Leveraging Overshadowing for Time-Delay Attacks in 4G/5G Cellular Networks: An Empirical Assessment

Virgil Hamici-Aubert, Julien Saint-Martin, Renzo E. Navas, Georgios Z Papadopoulos, Guillaume Doyen, Xavier Lagrange

Leveraging Overshadowing for Time-Delay Attacks in 4G/5G Cellular Networks: An Empirical Assessment

Virgil Hamici-Aubert
virgil.hamici-aubert@imt-atlantique.fr
IMT Atlantique
IRISA, UMR CNRS 6074
Rennes, France

Julien Saint-Martin
julien.saint-martin@imt-atlantique.fr
IMT Atlantique
IRISA, UMR CNRS 6074
Rennes, France

Renzo E. Navas
renzo.navas@imt-atlantique.fr
IMT Atlantique
IRISA, UMR CNRS 6074
Rennes, France

Georgios Z. Papadopoulos
georgios.papadopoulos@imt-atlantique.fr
IMT Atlantique
IRISA, UMR CNRS 6074
Rennes, France

Guillaume Doyen
guillaume.doyen@imt-atlantique.fr
IMT Atlantique
IRISA, UMR CNRS 6074
Rennes, France

Xavier Lagrange
xavier.lagrange@imt-atlantique.fr
IMT Atlantique
IRISA, UMR CNRS 6074
Rennes, France

ABSTRACT

Ensuring both reliable and low-latency communications over 4G or 5G Radio Access Network (RAN) is a key feature for services such as smart power grids and the metaverse. However, the lack of appropriate security mechanisms at the lower-layer protocols of the RAN–a heritage from 4G networks–opens up vulnerabilities that can be exploited to conduct stealthy Reduction-of-Quality attacks against the latency guarantees. This paper presents an empirical assessment of a proposed time-delay attack that leverages overshadowing to exploit the reliability mechanisms of the Radio Link Control (RLC) in Acknowledged Mode. By injecting falsified RLC Negative Acknowledgements, an attacker can maliciously trigger retransmissions at the victim User Equipment (UE), degrading the uplink latency of application flows. Extensive experimental evaluations on open-source and commercial off-the-shelf UEs demonstrate the attack's effectiveness in increasing latency, network load, and buffer occupancy. The attack impact is quantified by varying the bitrate representing different applications and the number of injected negative acknowledgments controlling the attack intensity. This work studies a realistic threat against the latency quality of service in 4G/5G RANs and highlights the urgent need to revisit protocol security at the lower-RAN layers for 5G (and beyond) networks.

CCS CONCEPTS

• **Security and privacy → Mobile and wireless security**; *Denial-of-service attacks*; Security protocols.

KEYWORDS

Radio Access Network, Overshadowing, Man on the Side, Deny of Service, Reduction of Quality, Latency, Time-delay

ARES 2024, July 30-August 2, 2024, Vienna, Austria

ACM ISBN 979-8-4007-1718-5/24/07
https://doi.org/10.1145/3664476.3670891

ACM Reference Format:
Virgil Hamici-Aubert, Julien Saint-Martin, Renzo E. Navas, Georgios Z. Papadopoulos, Guillaume Doyen, and Xavier Lagrange. 2024. Leveraging Overshadowing for Time-Delay Attacks in 4G/5G Cellular Networks: An Empirical Assessment. In *The 19th International Conference on Availability, Reliability and Security (ARES 2024), July 30-August 2, 2024, Vienna, Austria.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3664476.3670891

1 INTRODUCTION

The fifth-generation (5G) technology standard for cellular networks brings many improvements over its predecessor, including the reduction of latency on the Radio Access Network (RAN). Mobile network usages are evolving thanks to these improvements. For instance, the Metaverse [18, 23], which can be defined as a convergence of physical, augmented, and virtual reality in a shared online space, stands for one of the cornerstone emerging usages. Metaverse-based applications induce new traffic flow properties [9, 17], where latency is essential to the required Quality of Service (QoS). Besides, in several recent applications, including smart power grids and advanced manufacturing, latency must be guaranteed and protected against threats such as time-delay attacks [6, 14] which stands for a novel and stealthier form of Denial of Service (DoS).

5G networks already exhibit a certain level of protection on the RAN, including confidentiality and integrity. Those mechanisms are applied in the Packet Data Convergence Protocol (PDCP) protocol at the Layer 2 (L2) part of the 5G New Radio (NR) protocol stack. However, some vulnerabilities still exist [4]. Indeed, the PDCP mechanisms protect the upper RAN layers but leave those below out of comprehensive security mechanisms. Some research studies exploit this vulnerability through the use of a new attack vector named Man-on-the-Side (MotS), based on radio overshadowing [15, 21, 25]. Its principle consists in sending a stronger signal than the legitimate one to overwrite it. In that way, an attacker can threaten the User Plane (UP) without triggering network service interruption and stealthily conduct a DoS on the RAN.

In this paper, we propose and implement an attack which introduces a delay on the Radio Link Control (RLC) layer to degrade the

Uplink (UL) application traffic's latency. The attack principle consists in sending falsified RLC Negative Acknowledgement (NACK) inside a false RLC `STATUS report` by overshadowing, to trigger retransmissions at a victim User Equipment (UE). In compliance with the 3rd Generation Partnership Project (3GPP) definition of this protocol, the attack is seen by the UE as a legitimate message from the Base Station (BS, 5G Node Base station, gNB). Our attack mechanism stands for a variant of the overshadowing attack proposed on Cellular Internet of Things (IoT) networks [21] which drains IoT devices' battery by hijacking the RLC layer reliability mechanism—but, our variant allows to relax some hypothesis, making it plausible in a more realistic scenario.

We evaluate the impact on latency and other performance metrics of our proposed overshadowing-based attack as a function of two main input parameters: the bitrate, representing different application flows, and the number of injected NACKs, representing the stealthiness of the attacker. We performed an experimental evaluation campaign using an open-source-based and a Commercial Off-The-Shelf (COTS) UE and we demonstrated the effectiveness of the attack in substantially degrading the latency guarantee.

The contributions of this paper are the following:

(1) We propose a variant of the overshadowing attack proposed in [21] to build a time-delay attack. In our proposal, the attacker can control the number of injected NACKs within certain limits to generate a delay.
(2) We exhibit the behaviour of the RLC layer at the protocol level, with and without attack, over an experimental open-source-based UE.
(3) We conduct an empirical analysis of the UL application latency, UL RAN load, and the UL RLC buffer, to quantify the attack's impact on an experimental open-source-based and a COTS UE.

The rest of the paper is organized as follows. Section 2 provides the necessary 5G's technical background while Section 3 presents 5G protocol stack's vulnerabilities and related work on DoS on 5G. Section 4 describes the attack and Section 5 details the methodology of the experiment. Then, Section 6 presents the experimental scenarios and performance evaluation results. Finally, Section 7 concludes this paper and provides future perspectives.

2 TECHNICAL BACKGROUND

In this section, we present an overview of the 5G radio interface. We focus on the components and the mechanisms affected by the vulnerabilities we base our proposal on.

2.1 5G New Radio

The RAN is the air interface entry point of a UE. It is composed of numerous Base Station (BS). On the radio interface, each UE is identified with a Radio Network Temporary Identifier (RNTI). At every slot, the BS allocates radio resources to the UE. The allocation message is called Downlink Control Information (DCI). A DCI includes the number of resource blocks and the Modulation and Coding Scheme (MCS) to use. The UE can thus compute the Transport Block Size (TBS), which denotes the quantity of bytes allocated.

