NR-Scope: A Practical 5G Standalone Telemetry Tool

Haoran Wan
Princeton University
Princeton, NJ, USA
haoran.w@princeton.edu

Xuyang Cao
Princeton University
Princeton, NJ, USA
xyc@princeton.edu

Alexander Marder
Johns Hopkins University
Baltimore, MD, USA
amarder@jhu.edu

Kyle Jamieson
Princeton University
Princeton, NJ, USA
kylej@princeton.edu

Abstract

NextG cellular networks are designed to meet Quality of Service requirements for various applications in and beyond smartphones and mobile devices. However, lacking introspection into the 5G Radio Access Network (RAN) application and transport layer designers are ill-poised to cope with the vagaries of the wireless last hop to a mobile client, while 5G network operators run mostly closed networks, limiting their potential for co-design with the wider internet and user applications. This paper presents *NR-Scope*, a passive, incrementally-deployable, and independently-deployable Standalone 5G network telemetry system that can stream fine-grained RAN capacity, latency, and retransmission information to application servers to enable better millisecond scale, application-level decisions on offered load and bit rate adaptation than end-to-end latency measurements or end-to-end packet losses currently permit. Our experimental evaluation on various 5G Standalone base stations demonstrates NR-Scope can achieve less than 0.1% throughput error estimation for every UE in a RAN. The code is available at **https://github.com/PrincetonUniversity/NR-Scope**.

CCS Concepts

• **Networks → Wireless access points, base stations and infrastructure**.

Keywords

5G network; Telemetry; Wireless network; Network measurement.

ACM Reference Format:
Haoran Wan, Xuyang Cao, Alexander Marder, and Kyle Jamieson. 2024. NR-Scope: A Practical 5G Standalone Telemetry Tool. In *Proceedings of the 20th International Conference on emerging Networking EXperiments and Technologies (CoNEXT '24), December 9–12, 2024, Los Angeles, CA, USA.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3680121.3697808

1 Introduction

5G Standalone (SA) is becoming the dominant cellular Radio Access Network (RAN) technology, making up 25% of mobile data at the end of 2023 and forcasted to grow to 76% by 2029 [2]. With optimizations in the RAN and the mobile core, 5G SA provides higher throughput, higher capacity and lower latency than 4G, making a rich canvas for NextG novel applications, such as mobile cloud gaming, remote surgery and roadside and vehicular connectivity.

CoNEXT '24, December 9–12, 2024, Los Angeles, CA, USA

ACM ISBN 979-8-4007-1108-4/24/12
https://doi.org/10.1145/3680121.3697808

Wireless networks, however, suffer from highly variable end-to-end throughput due to user movement, channel fading, blockage, and interference. This variable throughput poses performance challenges for many applications, ranging from the congestion control needed for bulk data transfer [7, 10] to interactive video [8, 14, 31]. A natural approach to measuring throughput in 5G SA networks is to send packet probes over end-to-end connections, but the extensive buffering, and link-layer retransmission in the RAN (§5.4) [10, 16] limit the measurement resolution that end-to-end measurements can provide. Furthermore, the closed nature of the RAN inhibits co-design of applications and the cellular network.

These problems motivate the design of a telemetry tool to precisely estimate the throughput, channel condition, and cell's overall bandwidth, helping network application and end-to-end transport designers overcome these obstacles.

In this paper, we present **NR-Scope**, a RAN telemetry tool that decodes the necessary *Radio Resource Control* (RRC) messages and *Downlink Control Information* (DCI) from the 5G SA network. This information allows NR-Scope to determine a cell's configuration, traffic patterns for all of the user equipment (UEs) connected to the cell, and each UE's channel condition. NR-Scope works in tight synchronization with the 5G SA network, where it decodes the DCIs in every *Transmission Time Interval* (TTI), consisting of 1, 0.5, or 0.25 ms in 5G. NR-Scope also acquires the number of bits that are delivered in each TTI—providing an *exact* throughput estimation for each UE in the RAN—so that developers can improve the performance of applications and gain better understanding of the 5G RAN. The result is an open design that operates entirely independently of the 5G network operator and 5G mobile devices.

We have implemented NR-Scope on the USRP software-defined radio platform and evaluated it on three different 5G SA RANs: srsRAN, Aether Onramp 5G, the Amarisoft Callbox 5G, and T-mobile cells. Our experimental evaluation shows that NR-Scope produces highly accurate estimates of radio resource usage with a median throughput error of 1.01 kbps (Onramp) and 0 kbps (Amarisoft) compared to ground truth, and an average of 0.9% overall throughput estimation error. Further results demonstrate NR-Scope also detects packet aggregation, spare RAN capacity, modulation and coding rate, and retransmissions, informing applications of these critical network path changes in real time.

