Low-Layer Attacks Against 4G/5G Networks

Norbert Ludant
Northeastern University
Boston, MA, USA

Marinos Vomvas
Northeastern University
Boston, MA, USA

Stavros Dimou
Northeastern University
Boston, MA, USA

Guevara Noubir
Northeastern University
Boston, MA, USA

Abstract

Prior security analyses of 3GPP systems primarily focus on upper layers of the stack. Unfortunately, the physical and MAC layers are not as thoroughly analyzed, even though they are neither encrypted nor integrity protected. Furthermore, the latest 5G releases significantly increase the number of low-layer control messages and procedures. We conduct a systematic vulnerability analysis of these low layers, and demonstrate that current cellular systems are susceptible to passive attacks, and active spoofing of PHY/MAC messages. For instance, we find that sniffing beamforming information enables fingerprinting-based localization and tracking of users. We also show that signal spoofing is possible in 5G NR, and more efficient compared to LTE networks. We also evaluate active attacks against COTS UEs, showing it is possible to disrupt user communications by tricking connected UEs into acting as jammers, or by stealthily disconnecting active users. In our experiments we achieve user localization within 20-meters 96% of the time, user path tracking within 15 meters for 81% of the paths, and throughput reduction by over 95% within 2 seconds (by spoofing a 39-bit DCI).

CCS Concepts

• **Security and privacy → Mobile and wireless security**; **Security protocols**; • **Networks → Mobile networks**.

Keywords

Wireless Communications, Cellular Networks, 5G, Wireless Privacy, Localization, Denial of Service

ACM Reference Format:
Norbert Ludant, Marinos Vomvas, Stavros Dimou, and Guevara Noubir. 2025. Low-Layer Attacks Against 4G/5G Networks. In *18th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec 2025), June 30-July 3, 2025, Arlington, VA, USA*. ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3734477.3734725

1 Introduction

To defend against recent attacks against higher layers [1–4], 3GPP reinforced 5G security with the encryption of the initial NAS message, the introduction of the SUCI, and mandated integrity protection [5]. However, the demand for lower latency caused a push of control mechanisms to lower layers of the stack to allow faster reconfiguration of the network. This raises a security concern, as the lower layers are not encrypted nor integrity protected, and can be exploited by an adversary, targeting even authenticated users.

We study the security vulnerabilities of low layer procedures and control elements, and present new, passive and active attacks that span various procedures and control mechanisms. We experimentally demonstrate the first over-the-air injection of signals such as Downlink Control Informations (DCIs) or MAC CEs in 5G NR, and the feasibility of these attacks in our testbed. For ethical reasons, we evaluate our active attacks in an isolated testbed with COTS UE, and our passive attacks against our own devices using commercial operators with no impact on other UEs or the network.

Our passive attacks exploit beam management procedures enabling a passive attacker to localize and track users by exploiting the spatial beamforming configuration. We show that it is possible to localize users with an accuracy below 20 meters 96% of the time and track the movement of active users with a maximum deviation of less than 15 meters for 81% of the paths by monitoring random access exchange, and channel state reports respectively. These attacks require a single device and no hardware calibration, as they rely only on information reported by the target User Equipment (UE). We evaluate the attacks using three different phones and we observe consistent reports across all phones.

2 5G Low-Layer Background

The 5G Radio Access Network (RAN) follows many of the radio-design principles of 4G LTE, with an emphasis on flexibility, adapting to new types of UEs and services. Both technologies use OFDM, and the time and frequency resources are divided in a resource grid which accommodates multiple physical channels for different purposes. For channel estimation, LTE constantly broadcasts a set of pilots (CRS), while 5G pilots (DMRS) are *only* sent with transmissions, which allows for more efficient OTA injection in 5G.

UEs in the RAN are uniquely identified and addressed by their Radio Network Temporary Identifier (RNTI). The RNTI is assigned during initial access upon completion of the Random Access (RA) procedure, and updated if the Radio Resource Control (RRC) connection is re-established, for instance, after an inactivity period.

Finally, as beamforming technologies matured, additional adaptations to the physical layer were introduced in 5G to support its operation: For instance, beam measurement, beam reporting and beam tracking are included in the physical layer due to their low latency requirements.

The RAN control plane protocol stack is divided in three layers: (a) L1 which contains PHY, (b) L2 which contains Medium Access

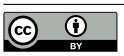

WiSec 2025, Arlington, VA, USA

ACM ISBN 979-8-4007-1530-3/2025/06
https://doi.org/10.1145/3734477.3734725

Control (MAC), Radio Link Control (RLC) and Packet Data Convergence Protocol (PDCP), and (c) L3 which contains RRC. The MAC layer is the link between the physical and higher protocol layers, and is in charge of resource scheduling and mapping between logical and transport channels. A single transport channel carries multiple logical channels, for instance, the Downlink (DL) shared channel can contain traffic or control data. To facilitate this, the MAC layer includes a header that indicates which type of logical channel is associated with its payload, using a field termed Logical Channel ID (LCID). Further, the MAC layer carries data and control plane signalling in the MAC header via MAC Control Elements (CEs). The purpose of MAC CEs is to dynamically modify the radio configuration, complementing the static radio configuration managed by the RRC. The LCID header field is therefore overloaded to also convey the type of MAC CE.

3 Low-layer attacks

We present the adversary model, the discovered vulnerabilities and attacks, then discuss the practical feasibility of performing said attacks. Our attacks can be used to passively localize and track user movements, or actively modify the state of a UE by injecting signaling commands. This is possible because these layers are not protected (no encryption neither integrity protection). This allows an attacker to enable or disable radio techniques at the UE, perform amplified distributed jamming attacks, or trigger Uplink (UL) wideband radio transmissions and RA procedures. This leads to battery draining, performance degradation, and network access disruption, with low power overhead for the attacker.