2.2 Air Interface Protocol Stack

Figure 1 shows the protocol stack of the 5G NR [3] RAN and we describe it in the following sections.

2.2.1 Physical Layer. The bottom part of the stack is the Physical Layer, which includes the physical channels, and especially the parameters broadcasted by a BS for the user's access to be synchronized in frequency and in time.

2.2.2 Layer 2. The second part is the L2, composed by the Medium Access Control (MAC), RLC, PDCP and Service Data Adaptation Protocol (SDAP) layers. The MAC layer ensures some reliability with the Hybrid Automatic Repeat reQuest (HARQ) mechanism. The RLC layer, which can exhibit several instances, is in charge of segmentation. It can also include an Automatic Repeat reQuest (ARQ) mechanism when the Packet Error Rate (PER) provided by the MAC one is not small enough to provide the QoS level required by the upper application. This is achieved with the Acknowledged Mode (AM), which is the most frequently used mode. The Unacknowledged Mode (UM) can also be leveraged but is restricted to Voice over IP (VoIP) essentially.

The MAC and RLC layers are very similar in 4G and 5G. The main difference lies in the concatenation and re-sequencing at the receiver side, which have been dropped in 5G to improve latency. The MAC layer permits the UE to periodically send a Buffer Status Report (BSR) that notifies the BS about the amount of bytes in the UE UL buffer. This buffer contains the RLC Service Data Unit (SDU) delivered by the PDCP entity and the Packet Data Unit (PDU) already transmitted but not explicitly acknowledged. Note that the buffer thus contains PDU that were negatively acknowledged and that wait for re-transmission [1].

The PDCP layer includes ciphering to ensure confidentiality of the Control Plane (CP) and UP. The integrity control is enforced in the CP and optional in the UP. Confidentiality and integrity control are provided to layers above the PDCP one. Finally, the SDAP layer includes a tag to manage different QoS levels.

2.2.3 Radio Resource Control and Non-Access Stratum. The last part of the 5G NR consists in the Radio Resource Control (RRC) and Non-Access Stratum (NAS) layers in the CP. The RRC layer exchanges radio management messages between a UE and a BS, with the Access Stratum (AS). The NAS layer is in charge of exchanging with other control entities of the network for mobility and authentication purposes, for instance. NAS packets are just forwarded by the BS.

2.3 Reliability on RAN

The RAN protocol layers have two types of reliability mechanisms. The first is provided by the MAC layer with HARQ, which stands for a parallel send-and-wait mechanism. The second is provided by the RLC AM, which is a selective-repeat protocol. The RLC layer depends on timers and configurable parameters that are defined by the 3GPP [1].

More specifically, the sender generally transmits several RLC PDU successively without waiting for an acknowledgement. It can ask the receiver to know which RLC PDU have been correctly received to selectively re-transmit those RLC PDU that have been lost. This is achieved with a `Polling bit` included in the RLC header. The receiver answers with a `STATUS report` frame that carries the

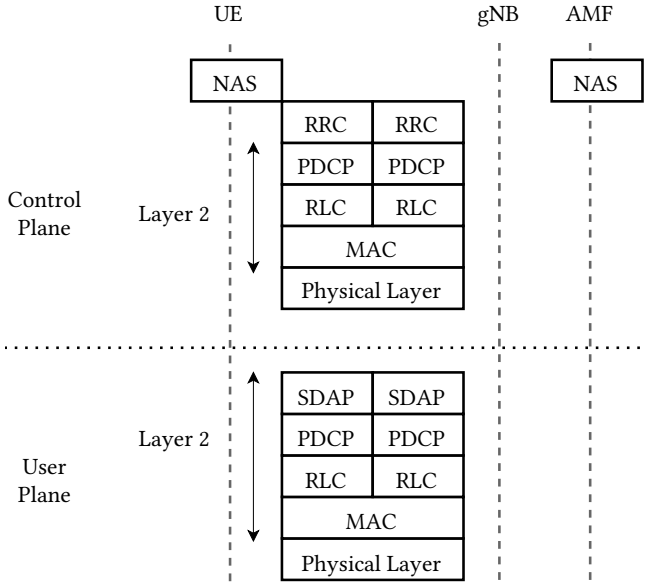

Figure 1: Protocol stack of 5G NR RAN

Sequence Number (SN) of the next expected RLC PDU denoted by `ACK_SN` and the list of all PDU that were not correctly received (i.e., `NACKs` list). The sender maintains a state variable `Tx_Next_Ack` and it updates its value when receiving a `ACK_SN` [1]. The only NACK values considered by the receiver are those between the previously received `ACK_SN` (`TX_Next_Ack`) and the highest SN already transmitted to the MAC layer [1]. The `Polling bit` is triggered when a given number of RLC PDU (`pollPDU`) or bytes (`pollByte`) sent has been reached or when a timer (`t-PollRetransmit`) expires. Indeed, each time the sender activates the `Polling bit`, it triggers a timer. Each of those parameters is cancelled as soon as one of them triggers the use of a `Polling bit`. At `Polling bit` reception, a timer is triggered (`t-StatusProhibit`), and any new `Polling bit` is ignored until it ends. At the end of `t-StatusProhibit`, a `STATUS report` is sent even if a `Polling bit` has not been received. Consequently, the receiver should process the re-transmissions even if the `Polling bit` has not been set. Figure 2 presents an example of a RLC AM communication. The PDU SN 1 and 3 are lost. The reception of the `STATUS report` triggers their re-transmission.

3 RELATED WORK

The placement of the PDCP layer in the 5G NR RAN induces some security issues. The confidentiality and reliability cannot be applied to the RLC and MAC layers. Consequently, the reliability on RAN can not be ensured. In this section, we describe how these vulnerabilities have be exploited in the literature and we provide an overview of DoS attacks in 5G networks.

3.1 Vulnerabilities

The attack presented in this paper uses two vulnerabilities. The first lies in the lack of cryptography, in particular confidentiality and integrity, for some RLC messages, , and the second in the reliability

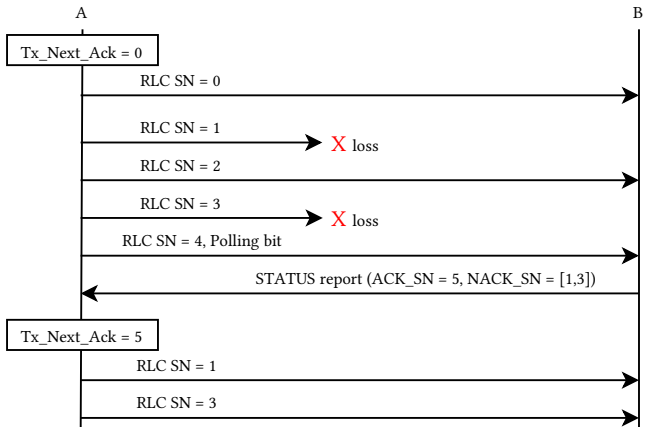

Figure 2: Example of an RLC AM re-transmission triggered by receipt of a STATUS report

definition of RLC AM. These flaws allow to exploit the reliability mechanisms of RLC and force re-transmissions.