2 Related Work

Cellular telemetry. Most existing telemetry tools are specific to 4G LTE where control information is not scrambled, and so they do not need to decode the RACH process (§3.1.2). Such tools cannot verify the correctness of the decoded information on their own ([12, FALCON], [17, LTESniffer], and [34, NG-Scope]), whereas NR-Scope can (§5.2). LTeye [19] does not handle MIMO, while OWL [9] lacks support for carrier aggregation. In contrast, NR-Scope targets

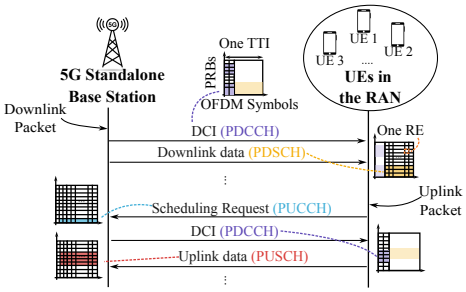

Fig. 1— An overview of downlink and uplink traffic scheduling in the 5G SA RAN.

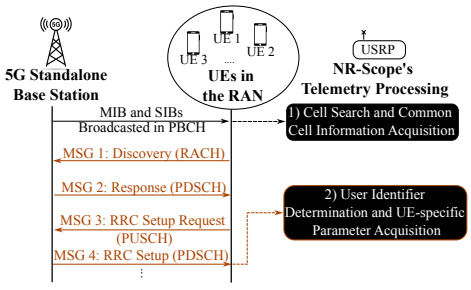

Fig. 2— *NR-Scope UE association tracking:* NR-Scope's steps have a black background.

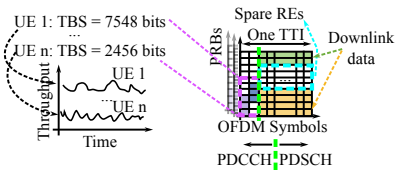

Fig. 3— *NR-Scope UE throughput estimation:* NR-Scope decodes the DCI for each UE within the TTI and estimates the throughput for each UE.

5G SA, with its more configurable control and data channels.

5GSniffer[22] develops a 5G control channel decoder for user activity prediction, however, it requires very detailed, low-level cell information from users, such as DCI format, the control channel's frequency domain position, and channel interleaving patterns, which are very hard for a normal user to acquire. 5GSniffer also only works on Frequency Division Duplexing (FDD) with 15 kHz subcarrier spacing (SCS), but commercial deployments often use Time Division Duplexing (TDD) with different subcarrier spacing (e.g., T-Mobile 5G SA in the U.S.). To the best of our knowledge, no existing open-source telemetry tool can decode 5G SA radio resource control messages and downlink control information for all the UEs in the 5G RAN without any low level cell information required. Industrial telemetry tools such as KeySight WaveJudge [18] and ThinkRF [4] for 5G are closed source, generally require expensive and specialized hardware, and stop short of integration with end-to-end internet protocols and applications, which is the focus of the present work. On the handset, MobileInsight [20] and QXDM [3] capture a variety of telemetry from different layers in the cellular stack from a *single* UE's perspective, making it difficult to infer other users' cell usage. Decoding the highly flexible 5G control channel is still an unmet need of the networking community for new RAN-aware designs, and NR-Scope bridges the gap.

Cellular-aware transport and RAN designs. PBE-CC [35] (working only in 4G) uses NG-Scope [34] to decode the RAN's control channel, estimate its capacity, and update a BBR-like [10] throughput estimation algorithm. ABC [15] marks the ECN bits in IP headers with a modified meaning, but stops short of a cellular implementation. CQIC [21] employs QXDM to decode the channel quality indicator (CQI) of an LTE RAN to estimate a user's potential throughput, but stops short of full RAN telemetry. piStream [32] and CLAW [33] utilize resource blocks from signal strength measurements to improve video workloads and web browsing workloads, respectively, but signal strength's predictive power is limited, as shown in their results. RadioSaber [11] uses physical channel information to perform better slicing management inside the 4G core, modifying the core stack. All of the above-mentioned transport protocols take 4G as their context, while NR-Scope instead provides 5G SA cell telemetry to future transport protocols and applications.