Adversarial model. We assume an attacker positioned within the coverage of the 4G/5G Base Station (BS), equipped with a Software Defined Radio (SDR). The attacker is also capable of decoding the plaintext BS transmissions, as shown previously in [6–11]. The attacker can generate and inject 4G/5G wireless signals over the air at desired time instants, as demonstrated by previous research for LTE [12–14]. We demonstrate OTA injection in 5G NR for the first time in Section 4.2.

Generic attack overview. We outline common aspects of our attacks. The attacker listens to the broadcast synchronization signals to synchronize with the BS. They then decode additional unprotected information from the physical channels, such as the Physical Downlink Control Channel (PDCCH). The attacker deduces the number of connected users and their RNTIs. The attacker has access to scheduling information pointing to user-specific data in the resource grid and can eavesdrop unprotected L1/L2 data [10].

As the RF spectrum is an open medium and L1 and L2 layers are unprotected, the attacker either eavesdrops the exchange of information between UEs and BS, or performs active attacks, such as spoofing messages to specific users, addressed by their RNTI. L1 messages do not contain headers of the protocol stack, and can be spoofed very efficiently by an attacker by injecting a signal on the physical channel, such as the PDCCH.

To perform an L2 injection in the DL, additional steps are required. The attacker crafts a packet containing L2 data, e.g., a MAC header, and encodes it in the Physical Downlink Shared Channel (PDSCH). They also craft a DCI, containing the RNTI of the user it is addressed to, encoded in the PDCCH. This DCI points to the location of the crafted L2 message in the PDSCH. The attacker, then, transmits the crafted PDCCH and PDSCH aligned with the BS resource grid at slightly higher power than the BS, ensuring overshadowing of the legitimate BS transmissions. UL attacks follow the same principle: the attacker monitors for UL grants in the DCI addressed to a given RNTI, and overshadows the UE transmission with the crafted L2 header on the Physical Uplink Shared Channel (PUSCH). When the target receives the crafted message it has no means to verify its legitimacy, and applies the indicated change.

Practical considerations. MAC CE spoofing is constrained by the encrypted RRC configuration, unknown to the attacker. However, attackers can infer UE states passively via open-source tools [10] or actively probe UE responses. Moreover, our targeted attacks require obtaining the UE temporary RNTI. Techniques for decoding active RNTIs exist in both 4G [6–8] and recently 5G [9, 11], and mapping RNTIs to persistent user identities has also been demonstrated [4, 9, 15]. Finally, decoding mmWave signals is challenging due to directionality. However, attackers co-located with base stations or exploiting mmWave sidelobes can reliably decode signals [16]. While open-source mmWave decoders are unavailable, attackers can use academic testbeds [16] or commercial equipment [17].

3.1 Identified vulnerabilities

3.1.1 Physical Downlink Control Channel (PDCCH). The PDCCH is used to convey resource scheduling information to the UE. This is carried in the DCI, which contains physical layer resource allocation for the downlink and uplink, along with the required data-encoding information. Additionally, DCIs contain fields to manage various control procedures of the UE, such as power control, Hybrid Automatic Repeat reQuest (HARQ) information, or trigger RA.

We analyze the impact of active attacks on DCI, that aim to disrupt communication by spoofing either resource scheduling, or control commands to the UEs. DCI spoofing is highly energy-efficient for the attacker, as it only occupies a small subset in the PDCCH region, as little as 72 subcarriers.

Attacks on resource scheduling. DCI spoofing against a UE can be used to fake the allocation of resources for either the downlink or the uplink. On the other hand, UL resource scheduling proves more interesting, as an attacker can force multiple UEs to transmit over the same resources, causing jamming of legitimate users and battery draining.

To perform the attack, the attacker selects one or more connected UEs, which we call Induced-Jammer UEs (IJ-UE), and injects an UL DCI in every time slot, allocating a set of resource blocks to them. This instructs the IJ-UEs to transmit data over the allocated RBs during every time slot. *We discovered that the UEs use all the allocated resources regardless of how much pending data they need to transmit.* This derives from constraints related to the design and structure of 3GPP L1 scrambling and interleaving. Therefore, the UEs pad their UL data to fill all allocated resources. Moreover, to maximize the impact, the attacker uses the Transmission Power Control (TPC) field in the same DCI to instruct the exploited UEs to transmit at the maximum power [18]. This causes the SINR of

other connected devices to drop close to or below 0, and severely impacts their throughput. We showcase this behavior for a COTS UE and evaluate the potential of this attack in Section 4.

PDCCH Order (PO). PO is a special DCI used to instruct a connected UE to initiate an RA procedure to re-establish synchronization in the UL (i.e., update the Timing Advance (TA) value). This control procedure is initiated by a DCI with predefined fields, indicating it is a PO message. PO is the only unprotected control procedure that can trigger RA for connected users, as other network-initiated RA procedures are triggered through the protected RRC layer. This makes PO particularly interesting as it provides an efficient and stealthy way to instruct a UE to perform RA. The RA procedure involves a 4-message exchange between UE and BS, as well as DL and UL resource allocations. A simple exploitation of this vulnerability leads to draining of the limited RA resources by injecting a single DCI, resulting in collisions. Additionally, it can be paired with other vulnerabilities to disconnect users and trigger localization attacks, as we discuss in Section 3.1.3.

3.1.2 Physical Uplink Control Channel (PUCCH).

Spoofing Scheduling Request. This uplink physical layer message is sent from the UE to the BS to request uplink resources. Upon reception of a SR, the base station will allocate resources for a user. Scheduling requests can be spoofed and leveraged by an attacker in three ways: i) maintaining users' RNTI connection active for long periods of time, bypassing the RRC inactivity timer, which enables long-term tracking, ii) requesting resources on behalf of multiple users that do not have pending uplink data, leading to congestion in network resources, and iii) to request an UL DCI for a specific user, and hijack the allocated UL grant to spoof higher-layer data on behalf of the user.