3.1.1 Cryptography. As shown in Figure 1, the location of the PDCP layer allows the application of cryptographic mechanisms solely on the upper layers of PDCP. The RLC and MAC layers implement neither confidentiality nor integrity protection. As such, an attacker can eavesdrop their messages and modify them. The RLC layer is consequenlty vulnerable to smart jamming [4, 8] which consists for an attacker to tamper a selected part of the legitimate signal. This attack is especially called *overshadowing* [7, 15, 25]. It is made possible thanks to the capture effect which implies that a UE receiving numerous signals on the same frequency listens to the strongest [24]. In addition, the attacker can eavesdrop the Downlink (DL) and the UL to recover the victims' `RNTI` and the operator's cell parameters [16].

3.1.2 Reliability. The MAC layer is the carrier of any RLC PDU. The RLC layer does not know if the MAC one has acknowledged some of its packets and keeps the concerned SN unacknowledged. In addition, several PDU can be sent between two `STATUS report`. As a consequence, an attacker can forge NACK list for a range of RLC SN even if the packets are already acknowledged by the MAC layer.

3.2 DoS Attacks on 5G

Given the scope of our Reduction of Quality (RoQ) attack on the RAN latency, we have studied the different methods an attack can follow to implement it. First, a False Base Station (FBS) can spoof a legitimate BS, copying its parameters. Second, a rogue UE attacks an entity of the network, impersonates a UE or sends any possible value in any field in an existing message by fuzzing. Third, MotS attack allows an attacker to overwrite a legitimate signal using overshadowing. Besides, some attacks target 4G or 5G instances where the vulnerabilities are errors in the implementation. Finally, some others target problems or vulnerabilities in the specification, which impact more implementations and are more challenging to

patch. In the following, we provide a small literature review on those attacks.

3.2.1 Physical Layer. The most popular DoS attack on the physical layer consists to eavesdrop on the broadcast of cell and system parameters to impersonate a BS and hijack a UE from network regular services by triggering a handover [5]. This attack constitutes the first step of other chained attacks on the layers above. Other attacks use the lack of cryptography to overshadow broadcast parameters from the BS, e.g., the Signaling Storm [25], which forces the UE to send a useless location update to the BS.

3.2.2 Layer 2. The attacks on L2 hijack the allocation of resources and reliability of the MAC and RLC layers on UP. They trigger useless re-transmission, send false acknowledgement, and overshadow the UE BSR to drain IoT batteries, break reliability, and drain radio resources [21], respectively. Other attacks use fuzzing (e.g., by sending malformed messages) trying to detect implementation errors, eventually triggering a crash from the BS [10].

3.2.3 RRC and NAS. Attacks on RRC use fuzzing to find integrity problems on failure messages with a FBS or to send malformed messages using a rogue UE, which leads to the disconnection of the UE or the BS [26], or the crash of the BS [10], respectively. Other attacks use a FBS to target RRC and NAS integrity specification problem to send malicious messages that affect the UE network operation [12], which in turn downgrade or deny the UE services [20] or detach the UE from the network [11]. NAS is also the target of replay on `authentication_request` from the network and malformed messages to the network, which can lead to a new authentication procedure from the UE [13] and a network entity crash, respectively. Even if some attacks are restricted to 4G, other research papers still consider them also valid on 5G [12, 19].

In conclusion, we do not identify in the literature time-delay attacks on 5G whose purpose is to stealthily degrade an application flow. Indeed, a FBS and a rogue UE disconnect the UE from the network and crash the network entities, respectively. The attacks introduced on the physical layer, RRC and NAS, as well as the attacks from [10] on L2, lead to the interruption of services which can be detected through standard monitoring. In addition, even if the attacks on NAS are launched at RAN, given the definition of NAS, those attacks target the network's core. Consequenlty, we selected a MotS strategy because it keeps UE connected to the network. We adapted the attacks on L2 from [21] given the possibility of hijacking protocols by the UP and to target the specification vulnerabilities. Those choices constitute the opportunity to conduct a RoQ to degrade services without interrupting the network operating, making the attacker stealthy [25] and more challenging to detect.

4 ATTACK OVERVIEW

In this section, we introduce our attacker model and describe the attack and its implementation.

4.1 Attacker Model

We consider that the attacker uses a Software Defined Radio (SDR) and he/she is under the coverage of a cell. He/she is synchronized in frequency and time, and knows the victim's `RNTI`. He/she can listen from the DL and send signals using overshadowing on the UL. Finally, he/she can not break the messages' ciphering and can not bypass the messages' integrity protection.

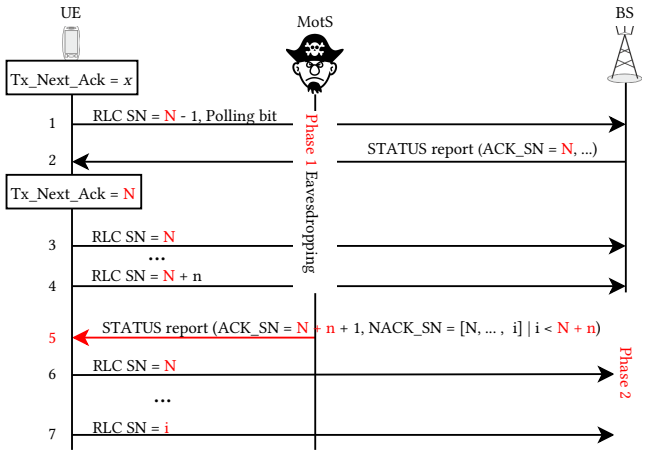

Figure 3: Overshadowing attack scenario

4.2 Attack Description

The objective of our attack consists in generating additional messages on the RAN between the victim UE and the BS, which can lead to an application latency increase. The attack operation is grounded by [21] which is deployed over IoT. The principle of [21] relies on avoiding the use of Discontinuous Reception (DRX), which allows a device to save battery power by managing listening cycles. To that aim, the attacker sends falsified RLC NACK to keep the devices in listening mode, draining their battery. By leveraging the same principle, our attack aims at triggering RLC re-transmissions, thus inducing latency increases for data packets.

Figure 3 describes the attack during a communication between a UE and a BS with a flow that uses RLC AM. During Phase 1, the attacker eavesdrops the UL and DL communication. The UE sends the $N-1-th$ SN and set the `Polling bit` (message 1). This triggers a `STATUS report`, brings the `ACK_SN` (N) and gives the `TX_Next_Ack` information to the attacker. Some additional SN are sent by the UE (messages 3 to 4). In Phase 2 the attacker knows a range of possible SN for NACK. The valid range starts from the previously observed `ACK_SN` in message 2 (`TX_Next_Ack`), N, to the $i-th$ SN. The attacker forges a valid DCI to trigger the UE decoding message and the `STATUS report` (message 5). The reception of this `STATUS report` triggers useless re-transmissions in the range of messages tagged as NACK by the attacker (messages 6 and 7).

One can notice that for an `ACK_SN` set to $N+n$ we choose to limit the highest `NACKS_SN` to $N+n-1$, because in a typical scenario the `STATUS report` is sent after the reception of a correct frame. Thus, we define the range of SN that an attacker can negatively acknowledge in a `STATUS report` as $[Tx\_Next\_ACk, ACK\_SN-1]$. The attacker can send a `STATUS report` before a `Polling bit` is set on the UE side or, in the worst case, when the valid `STATUS report` is sent by the BS.

