LC ID 4 stream 0 size 171_
↳bytes
+0.401082140s 2 [ CellId 3, bwpId 1, rnti 1] NrUeMac:DoReceivePhyPdu(): [INFO ]_
↳Received PHY PDU from LCID 4 of size 171 bytes.
+0.401332140s 2 [ CellId 3, bwpId 1, rnti 1] NrUeMac:DoReceivePhyPdu(): [INFO ]_
↳Received PHY PDU from LCID 4 of size 171 bytes.
+0.401582140s 2 [ CellId 3, bwpId 1, rnti 1] NrUeMac:DoReceivePhyPdu(): [INFO ]_
↳Received PHY PDU from LCID 4 of size 171 bytes.
```

It is possible to observe that in ~1,519.64 us, on the one hand, 243 bytes were correctly sent to the UE for the first BWP, whereas 513 bytes were sent for the second BWP. On the other hand, the first BWP took only ~283 us on average to transmit these TBs, while the second BWP took ~832 us. This aspect greatly highlights the trade-off that is taking place upon choosing the different BWPs.

5.9.2 Packet drops

Packets may be dropped by `NrUeMac` if the RNTI does not match with the expected one.

5.10 LteRlcUm - UE Side

When the RLC PDU is received at `LteRlcUm::DoReceivePdu()`, its header is removed. Its packet sequence number is checked to verify if (i) the data need reordering, (ii) the data is a duplicate of a previous one, and (iii) it has been received outside the reordering window. To fully understand this logic, please refer to [Section 5.1.2.2 of the LTE RLC protocol specification](#).

Once the SDU is reassembled correctly from the received PDUs, `LteRlcUm::ReassembleAndDeliver()` sends the packet up to the PDCP layer.

5.10.1 Packet latency

The delay at RLC layer can be tracked with the `RxPdu` trace. At the same time, it is possible to enable the `LteRlcUm` log component and filter by the `LteRlcUm::DoReceivePdu()` messages. It can be noticed that, for the low latency voice, the packets have an average latency of 234 us, whereas the high quality voice has 681 us. This data can be easily extracted with the following two commands:

```
$ grep '1 LteRlcUm:DoReceivePdu' output.log | \
  grep -oP '[0-9]+(?:=ns)' | \
  awk '{ total += $1; count ++ } END { print total/count/1e3 }'
233.976
$ grep '2 LteRlcUm:DoReceivePdu' output.log | \
  grep -oP '[0-9]+(?:=ns)' | \
  awk '{ total += $1; count ++ } END { print total/count/1e3 }'
680.866
```

At the same time, the aforementioned commands can be easily adapted to report the average bytes transmitted per packet. For the former, it is 132 bytes, while for the latter, 3210.

5.10.2 Packet drops

The RLC PDU is dropped if it was already received or its sequence number falls outside the reordering window, as per Section 5.1.2.2 of the [LTE RLC protocol specification](#).

5.11 LtePdcP - UE Side

Like what we have seen at RLC layer, the incoming PDCP PDUs are received at `LtePdcP::DoReceivePdu()` method. The header is removed and the packet is then transmitted to the `LteUeRrc`, which straightly pass the packet to `EpcUeNas` and then `NrUeNetDevice`.

5.11.1 Packet latency

Latency can be evaluated with the same method used for Section 5.10.1. `LteUeRrc` and `EpcUeNas` acts transparently and do not add any additional latency.

5.11.2 Packet drops

Packets cannot be dropped at this layer, neither at RRC and UE NAS.

5.12 NrUeNetDevice

The packet is received by the `EpcUeNas` to `NrUeNetDevice::DoRecvData()` which does an up-callback to `NrUeNetDevice::Receive()`. The implementation of such callback is the same for both UE and gNB, and thus it can be found in `NrNetDevice`, i.e., `contrib/nr/model/nr-net-device.cc`.

If the `NrNetDevice` log component is enabled, the simulation produces the following message:

```
+0.400533031s 1 NrNetDevice:Receive(): [INFO ] Received 128 bytes on
↳ 00:00:00:00:00:09. IPv4 packet from 1.0.0.2 to 7.0.0.2
```

Finally, such packet is then available for the upper layers to consume, i.e., transport and application ones, thanks to the `m_rxCallback` up-call.

5.12.1 Packet latency

Packets cannot incur latency in the `NrUeNetDevice`.

5.12.2 Packet drops

Packets cannot be dropped in the `NrUeNetDevice`.

CONCLUSIONS

In this tutorial, we have shown how the program `cttc-nr-demo` works, which is part of the `ns-3 nr` module. As we have provided a detailed description of what happens layer-by-layer, the lifecycle of downlink packets has been shown for two applications, one that focuses on low latency communications, and the other which focuses on high quality voice. These two flows were characterized by different type of BWPs, which were handled differently by the NR RAN. At each layer, the packet latency and the possibility of their drops were discussed.

We hope that this tutorial was useful in getting started with the `nr` module and better understand how this implementation of the NR standard works. If you feel the need for any clarification or you want to update this document, please contact us!