HARQ-Attack. Message acknowledgments (ACKs) are a crucial part of communications. In the 5G uplink channel, ACKs are scheduled dynamically and can be aggregated in the same UCI. The ACK timing is specified in the DL DCI PDSCH-to-HARQ Feedback Timing Indicator field, that informs the UE of the uplink time slot where the ACK should be reported. The UE reports a dynamic sized bitmap, using one bit for each aggregated ACK. However, if the UE fails to decode a DL DCI, the bitmap size would be mismatched between the BS and the UE. In order to address this scenario, the DL DCI includes a counter (Downlink Assignment Index (DAI)), used by the UE to determine whether any transmissions were lost, and adjust the bitmap size accordingly (Figure 1, left). However, these mechanisms do not account for active attackers. An attacker aiming to disrupt the communication of a specific user, spoofs a DCI with a modified DAI counter that breaks the synchronization between transmitted and received packets at the UE, as shown in Figure 1 (right). In this case, when the UE reports the ACK bitmap, it does not match the size expected by the BS, and the BS is unable to determine the ACK-packet mapping, leading to a HARQ failure. We demonstrate this attack in our testbed using a COTS UE (Section 4).

Sniffing Channel State Information (CSI) Reports. UEs perform measurements of the quality of the downlink channel, and report the measured CSI to the BS. The BS uses the CSI parameters as input for resource scheduling, beam management, and MIMO.

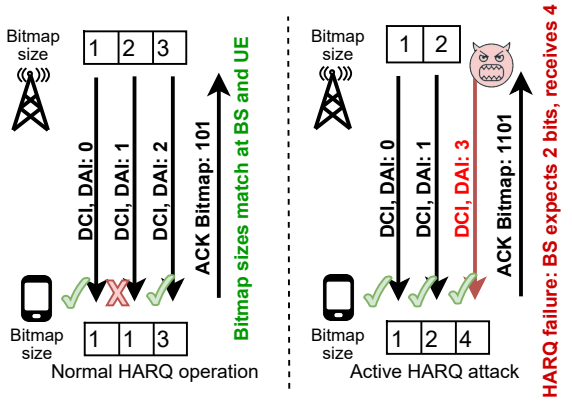

Figure 1: Two examples of the HARQ procedure during a missed allocation in normal operation (left) and under an active attack (right). The DAI counter implicitly indicates missed transmissions, and can be leveraged by an attacker.

CSI reports are carried by UCI but can be scheduled by MAC layers. We describe a user tracking attack using CSI reports in Section 3.1.5.

3.1.3 Physical Random Access Channel (PRACH). The PRACH is used for the RA procedure, which is the first step to attach to the network. This procedure is unprotected and therefore susceptible to eavesdropping and spoofing. We identify a set of attacks that target the RA procedure.

Blocking initial cell access. The RA parameters are broadcast by the BS in the System Information Block (SIB), and enable UEs to connect to the network. An attacker can overshadow the SIBs and modify the RA configuration at the UE side. Specifically, the attacker modifies `ra-ResponseWindowSize` to the minimum value, (i.e., `sf2`) to shorten the window. This forces the UE to only monitor for a Random Access Response (RAR) message within 3 + `ra-ResponseWindowSize` subframes as described in the 3GPP standard. When this timer elapses, the UE deems the RA unsuccessful and retries for a maximum of `preambleTransMax` attempts. To increase congestion, the spoofed SIBs also set the `preambleTransMax` to the maximum value (i.e., 200).

This attack causes the RA to fail for all UEs that are trying to connect to the network, but does not affect already connected UEs, which do not monitor the SIBs. To target already connected users, the attacker injects an SIB paging, which instructs all connected UEs to monitor the SIB for updates, as previously shown in [14]. The attacker can now inject a PDCCH Order (PO) DCI to a target UE, which triggers the RA procedure.

This attack is based on the assumption that the BS, unaware of the change in RA parameters, does not prioritize the RAR, leading to RA failure due to RAR timer expiration. Further, POs to multiple users can be injected at the same time, since the PDCCH can fit many DCIs, which amplifies the collision effect during RA. The repeated RA attempts flood the random access channel, creating overhead at the BS, and hindering the access for new users. Additionally, the UE ramps up the power after every failed RA causing increased battery consumption. As a result, by paging active users and overshadowing the SIB, the attacker controls the time of the

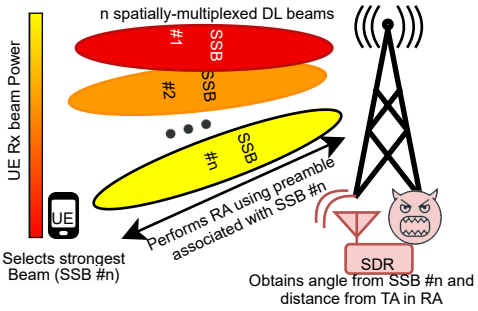

Figure 2: The UE measures the strongest beam and implicitly reports it during RA. An attacker passively listening to the exchange, which includes TA, can then localize the user.

attack and operates in a power-efficient and stealthy manner, instead of continuously overshadowing the SIB. In Section 4.2, we validate our assumption with a real-world passive measurement campaign of RAR times of three major cellular Mobile Network Operators (MNOs) and also demonstrate the active attack on our isolated testbed.

SSB-RA fingerprinting localization attack. 5G NR uses beamforming techniques to create directional transmissions (i.e., *beams*) that focus the transmitted signals to specific locations. New mechanisms are introduced in 5G to support beamforming. Specifically, the RA procedure is modified to include an implicit beam-reporting mechanism that facilitates beamforming since the initial message exchange. This procedure is depicted in Figure 2. The BS broadcasts the Synchronization Signal Blocks (SSBs) over different beams, each identified by a unique index in the cell. The UE measures the received power for each SSB, and determines the strongest beam. 3GPP defines a one-to-one mapping between the SSB beam indexes and RA occasions, such that the BS can determine the optimal beam for the UE based on the RA occasion used by the UE.