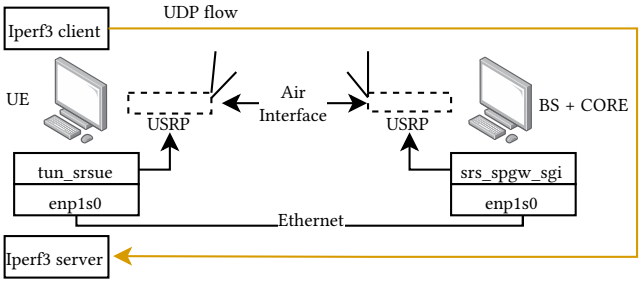

Figure 4: Experimental testbed used to measure the UL latency of User Datagram Protocol (UDP) flows

5 EXPERIMENTAL SETUP AND METRICS

In order to assess the attack's effect and validity, we implemented it as an emulation on an SDR-based platform. The 4G and 5G RAN protocol stacks are fundamentally the same. The differences lie in the re-sequencing and the concatenation done by RLC in 4G, and the absence of integrity control on UP in 4G. Due to the maturity of the 4G network's existing implementation as compared to 5G, we chose to emulate our attack on a 4G network. In this section, we describe our experimental testbed and the measurement methodology we considered to assess the attack impact.

5.1 Experimental Testbed

5.1.1 Implementation. The implementation we selected is srsLTE[1], with their UE, BS (evolved Node Base station, eNB), and core (Evolved Packet Core (EPC)). The attacker emulation is achieved directly in the code of the stack inside the RLC AM Long Term Evolution (LTE) part, in the function called at a `Polling bit` reception. The first operation is to check if the PDU will be sent to the UP. In this case, the subsequent operation is to fill the `STATUS report` with the amount of falsified NACK requested in the execution within the correct range. The difference with the scenario is that the attack is launched at the `Polling bit`'s reception. Then, the correct range of NACK is located between the previous `ACK_SN` (`TX_Next_Ack`) and the current `ACK_SN` minus one. The last operation is to keep the current value of the `ACK_SN` in memory. In addition, some modifications in the main function of the evolved Node Base station (eNB) allows to activate and manage the emulation through some options. Those options enable the control of the falsified NACK limit and a counter to avoid the attack at the start of the communication.

5.1.2 Material. Figure 4 presents our experimental testbed. The UE and the mobile network (BS + CORE) infrastructure are separated on two machines. The UE is executed in a Dell Precision 3650 Tower with an Intel(R) Xeon(R) w-1270 CPU@3.40 GHz with Ubuntu 22.04.3 LTS. The BS and core are executed in a DELL Precision 5826 Tower with an Intel (R) Xeon (R) w-2245 CPU@3.90 GHz with Ubuntu 22.04.4 LTS. The radio modules are SDR Ettus USRP B210[2]. To avoid radio link interference we use two coaxial cables in place of antennas.

5.1.3 Configuration. The bandwidth is set to 5 MHz and the UL and DL use Frequency Division Duplex (FDD), thus providing 25 Physical Ressource Block (PRB) on UL and DL with a maximum MCS of 13 for the UL. The maximum TBS that can be scheduled for the UE in this configuration is 1335 bytes [2]. The theoretical limit of the throughput is $1335 \times 8 \times 1000 = 10.7$ Mbps. The limit of the UL buffer is 400k bytes. The DL and UL frequencies we considered are 2680 MHz and 2560 MHz, respectively. Finally, the RLC AM parameters `t-PollRetransmit`, `pollPDU`, `pollByte` and `t-StatusProhibit` are set to 85 ms, 128, 125 kB and 60 ms, respectively.

5.1.4 Traffic generation. We employ Iperf3.9[3] to generate the UL flows, with UDP as a transport protocol. The UE's machine executes the client and the server, the client generates the UDP flow to the server.

5.1.5 UL Latency Measurement. As shown in Figure 4, a closed loop permits to observe the traffic sent and received on the same machine. The UE sends the UDP flow through the air interface (`tun_srsue`) which is then received on the Ethernet interface (`enp1s0`). We use Tshark, a network analyzer, to tag the packets with a timestamp. The packet's latency starts from the air interface output and ends with the Ethernet arrival, which allows measuring the UL latency samples for one communication.

5.2 Metrology

Our experiments are executed in a clear box because we can access to all the metrics provided by srsLTE on the UE side including the radio link, the load on RAN, and the scheduling. In addition, we have access to all LTE protocol levels from srsLTE logs. For each test, the radio link quality is checked with the Signal-to-Noise Ratio (SNR) and Reference Signal Receive Power (RSRP). We ensure that the SNR is larger than 15 dB and the RSRP is in the [-82 dBm;-79 dBm] interval.

5.2.1 Metrics. We use several logs from the different processes of the eNB, mainly at the RLC and MAC levels. In all the following, we use i as a time index (each time, a new log value is captured, i is incremented).

Each time t_i a new RLC frame is transmitted on the radio interface, SN is increased. We denote the frame number as Tx_sn_i. Similarly, we denote the sequence number of frames that are retransmitted as rTx_sn_i. We use $Nacks\_sn_i$ to denote the SN for which the UE received the ith NACK in a `STATUS report`. At the MAC level, we denote the ith allocation occurrence for UE's UL as $UL\_allocation_i$. Finally, at radio interface, we denote the bitrate generated by the UE application on UL RAN as Ld (for load). The load is given in Mbps.

5.2.2 Indicators. Let T_i^A and T_i^E be the time at which a packet i is transmitted on the air interface and received on Ethernet, respectively. The measured latency of a packet is given by:

$$L_i = T_i^E - T_i^A \tag{1}$$

[1] srsLTE - 4G - https://www.srslte.com/4g
[2] Ettus USRP B210 - https://www.ettus.com/all-products/ub210-kit/.
[3] iperf3.9 - https://iperf.fr/iperf-doc.php

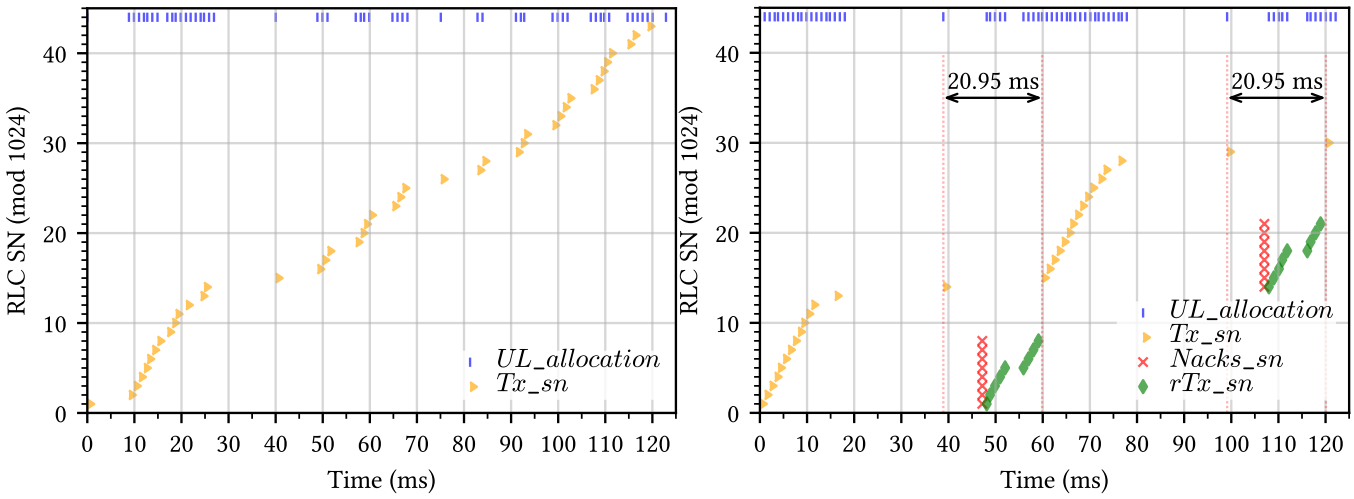

(a) Without attack **(b) Under attack with *Nacks* fixed to 8**

Figure 5: Experiment 2, protocol level view at RLC with UL transmission at 2 Mbps