This process can be exploited by an attacker, using a fingerprinting-based approach, to localize 5G users. The attacker fingerprints the beam configuration of a cell, and creates a map with the precise locations of the BS and every beam in that cell. Then, the attacker monitors the RA channel and deduces the beam chosen by the UE from the RA occasion. Additionally, the attacker obtains the TA value from the RA response sent by the BS. These values are used to estimate the azimuth of the UE and the distance from the BS respectively, resulting in an estimate of the UE location. Even though RA is primarily used for initial access, the attacker can also target already connected users by performing a PO spoofing attack to trigger the RA procedure, and estimate their location. We note that our attack does not require calibration or multiple devices for triangulation, and relies solely on measurements reported by the UEs. In Section 4.2 we evaluate this attack by fingerprinting the beam locations in a real-world deployment, and measuring the localization error against our test device.

3.1.4 Attacks on Carrier Aggregation. Carrier Aggregation (CA) has been widely deployed to provide high throughput gains by aggregating bandwidths at multiple cells operating at different center frequencies, denoted Secondary Cells (SCells). One of the main drawbacks of CA is the considerable energy consumption at the UE, which is monitoring the PDCCH simultaneously in multiple frequencies, and transmitting periodic CSI reports. CA energy consumption measurements show that UEs experience an average power consumption increase of 79% when activating one additional SCell [19]. To tackle this issue, SCells are activated and deactivated dynamically at the MAC layer by the MAC CE *SCell Activation/Deactivation*. This MAC CE consists of a bitmap which indicates the SCells to be activated or de-activated at the UE. The deactivation process can be performed either by the MAC CE, or by expiration of the *sCellDeactivationTimer*, ranging from 20ms to 1.28s. If unspecified, the de-activation timer is set to infinity [20].

As a result, a malicious SCell Activation MAC CE against a UE stealthily forces it to use a considerably higher amount of energy. Unaware of this change, the BS does not de-activate the active SCells through a DL MAC CE. In Section 4.2 we experimentally measure the potential of this attack in a real world setup by surveying the *sCellDeactivationTimer* configuration for three major MNOs.

Moreover, an attacker can de-activate SCells for a device that uses CA without the knowledge of the BS. This incurs a drastic throughput reduction at the physical layer, and even higher in the application layer. For instance, if the traffic is delivered over TCP, the data is multiplexed at the MAC layer, and the throughput plummets due to retransmissions and out-of-sequence delivery.

3.1.5 Tracking and localization using reference signals. The Channel State Information Reference Signal (CSI-RS) is transmitted periodically in the downlink by the base station, and the UE reports the measured CSI back to the BS for scheduling purposes, as described in Section 3.1.2. Particularly in beamforming scenarios, the CSI includes a pair of values, $\{B_{idx}, RSRP\}$, where B_{idx} is the beam index identifier of the strongest measured beam, and $RSRP$ is the measured signal strength for that beam. This enables swift beam management, as the BS will use this information to transmit the downlink information to a UE using the strongest beam. The CSI report is sent in the clear, carried by the PUCCH, and contains the RNTI of the UE.

This information can be used by an attacker to track users' location in three steps: First, an attacker fingerprints the static cell beam configuration, i.e., measures the physical area covered by each beam index. This step is common with our beamforming localization attack described in Section 3.1.3. Second, the attacker decodes the PUCCH, retrieving the pair $\{B_{idx}, RSRP\}$ from the CSI reports and the RNTI of the UE. Finally, the attacker estimates the GPS coordinates that describe the UE path from the beams reported by the UE, using our path inference algorithm, described in `BeamToPath` (Algorithm 1). Our algorithm compensates for various factors that affect wireless propagation, such as reflections, blockages and non-line-of-sight, using a three-step filtering process.

First, Power Filtering identifies outliers in RSRP values, e.g., high variations in a very short time, due to blockages. Second, the Smoothing filter measures the consecutive occurrences of each reported beam, and discards infrequently observed beams. This can be either due to sporadic reflections, or due to transitions between two contiguous beams, which results in alternating reported beam indexes. Third, the algorithm discards unrealistic beam transitions, e.g., crossing to a non-neighboring beam. Finally, the algorithm

estimates the path using the fingerprinted centroid coordinates of each reported beam, and applies linear interpolation on the estimated coordinates. In Section 4.2 we present our attack evaluation in an urban scenario by tracking the movement of our test device connected to an operator network.

BeamToPath 1 Infer path from CSI reports

Input: $\{\{(B_{idx,i}, RSRP_i)\}_{1..N}, RSRP_{base}\}$
Output: Set of latitude/longitude coordinates (path)
$count \leftarrow 0$ for all idx;
for $n = 0, \ldots, N$ **do**
  if $|RSRP_n - RSRP_{base}| > P_{thres}$ **then**
    continue;
  if $B_{idx,n} == B_{idx,n+1}$ **then**
    $count++$;
  else
    if $(count > C_{thres})$ && areContiguous(B_{idx}, $Path[-1]$) **then**
      addToCoordinates(GPSCoords(B_{idx}));
      $count \leftarrow 0$;
return linearInterpolation(Coordinates);

4 Experimental Evaluation

In this section we present our experimental results that evaluate attacks described in Section 3. We note that active attacks are evaluated in our own isolated anechoic chamber testbed. Additionally, we carry out a passive-measurements survey of MNO network configurations across three different countries to validate if operators use configurations that make users more susceptible to our attacks.

4.1 Evaluation setup

Our evaluation setup consists of three distinctive scenarios: passive UE localization attacks, active injection attacks, and network configuration surveying. We evaluate our passive localization attacks with three mmWave COTS UE phones: Google Pixel 5, OnePlus 8 5G UW, and LG Velvet 5G UW. We connect our devices to the operator's network in an urban setting, and we log relevant, unprotected transmitted information, using QXDM [21]. For active injection attacks, we first demonstrate OTA spoofing of physical signals to a COTS UE connected to a 5G srsRAN gNB [22]. We then perform our measurements using a modified version of srsRAN, where we integrate the attacker in the base station code to ensure consistency and repeatable results. For UE devices we use a COTS UE, Pixel 5 with custom USIM cards, as well as the srsRAN UE implementation. For SDR we use the Ettus USRP B210 and X310. To survey the current configuration of MNOs, we use our COTS UE with USIM cards of the different MNOs, and obtain the network information during UE operation.

4.2 Experimental Results

We evaluate our beamforming-based localization and tracking attacks, DCI spoofing, UE-jammers, HARQ failure attack, and Carrier Aggregation attacks. We start by validating that the Commercial off-the-Shelf (COTS) UE is receiving the MAC CEs and modifies its state accordingly.

4.2.1 **Passive user localization leveraging beamforming leakages.** In this subsection we evaluate our attacks on beamforming that lead to passive user localization (Section 3.1.3) and user movement tracking (Section 3.1.5). We find that even minimal data usage, such as retrieving email or transmitting one stealthy Signal or Telegram message [9] triggers the use of mmWave, making users highly susceptible to the attack.

Fingerprinting beam locations. To perform our attacks, we initially fingerprint the locations of the beams used in an urban environment by a 5G operator, shown in Figure 3. The 5G BS (situated on top of a building, at 23 meters height) operates in the 28 GHz band, with a subcarrier spacing of 120 KHz, and uses 48 beams to cover the area within an angle of approximately 120 degrees. To fingerprint the beamforming configuration, we connect our three different COTS UEs to the BS at multiple locations across the entire area. For each connection, we record the beam index reported by the phone, and the TA value reported by the BS. We also monitor the transition of the phone from one beam to another, while walking throughout the cell coverage area. This indicates which regions of the map are covered by each beam, as shown in Figure 3 by the shapes in different colors. Our fingerprinting comprises the beam areas, the distribution of TA values within each beam area, and one line for each beam area that starts from the BS location and crosses the area's centroid. We note that in the absence of line of sight, beamforming coverage is either lost, or the reflections of other beams are dominant and serve the UE (e.g., Figure 3, beams 30, 31). Finally, we verified that this beam configuration is static, i.e., remained the same throughout all experiments conducted over more than one year period and are independent of the 5G UE. We use the output of this fingerprinting phase for both our localization and movement tracking attacks described below.

SSB-RA localization. In order to evaluate our localization attack, we position our COTS UEs in the coverage of the BS and we connect to the cell. We log the beam index used during RA, the TA value sent by the BS in the RA response, and the exact GPS location. We take measurements with 3 phones, every 2 meters, across the beamforming coverage of the BS, and obtain a dataset with a total of 1835 measurements taken at 320 different points. For each measurement, we compute the estimated user location based on our fingerprinting, using the beam index and the reported TA. First, the reported beam index indicates the azimuth angle of our location estimate, i.e., the line that starts from the BS and splits the beam area in half. When there are multiple candidates (e.g., beam 30 in Figure 3), we distinguish the areas based on our fingerprinted TA distribution (note that, when reflected, a beam has longer propagation time and therefore higher TA value). Second, we translate the reported TA value to an estimate of the UE distance from the BS using the formula $T_{TA} = f(TA)$ [23, Section 4.3.1]. Based on the cell configuration in our experiment, one TA increment corresponds to a distance of 9.77 meters. Due to the discretized design of TA values, instead of a single distance estimate we obtain the range of distances $[f(TA), f(TA+1)]$, which defines an annulus (or ring shape) around the BS. We compute the line segment that is the intersection of the beam line and the TA annulus, and output its midpoint as our location estimate. Finally, we compute the localization error as the distance between our estimate and the ground

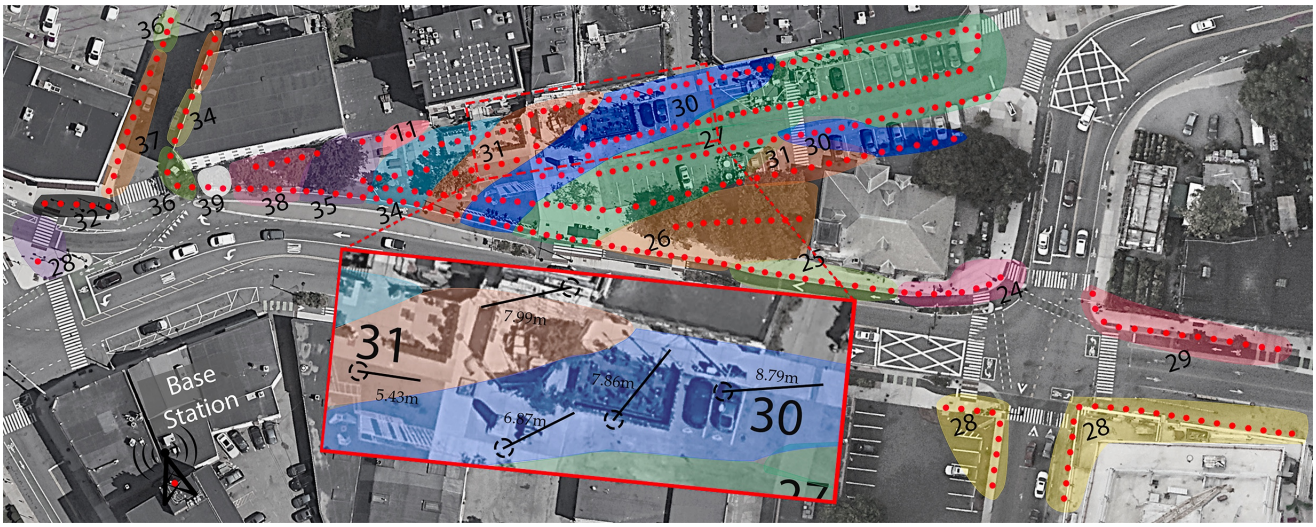

Figure 3: Fingerprinting of static beams within a cell. Red points denote the GPS locations of every measurement and the BS. Colored overlays denote the fingerprinted beam locations.