Table 1: Experiment 1, Metrics and Indicators without Attack (baseline)

bitrate (Mbps)	L (ms)		Ld (Mbps)	Buffer B (bytes)	BSR_{min} (bytes)	BSR_{max} (bytes)	IR (count)		$Delay$ (ms)
	mean	stdev	mean				mean	max	mean
0.5	19	8	0.70	694	27	1817	3	14	60
1	15	7	1.5	1204	37	3995	4	19	60
2	15	6	2.8	2513	1	8787	9	31	60
4	15	7	4.7	5316	1	16507	18	40	60
8	11	1	8.4	6411	1817	19325	46	111	117

We denote the bytes quantity in a UE BSR as BSR. The average of UE BSR and thus the UE buffer size average is given by:

$$B = \frac{\sum_{i=0}^{i_{max}} BSR_i}{i_{max}} \tag{2}$$

Let $Bytes^T$ and $Bytes^R$ be the amount of bytes transmitted and re-transmitted by the UE, respectively. The re-transmission rate is given by:

$$rTx\_rate = \frac{Bytes^R}{Bytes^T} \times 100 \tag{3}$$

Let T_i^R be the time at which a `STATUS report` is sent by the BS. The delay between two consecutive `STATUS report` is given by:

$$Delay_i = T_i^R - T_{i-1}^R \tag{4}$$

The number of SN that an attacker can negatively acknowledge in one `STATUS report` message is called the Injection Range and given by:

$$IR = (ACK\_SN - Tx\_Next\_Ack - 1) \mod 1024 \tag{5}$$

5.3 Experiment Configurations and Parameters

We generate a UL communication of 60 seconds for each execution of an experimental test. The UE's MCS is dynamic. However, the BS's MCS is fixed, to counter-balance the absence of DL flow generation in our experiment. After some tests the BS's MCS is set to 25 for the DL.

We consider three parameters to form the space of our experimental campaign: the bitrate and payload size set in Iperf, and $Nacks$. We define $Nacks$ as the limit of NACK in each `STATUS report` sent by the attacker. In other words, NACK are constrained in $[Tx\_Next\_Ack\ ,\ \min(Tx\_Next\_Ack + Nacks, ACK\_SN)]$. One can notice that some preliminary tests revealed that the variation of the payload size does not impact any of the collected results. It was therefore fixed to 1024 bytes which is the power of 2 closest to the Maximum Transmission Unit (MTU).

6 EXPERIMENTAL SCENARIOS AND PERFORMANCE EVALUATION RESULTS

This section presents our experimental campaign. Detailing, for each of the five experiments, the scenario, and its results and interpretations. All of the metrics and indicators we consider are those presented in Subsection 5.2.

6.1 Results Computation

To observe the UE's flow during the attack, we truncated the communication between 10 and 50 seconds for each experimental test

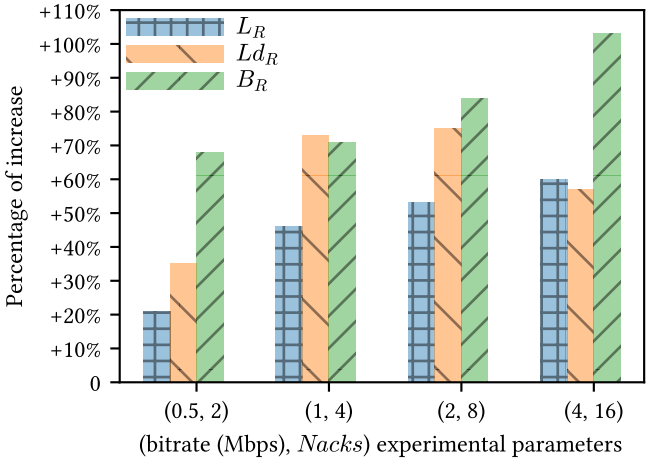

Figure 6: Experiment 3, increase of the average latency (L_R), Load (Ld_R) and Buffer size (B_R) as compared to the baseline of Experiment 1

results. The experimentation was conducted numerous times for each parameter, and observing the confidence interval ensured that we have no non-mastered random phenomena which could bias the results. However, in order to also avoid any noise of measurement, the evaluation results we present below are the mean of five to ten repetitions of each experimentation with a given set of parameters.

6.2 Experiment 1: Testbed Operation without Attack (baseline)

6.2.1 Scenario. In this first experiment, the bitrate parameter takes values from 0.5, 1, 2, 4, to 8 Mbps, whose upper bound value avoids reaching the bitrate limit of the radio interface. The goal here is to measure the regular metrics and indicators on different types of flows in a baseline scenario, allowing to compare it with the measures of the subsequent experiments under attack. We performed this experiment five times.

6.2.2 Results and Interpretations. Table 1 presents the baseline values of the *Ld* metric, and the *L*, *IR*, *B*, *BSR* and *Delay* indicators. The latency *L* and its standard deviation decrease as the bitrate increases, and we also observe a flat shape for the 1, 2 and 4 Mbps *L* mean. On the contrary, the *B*, BSR_{max}, *IR* and *Delay* increases. The *IR* values demonstrate that the attacker power possibility increases as a function of the bitrate. The BSR_{min} values denote a constant occupation of the UL buffer, which is significant for 8 Mbps. The additional bytes between the bitrate and the load implied by the application packets encapsulation increases to 2 Mbps and start to decrease at 4 Mbps. The load values and BSR_{max} demonstrate that our experimental testbed does not reach the link capacity limit as defined in Subsection 5.1. The *Delay* follows the value of `t-StatusProhibit` except for 8 Mbps. In addition, the 8 Mbps present outliers on *L* standard deviation, *IR* max and *Delay* mean.

The decrease of *L* mean, and its huge difference between 0.5 and 8 Mbps, as well as the *L* standard deviation outlier from 8 Mbps could be explained by the allocation resource adaptation from the BS triggered by the UL buffer occupation increase and notified to the BS by the BSR. The *IR* and *Delay* outliers are due to the constant UL buffer occupation, which might avoid the BS to send the `STATUS report` in time and let the *IR* grow. The *IR* values allows us to the set the *Nacks* variation in the subsequent experiments.

6.3 Experiment 2: Attack Impact on the RLC Layer

6.3.1 Scenario. In this experiment, we observe the behaviour at the RLC layer without attack and under attack with *Nacks* set to 8. The bitrate is set to 2 Mbps.

6.3.2 Results and Interpretations. Figure 5a and Figure 5b present the SN values as a function of time. Figure 5a presents the protocol overview of RLC without attack. We observe an irregularity of $ULallocation$, which impacts the delay between Tx_sn. Thus, the Tx_sn irregularity could explain the standard deviation presented in the results of Experiment 1 in Table 1. Figure 5b introduces the protocol overview under attack. This figure demonstrates that the reception of the $Nack\_sn$ value triggers contiguous retransmissions. We observe the impact on the delay between two SN pairs ($Tx\_xn = \{14, 15\}$ and $\{29, 30\}$) by 20.95 milliseconds, that reduces the UE regular transmissions (Tx_sn) as compared to Figure 5a.

The attack triggers a contiguous delay due to the priority of the re-transmission induced by the re-sequencing of RLC SN (Tx_xn) in 4G.

6.4 Experiment 3: Impact of the Bitrate Variation

6.4.1 Scenario. In Experiment 3, we evaluate the impact of different bitrates under attack, and we compare it with the baseline values in Table 1. We consider some couples of bitrate and *Nacks*, the latter being fixed at a power of 2 closest to the minimum of bitrate's *IR* as provided in Table 1. We repeated the experiment nine times.

Let l_N and l_A the average of *L* as defined in Subsection 5.2 and measured for the tests without and with attack, respectively. We define the increase rate as follows:

$$L_R = \frac{l_N - l_A}{l_N}$$

Similarly, we compute the relative increase of the Load and the Buffer size as defined in Subsection 5.2. We define them as Ld_R and B_R, respectively.

6.4.2 Results and Interpretations. Figure 6 shows the observed *L* (L_R), Load (Ld_R) and Buffer (B_R) percentage of increase according to the bitrate under attack, as compared to Experiment 1 in Table 1. These results demonstrate that the re-transmissions triggered by the NACK reception imply an increase of all observed metrics and indicators and they impact each observed bitrate. The comparison between each tuple of results denotes the correlation between the increase of the re-transmission triggered by *Nacks* and the rise expansion of $Buffer$ (B_R) and *L* (L_R). The Load increase follows the same trend observed on results from Experiment 1 with an increase between each Ld_R up to 2 Mbps and a decrease starting from 4 Mbps.

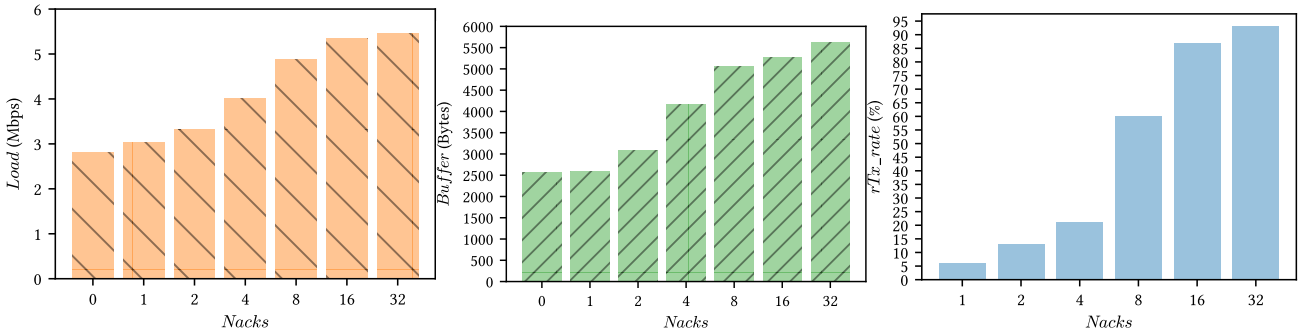

(a) Load average increases **(b) Buffer increases** **(c) rTx_rate average increases**

Figure 7: Experiment 4, impact on srsLTE UE Ld and B, and rTx_rate related to $Nacks$ variation

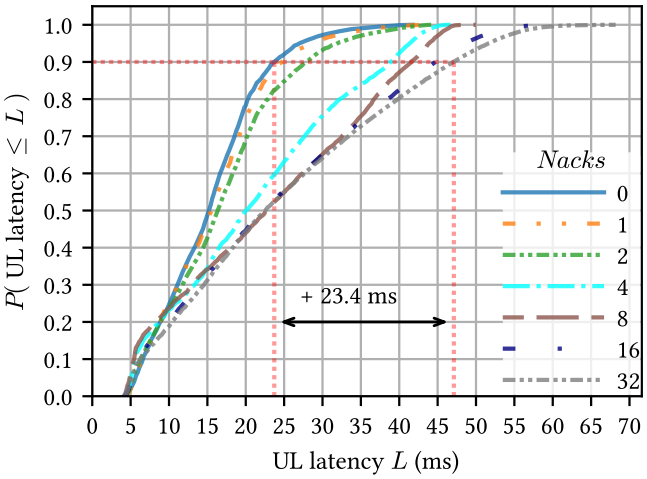

Figure 8: Experiment 4, srsLTE UE UL latency (L) empirical CDF related to $Nacks$ variation

We do not present the 8 Mbps results in this experiment because the attack triggers a saturation of bandwidth. Regarding the latency, we exhibited in Subsection 6.3 that the re-transmissions are contiguous, so the delay should increase as a function of the re-transmission number.

6.5 Experiment 4: Impact of the $Nacks$ Increase

6.5.1 Scenario. In Experiment 4, we evaluate the attack impact when the $Nacks$ parameter increases. The bitrate value is fixed to 2 Mbps and represent a low video streaming flow. The $Nacks$ is now a variable parameter; its start from 0, up to the max IR observed from Table 1 for this bitrate value, rounded by the power of 2.

6.5.2 Results and Interpretations. Figure 8 presents the empirical CDF of L where each curve denotes the $Nacks$ value used. We observe that the $Nacks$ increase raises the impacted UDP traffic proportion, which is almost between 70% and 90 %. The latency degradation follows the same trend. Let L_9^{Nacks} be the last decile of L with its $Nacks$ parameter value. The attack increases L of 23.4 ms, between L_9^0 and L_9^{32}, which denotes an efficient degradation of latency guarantee.

Figure 7a, Figure 7b and Figure 7c present the continuous increase of Ld, B, and rTx_rate, respectively, as a function of $Nacks$. Their raise are correlated with the $Nacks$ increase and represents the side effects implied by an overuse of bandwidth and RLC buffer triggered by the re-transmissions. Figure 7c denotes a gap in the amount of traffic re-transmitted between 4 and 8 $Nacks$. In addition, the buffer does not suffer from the 1 $Nacks$.

Figure 8 and Figure 7c show that the impact starts to be significant from 4 $Nacks$ with 21 % of traffic re-transmitted. It is therefore optional to re-transmit all the traffic to create a significant impact.

6.6 Experiment 5: Validation on a COTS UE

Finally, we performed an evaluation similar to Experiment 4 (Subsection 6.5) to validate and evaluate the impact of the overshadowing attack on a COTS UE. We executed this experimental evaluation three times.

6.6.1 Change in setup. The testbed changed on the UE side with an Asus Zenfone 8 that comes with Android 11, a Central Processing Unit (CPU) @2.8 GHz, and a System on a Chip (SoC) Qualcomm Snapdragon 888. The change on the BS side is the use of an omni-directional antenna with LTE capabilities in place of cables, inside a Faraday cage.

We use `tethering` to measure the UL latency. We still have access to Ld load from the srsLTE's BS side and to our latency (L) measurement.