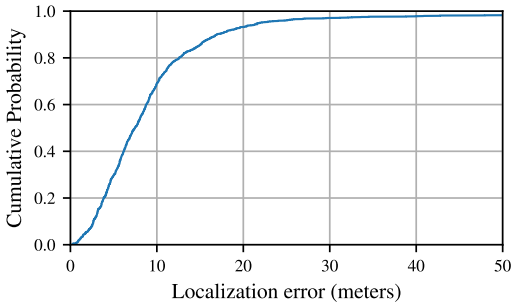

Figure 4: Empirical cumulative distribution function of the SSB-RA localization attack error across all data points.

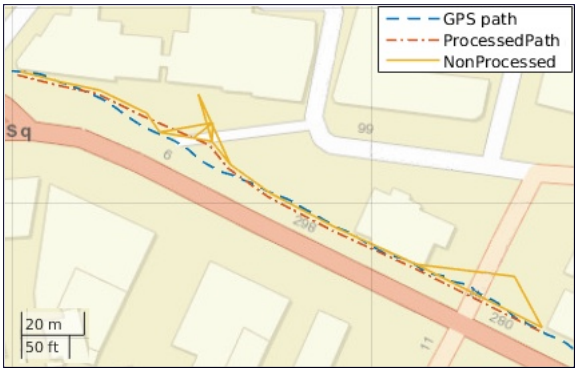

Figure 5: Comparison of the GPS path and the estimated path using CSI Reports, with and without our filtering algorithm.

truth location, and show the ECDF of the error in Figure 4. Our results show that an attacker can localize a user with an error of less than 10 meters 70.34% of the time, and less than 20 meters 94.28% of the time, regardless of their location in the cell, by passively sniffing the RA message exchange. Only 2% of our collected data points fail localization due to missing or inconclusive beam indexes in our fingerprinting. Additionally, we quantify the distribution of all reported TA data across measurements for each location. We obtained a total of 1605 TA reports from all three phones that differ by 0.15 from their location-specific mean TA value, with a standard deviation of 0.87. This indicates a consistent distribution of TA values across the three phones for the same locations.

User movement tracking with CSI reports. To validate and evaluate our tracking attack by leveraging leakages from CSI reports, we connect all the COTS UEs (targets of the attack) to the operator network and walk in different path patterns, including changes of direction, covering the area that the mmWave cell serves. Our total dataset comprises more than 3.5 km in path lengths ranging from 20 meters to 150 meters, with a total of 60 paths that cover all the pedestrian area. We log both the GPS position and the CSI reports sent by the UEs from the chipset. We use the sequence of CSI reports as input to the algorithm described in Section 3.1.5, which outputs the estimated path that the UE traversed. Figure 5 displays the GPS (ground truth) path alongside our estimated path derived from CSI reports, both before and after processing by our algorithm. We find that the non-filtered paths present high fluctuations, especially close to areas where the density of the beams increases. More importantly, we find that using the raw CSI reports results in rapid alternations between points when the user is transitioning from one beam to the next. This makes the non-filtered path considerably larger than the ground truth distance, and does not reflect the actual path taken by the phone. In contrast, the filtered path closely aligns with the GPS path, as shown in Figure 5, remaining within its boundaries and allowing an attacker to effectively track a specific user's movements within a cell. In order to evaluate the path estimation accuracy, we compute the maximum distance deviation between the estimated path and the ground truth GPS

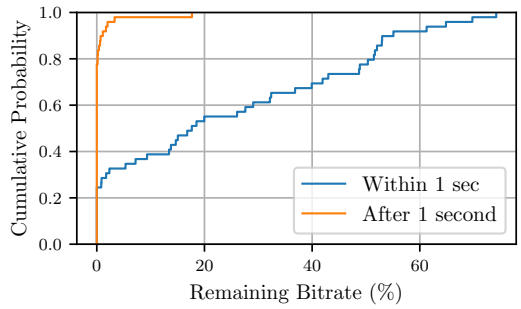

Figure 6: Empirical cumulative distribution function of Bitrate DoS with DCI Spoofing.

path for all our recorded paths. Our algorithm achieves a maximum path deviation lower than 15 meters in 80.88% of the paths, with an average path deviation of 9.92 meters across all paths. Moreover, we measure the CSI report configurations in cells supporting beamforming across six MNOs spanning three different countries. We find that, to support beamforming, all operators use periodic CSI configurations, every 20 to 40 ms, which provide fine granularity.

4.2.2 ***Over the air injection in 5G NR.*** Previous work demonstrated that it is possible to spoof physical signals to connected UEs in LTE [12, 14]. We carry out the first demonstration of 5G NR over the air signal spoofing attacks of PDCCH and PDSCH channels. Our analysis and results reveal that physical signal spoofing in 5G NR is more efficient due to changes in the 5G physical layer and has relaxed time synchronization requirements. Namely, 5G reduces the number of control and always-on signals. In this way, 5G does not transmit a PCFICH with each subframe, or continuous pilots as LTE for channel estimation. Instead, it only transmits pilots when transmitting data, which constrains the transmission to a small bandwidth in frequency (as low as 1.08 MHz), compared to the required wide bandwidth in 4G. This allows an attacker to inject PDCCH and PDSCH signals in a standalone way, and carry out opportunistic injection in empty pockets of spectrum, as constantly overshadowing pilots as in LTE to reflect the correct attacker channel is no longer needed [14]. We verify OTA injection of DCIs and MAC CEs using an USRP X310 SDR against COTS UEs connected to a 5G gNB testbed in an anechoic box. The attacker generates locally the IQ samples containing PDCCH/PDSCH signals addressing UEs by their RNTI, and transmits the signals OTA using an SDR. We provide a demonstration video of DCI (DL/UL grants, along with TPC, and PDCCH Order) and MAC CE injection against a COTS UE[1]. We highlight that the changes in the 5G PHY allow, not only synchronized injection, in the same way as 4G, but also non-synchronized injection due to the relaxed requirements. To showcase this, we transmit IQ samples containing a DCI repeatedly to a specific user, and we measure the probability of successful DCI injection. We find that, even without time synchronization, the DCI is successfully injected at a rate of 7.54 DCIs per second on average when the power is 3dB above the gNB.