6.6.2 Results and Interpretations. Figure 9a and Figure 9b present the L ECDF and the Ld average increase both as described in experiment 6.5, respectively, and bring the same meaning as Figure 8 and Figure 7a from Experiment 4. We observe almost the same range of UDP traffic impacted as in Experiment 6.5 on L. The increase between L_9^0 and L_9^{32}, is less from srsLTE, with a value of 17.66 ms. However, it still demonstrates the attacker's capacity to significantly degrade the latency guarantee. We also observe side effects on the $Load$, with its continuous increase as a function of $Nacks$, which denotes an overuse of the UE resources, as was the case in Experiment 6.5. The curves and bars for 16 and 32 $Nacks$

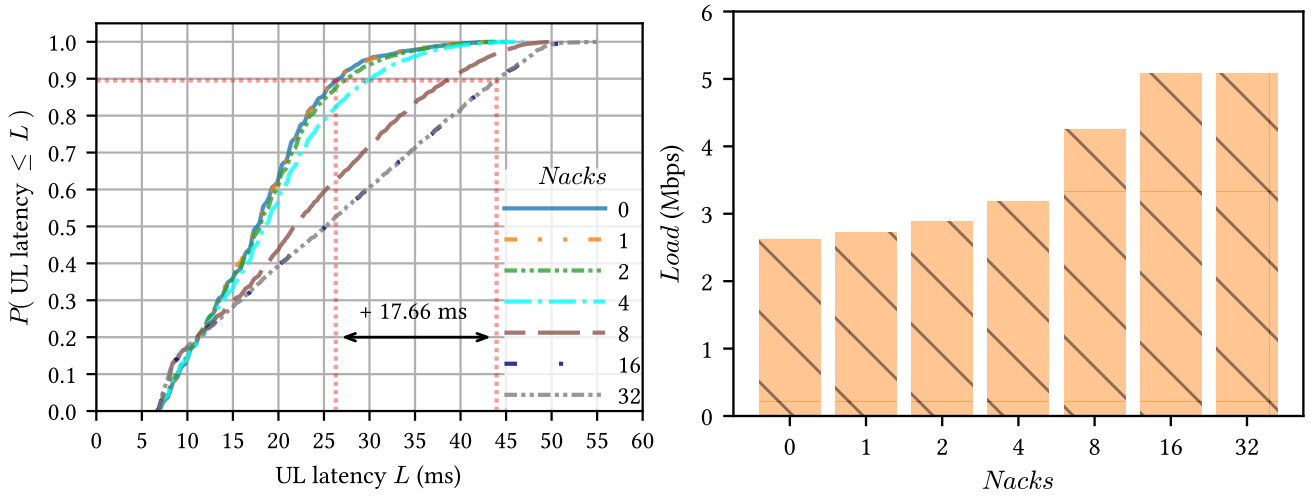

(a) UL latency (L) empirical CDF

(b) Load Ld average

Figure 9: Experiment 5, impact on COTS UE related to *Nacks* variation

in Figure 9a and Figure 9b demonstrate a limit in the impact of the increase.

Overall, the attack impacts negatively the UE COTS, which is vulnerable to this attack. In addition, the impact is raised by the increase of *Nacks*. However, the impact on L and Ld are less significant than on the srsLTE implementation. The significant impact on srsLTE starts from 4 NACK, while from 8 NACK on the UE COTS. Considering the time-delay objective, this slight decrease of the attack impact is not an issue for the attacker given that the impact is still sufficient to degrade the latency required for specific flows.

6.7 Limitations

The previously presented results validate the possibility to conduct a RoQ on latency with the hijacking of the RLC AM. However, our contribution comes with limitations in terms of assumptions. Indeed, we implemented the attack as an emulation on the BS side. Thus we do not proceed to a real overshadowing and the eavesdropping of prerequisites on RLC. In addition, we described the vulnerabilities of 5G, but we conducted the attack on 4G; even though the protocol stack is almost the same, our empirical approach does not strictly validate the feasibility on 5G networks. In addition, the 5G NR may not prioritise re-transmission without re-sequencing at RLC, and this changes the impact of our attack on the latency.

7 CONCLUSION AND FUTURE WORK

The results we exposed in this paper show that the latency increase impacts all of the observed bitrates of our experimental testbed. We also demonstrate that an attacker can increase the impact on the latency guarantee by controlling the NACKs quantity. The experimental tests highlight an overuse of the RLC buffer and the load generated over the RAN. In addition, it shows that we do not have to set the power of the attack at the maximum (i.e., highest number of NACKs) to obtain a notable degradation of the latency. The attack against a COTS UE shows that it is also vulnerable and impacted by the attack on the latency guarantee and the load.

The following steps in our line of work consists in understanding the impact of our attack on a real application. We consider a Metaverse-like flow to extract a latency threshold from the required QoS, and verify if the impact on latency from our attack is a real threat. This should be further studied using a COTS UE running Metaverse-based applications and quantifying the latency-increase impact on application-oriented metrics. Furthermore, we plan to select a 5G implementation to check if our attack is still relevant. Finally, we also plan to conduct those experiments with other UEs in the same cell to check if the observed side effects are also a threat.

Regarding countermeasures, researchers have created a tool that monitors the UE flows and uses Deterministic Finite Automata (DFA) to check the correct behaviour between MAC and RLC layers [22]. For example, if the frame has been acknowledged by the MAC layer, it cannot get a NACK at the RLC layer. However, this solution is external and requires additional equipment. Consequently, a perspective of our works would be to fix the L2's reliability's vulnerability exposed in this paper directly on the 5G (or beyond) standard to offer this protection to most users.

ACKNOWLEDGMENTS

This work was carried out in the context of 5GMetaverse, a project funded by the French government as part of the economic recovery plan, namely "France Relance" and the investments for the future program.

REFERENCES

[1] 3GPP. 2023. *3rd Generation Partnership Project; Technical Specification Group Radio Access Network; NR; Radio Link Control (RLC) protocol specification (Release 18).* Technical Report TS 38.322 V18.0.0 (2023-12). 3GPP. https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3195

[2] 3GPP. 2024. *3rd Generation Partnership Project; Technical Specification Group Radio Access Network; Evolved Universal Terrestrial Radio Access (E-UTRA); Physical layer procedures (Release 18).* Technical Report TS 36.213 V18.2.0 (2024-03). 3GPP. https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=2427

[3] 3GPP. 2024. *3rd Generation Partnership Project; Technical Specification Group Radio Access Network; NR; NR and NG RAN Overall Description; Stage 2 (Release 18).* Technical Report TS 38.300 V18.1.0 (2024-03). 3GPP. https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3191

[4] Youness Arjoune and Saleh Faruque. 2020. Smart Jamming Attacks in 5G New Radio: A Review. In *2020 10th Annual Computing and Communication Workshop and Conference (CCWC).* 1010–1015. https://doi.org/10.1109/CCWC47524.2020.9031175

[5] Evangelos Bitsikas and Christina Pöpper. 2021. Don't hand it Over: Vulnerabilities in the Handover Procedure of Cellular Telecommunications. In *Proceedings of the 37th Annual Computer Security Applications Conference* (<conf-loc>, <city>Virtual Event</city>, <country>USA</country>, </conf-loc>) *(ACSAC '21).* Association for Computing Machinery, New York, NY, USA, 900–915. https://doi.org/10.1145/3485832.3485914

[6] Chunyu Chen, Yang Chen, Kaifeng Zhang, Ming Ni, Shushan Wang, and Rui Liang. 2021. System Redundancy Enhancement of Secondary Frequency Control Under Latency Attacks. *IEEE Transactions on Smart Grid* 12, 1 (2021), 647–658. https://doi.org/10.1109/TSG.2020.3012977

[7] Simon Erni, Martin Kotuliak, Patrick Leu, Marc Roeschlin, and Srdjan Capkun. 2022. AdaptOver: adaptive overshadowing attacks in cellular networks. In *Proceedings of the 28th Annual International Conference on Mobile Computing And Networking* (Sydney, NSW, Australia) *(MobiCom '22).* Association for Computing Machinery, New York, NY, USA, 743–755. https://doi.org/10.1145/3495243.3560525

[8] Maya E. Flores, Devon D. Poisson, Colin J. Stevens, Adriyel V. Nieves, and Alexander M. Wyglinski. 2023. Implementation and Evaluation of a Smart Uplink Jamming Attack in a Public 5G Network. *IEEE Access* 11 (2023), 75993–76007. https://doi.org/10.1109/ACCESS.2023.3296701

[9] Margarita Gapeyenko, Vitaly Petrov, Stefano Paris, Andrea Marcano, and Klaus I. Pedersen. 2023. Standardization of Extended Reality (XR) over 5G and 5G-Advanced 3GPP New Radio. *IEEE Network* 37, 4 (2023), 22–28. https://doi.org/10.1109/MNET.003.2300062

[10] Matheus E. Garbelini, Zewen Shang, Sudipta Chattopadhyay, Sumei Sun, and Ernest Kurniawan. 2022. Towards Automated Fuzzing of 4G/5G Protocol Implementations Over the Air. In *GLOBECOM 2022 - 2022 IEEE Global Communications Conference.* 86–92. https://doi.org/10.1109/GLOBECOM48099.2022.10001673

[11] S. Hussain, O. Chowdhury, S. Mehnaz, and E. Bertino. [n. d.]. LTEInspector: A Systematic Approach for Adversarial Testing of 4G LTE. *Network and Distributed Systems Security (NDSS) Symposium 2018* ([n. d.]). https://par.nsf.gov/biblio/10055689

[12] Syed Rafiul Hussain, Mitziu Echeverria, Imtiaz Karim, Omar Chowdhury, and Elisa Bertino. 2019. 5GReasoner: A Property-Directed Security and Privacy Analysis Framework for 5G Cellular Network Protocol. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security* (London, United Kingdom) *(CCS '19).* Association for Computing Machinery, New York, NY, USA, 669–684. https://doi.org/10.1145/3319535.3354263

[13] Imtiaz Karim, Syed Rafiul Hussain, and Elisa Bertino. 2021. ProChecker: An Automated Security and Privacy Analysis Framework for 4G LTE Protocol Implementations. In *2021 IEEE 41st International Conference on Distributed Computing Systems (ICDCS).* 773–785. https://doi.org/10.1109/ICDCS51616.2021.00079

[14] Xin Lou, Cuong Tran, Rui Tan, David K. Y. Yau, Zbigniew T. Kalbarczyk, Ambarish Kumar Banerjee, and Prakhar Ganesh. 2020. Assessing and Mitigating Impact of Time Delay Attack: Case Studies for Power Grid Controls. *IEEE Journal on Selected Areas in Communications* 38, 1 (2020), 141–155. https://doi.org/10.1109/JSAC.2019.2951982

[15] Norbert Ludant and Guevara Noubir. 2021. SigUnder: a stealthy 5G low power attack and defenses. In *Proceedings of the 14th ACM Conference on Security and Privacy in Wireless and Mobile Networks* (Abu Dhabi, United Arab Emirates) *(WiSec '21).* Association for Computing Machinery, New York, NY, USA, 250–260. https://doi.org/10.1145/3448300.3467817

[16] Norbert Ludant, Marinos Vomvas, and Guevara Noubir. 2024. Unprotected 4G/5G Control Procedures at Low Layers Considered Dangerous. arXiv:2403.06717 [cs.CR]

[17] Qinpei Luo, Hongliang Zhang, Minrui Xu, Boya Di, Anthony Chen, Shiwen Mao, Dusit Niyato, and Zhu Han. 2023. An Overview of 3GPP Standardization for Extended Reality (XR) in 5G and Beyond. *GetMobile: Mobile Comp. and Comm.* 27, 3 (nov 2023), 10–17. https://doi.org/10.1145/3631588.3631592

[18] Sang-Min Park and Young-Gab Kim. 2022. A Metaverse: Taxonomy, Components, Applications, and Open Challenges. *IEEE Access* 10 (2022), 4209–4251. https://doi.org/10.1109/ACCESS.2021.3140175

[19] Roger Piqueras Jover and Vuk Marojevic. 2019. Security and Protocol Exploit Analysis of the 5G Specifications. *IEEE Access* 7 (2019), 24956–24963. https://doi.org/10.1109/ACCESS.2019.2899254

[20] Altaf Shaik, Ravishankar Borgaonkar, N. Asokan, Valtteri Niemi, and Jean-Pierre Seifert. 2017. Practical Attacks Against Privacy and Availability in 4G/LTE Mobile Communication Systems. arXiv:1510.07563 [cs.CR]

[21] Zhaowei Tan, Boyan Ding, Jinghao Zhao, Yunqi Guo, and Songwu Lu. 2022. Breaking Cellular IoT with Forged Data-plane Signaling: Attacks and Countermeasure. *ACM Trans. Sen. Netw.* 18, 4, Article 59 (nov 2022), 26 pages. https://doi.org/10.1145/3534124

[22] Zhaowei Tan, Jinghao Zhao, Boyan Ding, and Songwu Lu. 2023. CellDAM: User-Space, Rootless Detection and Mitigation for 5G Data Plane. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23).* USENIX Association, Boston, MA, 1601–1619. https://www.usenix.org/conference/nsdi23/presentation/tan

[23] Hang Wang, Huansheng Ning, Yujia Lin, Wenxi Wang, Sahraoui Dhelim, Fadi Farha, Jianguo Ding, and Mahmoud Daneshmand. 2023. A Survey on the Metaverse: The State-of-the-Art, Technologies, Applications, and Challenges. *IEEE Internet of Things Journal* 10, 16 (2023), 14671–14688. https://doi.org/10.1109/JIOT.2023.3278329

[24] K. Whitehouse, A. Woo, F. Jiang, J. Polastre, and D. Culler. 2005. Exploiting the capture effect for collision detection and recovery. In *The Second IEEE Workshop on Embedded Networked Sensors, 2005. EmNetS-II.* 45–52. https://doi.org/10.1109/EMNETS.2005.1469098

[25] Hojoon Yang, Sangwook Bae, Mincheol Son, Hongil Kim, Song Min Kim, and Yongdae Kim. 2019. Hiding in Plain Signal: Physical Signal Overshadowing Attack on LTE. In *28th USENIX Security Symposium (USENIX Security 19).* USENIX Association, Santa Clara, CA, 55–72. https://www.usenix.org/conference/usenixsecurity19/presentation/yang-hojoon

[26] Jingda Yang, Ying Wang, Tuyen X. Tran, and Yanjun Pan. 2023. 5G RRC Protocol and Stack Vulnerabilities Detection via Listen-and-Learn. In *2023 IEEE 20th Consumer Communications & Networking Conference (CCNC).* 236–241. https://doi.org/10.1109/CCNC51644.2023.10059624