4.2.3 **DCI Spoofing.** We evaluate our DCI spoofing attacks, that lead to UE collisions, service degradation, and denial of service. We craft DCI messages and inject them to active UEs, which cannot detect the injections since the messages are not integrity protected.

Tricking legitimate UEs to jam other UEs. To validate the theoretical attack, we initially inject DCIs with UL grants of various sizes to active UEs that do not have any data to transmit, and are not requesting any resources. We find that the UEs use all the resources allocated to them, and fill the allocated space with padding. We setup the attack evaluation by having two COTS UEs connected to our 5G testbed. The first UE generates Iperf3 traffic on the uplink at maximum bandwidth, and will be the target of our attack. The second UE, referred to as Induced Jammer UE (IJ-UE), will be tricked to jam. We inject a crafted DCI message to the IJ-UE and set the UL Grant allocation to include all the physical resource blocks of the PUSCH during every time slot. Additionally, we set the TPC parameter to the highest value, which instructs the UE to increase the transmit power to the maximum. We monitor the throughput reported by Iperf throughout the duration of the attack over 1-second intervals and we measure the average throughput: before the attack, during the first second after the attack, and during the remaining duration of the attack after the first second. Figure 6 shows the ECDF of the affected throughput as a percentage of the original throughput of the UE before the attack. We notice that 77% of the time the target UE achieves less than 50% of its original throughput within one second of the attack and less than 0.1% throughput for the remainder of the attack. Therefore, an adversary is able to trick UEs against each other, blocking all communications, at the cost of spoofing very small DCI messages.

HARQ-Attack. We validate the attack described in Figure 1 against a COTS UE, by injecting a DCI with an out-of-sequence DAI, and analyzing the reported ACK bitmap. We find that using a counter n above the current counter value, the UE reports an ACK report with n more bits than the base station expects, accounting for the n missed transmissions. For instance, by injecting a DAI with value 2, the UE reports 3 bits, with two NACKs corresponding to DAIs 0 and 1. To evaluate our attack, we connect our COTS UE to our 5G srsRAN BS, and we spoof DCIs with a DAI value above the current DAI. We observe that the base station is unable to match the received ACK bitmap to the correct transmissions, and reaches HARQ failure. Consecutive missed ACKs lead to a radio link failure after two seconds, and breaks the connectivity between the COTS UE and the base station. For ethical reasons, we do not evaluate the attack on an operator network, but we highlight that bitmap mismatches are not handled by the 3GPP standard.

4.2.4 **Carrier aggregation attacks.** We configure our LTE BS to operate over two 5 MHz carriers in LTE bands 3 and 7. We connect our COTS UE with enabled CA over two carriers. We generate downlink TCP traffic using iPerf3 and monitor the throughput between the UE and BS. We then spoof a SCell Deactivation DL MAC CE to the COTS UE at time instant 3 seconds, instructing the UE to deactivate the SCell. The measured iPerf TP is depicted in Figure 7. The UE maintains a throughput of 10.5 Mbps using CA for the initial 3 seconds, but this value rapidly decreases to as low as 1.04 Mbps, at the time of the attack. The observed TCP retransmissions and the throughput drop to below 50%, validate

[1]Video link: Dropbox link

	MNO	SCellDeact.	Max CA	Max RA	RARWind
Country A	MNO 1	Infinite	3CA	n10	sf10
	MNO 2	Infinite	4CA	n5	sf10
	MNO 3	Infinite	5CA	n10	sf10
Country B	MNO 1	Infinite	3CA	n10	sf10
	MNO 2	Infinite	3CA	n10	sf10
Country C	MNO 1	Infinite	3CA	n10	sf10

Table 1: Observed values of LTE RRC parameters related to CA and RA from six major MNOs across three countries.

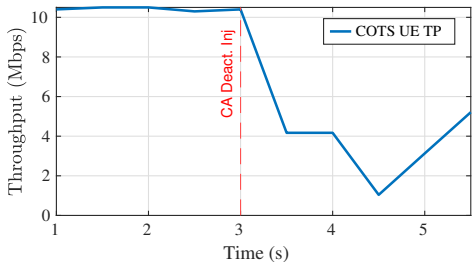

Figure 7: Measured Iperf3 throughput (TP) from a UE during a SCell Deactivation MAC CE injection at t=3.

that the UE no longer monitors the secondary cell that the BS is still using to transmit iPerf traffic.

To assess the efficiency of our attacks in real-world scenarios, without targeting deployed networks, we analyze the current configuration of three MNOs serving over 100 million subscribers each. We use QXDM to extract the configuration parameters of interest for RA from SIB and CA from the RRC by connecting a phone with the MNO's SIM card, and include our findings in Table 1.

One factor that limits the impact of MAC CE injection attacks is the presence of timeouts for various UE states, such as the `sCellDeactivationTimer`. Our goal with the SCell activation attack is to keep the attacker-enforced SCells active for an extended periods, reducing the need for frequent injections. As shown in Table 1, all three tested MNOs omit this timer value in the RRC configuration, defaulting to infinity. This simplifies control for the BS, granting it full control of SCell activation/deactivation via MAC CEs, but increases UEs vulnerability and leads to significant battery drain. We experimentally confirm that UEs activate configured SCell receiving a spoofed SCell Activation DL MAC CE. Additionally, we find that all tested MNOs support 3-5 aggregated carriers, and prior work shows that even one active secondary carrier increases power consumption by an average of 79% [19].

5 Mitigations and Conclusions

Low-layer encryption. A possible mitigation is to derive a physical-layer key K_{PHY} post-authentication, similar to RRC-layer keys [5, p. 48], and modify existing scrambling blocks of physical channels (e.g., PDCCH/PUSCH) to cryptographically encrypt control information and CSI reports with sequences derived from K_{PHY}, system frame number `SFN`, subframe `sf`, and resource block `RB` indices.

Lightweight message authentication. Alternatively, digital signatures can be incorporated within Non-Critical Extensions (NCE) of SIBs, protecting critical parameters. Additionally, integrity checks for MAC CEs/DCI messages can be integrated into protected upper layers, with some overhead and latency trade-offs.

Conclusions. We systematically analyzed the attack surface in 4G/5G physical channels, which lack encryption and integrity protection. In this paper, we experimentally demonstrate OTA spoofing attacks in 5G, active PHY/MAC attacks against COTS UEs in an isolated testbed, and passive high-accuracy localization exploiting mmWave beam management leakages in commercial networks.

References

[1] H. Kim, J. Lee, E. Lee, and Y. Kim, "Touching the untouchables: Dynamic security analysis of the LTE control plane," in *2019 IEEE Symposium on Security and Privacy (SP)*, 2019.
[2] C. Yu, S. Chen, Z. Cai, and J. Díaz-Verdejo, "LTE Phone Number Catcher: A Practical Attack against Mobile Privacy," *Security and Comm. Networks*, 2019.
[3] A. Shaik, R. Borgaonkar, N. Asokan, V. Niemi, and J.-P. Seifert, "Practical Attacks Against Privacy and Availability in 4G/LTE Mobile Communication Systems," in *NDSS*, 2016.
[4] D. Rupprecht, K. Kohls, T. Holz, and C. Pöpper, "Breaking LTE on layer two," in *IEEE Symposium on Security & Privacy (SP)*, 2019.
[5] "5G; Security architecture and procedures for 5G System (3GPP TS 33.501 version 16.9.0 Release 16)," 2022.
[6] N. Bui and J. Widmer, "OWL: A Reliable Online Watcher for LTE Control Channel Measurements," in *Proceedings of the 5th Workshop on All Things Cellular: Operations, Applications and Challenges*, 2016.
[7] R. Falkenberg and C. Wietfeld, "FALCON: An Accurate Real-Time Monitor for Client-Based Mobile Network Data Analytics," in *2019 IEEE Global Communications Conference (GLOBECOM)*, 2019.
[8] S. Kumar, E. Hamed, D. Katabi, and L. Erran, "LTE Radio Analytics Made Easy and Accessible," in *Proceedings of the ACM Conference on SIGCOMM*, 2014.
[9] N. Ludant, P. Robyns, and G. Noubir, "From 5G Sniffing to Harvesting Leakages of Privacy-Preserving Messengers," in *2023 IEEE Symposium on Security and Privacy (SP)*, 2023.
[10] T. D. Hoang, C. Park, M. Son, T. Oh, S. Bae, J. Ahn, B. Oh, and Y. Kim, "LTESniffer: An Open-Source LTE Downlink/Uplink Eavesdropper," in *Proceedings of the 16th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, 2023.
[11] H. Wan, X. Cao, A. Marder, and K. Jamieson, "NR-Scope: A Practical 5G Standalone Telemetry Tool," in *Companion of the 20th International Conference on Emerging Networking EXperiments and Technologies (CoNEXT)*, 2024.
[12] S. Erni, M. Kotuliak, P. Leu, M. Roeschlin, and S. Capkun, "Adaptover: Adaptive overshadowing attacks in cellular networks," in *Proceedings of the 28th Annual International Conference on Mobile Computing And Networking*, 2022.
[13] M. Kotuliak, S. Erni, P. Leu, M. Roeschlin, and S. Capkun, "LTrack: Stealthy tracking of mobile phones in LTE," in *31st USENIX Security Symposium*, 2022.
[14] H. Yang, S. Bae, M. Son, H. Kim, S. M. Kim, and Y. Kim, "Hiding in plain signal: Physical signal overshadowing attack on LTE," in *28th USENIX Security Symposium (USENIX Security 19)*, 2019.
[15] R. P. Jover, "LTE security, protocol exploits and location tracking experimentation with low-cost software radio," 2016.
[16] I. K. Jain, R. Subbaraman, T. H. Sadarahalli, X. Shao, H.-W. Lin, and D. Bharadia, "mMobile: Building a mmWave Testbed to Evaluate and Address Mobility Effects," in *Proceedings of the 4th ACM Workshop on Millimeter-Wave Networks and Sensing Systems*, 2020.
[17] Keysight, "WJ1000A mmWaveJudge Remote Receiver," https://www.keysight.com/us/en/product/WJ1000A/mmwavejudge-remote-receiver.html, 2020.
[18] "5G; NR; Physical layer procedures for control (3GPP TS 38.213 version 16.3.0 Release 16)," 2020.
[19] R. Sanchez-Mejias, Y. Guo, M. Lauridsen, P. Mogensen, and L. A. Maestro Ruiz de Temino, "Current consumption measurements with a carrier aggregation smartphone," in *2014 IEEE 80th Vehicular Technology Conference*, 2014.
[20] "LTE; Evolved Universal Terrestrial Radio Access (E-UTRA); Radio Resource Control (RRC); Protocol specification (3GPP TS 36.331 Release 17)," 2022.
[21] Qualcomm, "QxDM Pro Qualcomm eXtensible Diagnostic Monitor." 2022.
[22] SRS, "Software Radio Systems. Open source SDR 4G/5G software suite," https://github.com/srsran/srsRAN, 2020.
[23] "5G; NR; Physical channels and modulation (3GPP TS 38.211 version 16.4.0 Release 16)," 2021.