3 Design

The high-level goal of NR-Scope is to decode and interpret DCI, fine-grained telemetry information that the 5G RAN broadcasts onto the airwaves. The DCI also tells each UE where in the 5G SA physical data channels to receive or send its data, and thus how much data is actually present on the downlink and uplink channels. Furthermore, NR-Scope decodes all the unencrypted RRC messages delivered from the gNB to the UEs, which are not provided in previous tools [9, 34], providing us thorough layer 2/1 cell configurations, such as number of MIMO layers used for each DCI, re-transmission patterns and slot allocation for downlink/uplink transmission.

Preliminaries. The RAN divides time and frequency into *Physical Resource Blocks* (PRBs), equally-sized frequency slices that carry information to or from the base station, also known as a *cell*, or *gNodeB* (gNB). Time is divided into *slots*, or TTIs which are also the units of downlink-uplink transmission switching. Within each slot, there are fourteen OFDM symbols. One PRB and one OFDM symbol make up the *Resource Element Group* (REG) – minimal scheduling unit within one TTI. Unlike 4G, 5G has three possible subcarrier spacings including 15 kHz (same as 4G), 30, and 60 kHz, which result in TTIs of 1, 0.5, and 0.25 ms.

Two key channels contain the relevant information: the *Physical Downlink Control Channel* (PDCCH), for control traffic from the gNB to the UE, and the *Physical Downlink Shared Channel* (PDSCH) for traffic from gNB to UE. DCI information in the PDCCH points to PRBs in the PDSCH or *Physical Uplink Shared Channel* (PUSCH), containing both cell configuration information and network traffic. Fig. **1** shows the scheduling process of gNB and UE when there is downlink or uplink traffic, with numbers indicating the order: we decode both downlink and uplink scheduling information by decoding the PDCCH. Since NR-Scope only needs to decode the downlink control channel for both downlink and uplink DCIs (scheduling result), UE and uplink channel status have no effect on it.

The *Cell Radio Network Temporary Identifier* (C-RNTI) is a unique identifier used in the RAN to identify a specific mobile device or UE. The cell assigns the UE a C-RNTI through the *Random Access Channel* (RACH) Process, where the UE and gNB exchange information to establish an RRC session between UE and gNB [28]. RRC is a Layer-3 protocol whose major functions include establishing and releasing connections, as well as broadcasting system information.

3.1 UE Association Tracking

Unlike 4G, 5G DCI is scrambled with C-RNTI derived sequence, and we can't decode the DCI without knowing it. We track UE's association for C-RNTI. Our UE association tracking has two steps: **1)** cell search and common cell parameter acquisition and **2)** user identifier

determination and UE-specific parameter acquisition (Fig. **2**).

3.1.1 Cell Search and Parameter Acquisition. The goal of this step is to mimic how a UE discovers a 5G cell and extracts its channel configuration parameters, which describes the structure of the channel that a UE uses to associate with the cell (§3.1.2).The base station periodically broadcasts basic information about itself in a *Master Information Block* (MIB), which includes a time index (the *system frame number*[1]) and configuration parameters of initial association. A UE would use the information in the MIB to locate PRBs that contain the *Control Resource Set* (CORESET) 0, a part of the PDCCH. CORESET 0 carries a DCI that points to PRBs in the PDSCH containing *System Information Block* (SIB) 1 in the PDSCH (lower left corner of Fig. **2**). 5G SA cells broadcast the initial signals for UE, so our tool can decode them without extra effort. To note, such information is not broadcasted under NSA 5G mode but provided through the anchored LTE cells under encryption. In such case, NR-Scope requires manual input of 5G cell information.

SIB 1 carries common[2] information about the cell, including physical channel configuration and all the information a UE may need for the RACH processing described next in §3.1.2. This information includes the subcarrier spacing used in the RACH process, the parameter and time-frequency position for MSG 1 in RACH, and the PDCCH parameters, obviating the blindly searching [12, 34].

3.1.2 User Identifier and Parameter Discovery. The goal of this step is to acquire users' C-RNTI, and the parameters for the channel containing the telemetry information. NR-Scope must decode the RACH process to acquire the C-RNTIs, as all DCIs after the RACH are scrambled by a sequence derived from the C-RNTIs.

There are four messages exchanged between the gNB and a UE (Figure **2**): Preamble Transmission (MSG 1), Random Access Response (MSG 2), RRC Setup Request (MSG 3), and RRC Setup (MSG 4). SIB 1 tells the UE where and when to transmit MSG 1 to the cell, which begins the process.[3] MSG 4 contains most of the UE-specific information required for mobile communication and for telemetry, namely the PDCCH for the UE. NR-Scope only needs to decode MSG 4, whose DCI are plain text with a TC-RNTI-scrambled CRC of the DCI appended, the same as in 4G [5]. Knowing this, we can calculate the C-RNTI by an exclusive-or of a CRC of the received DCI plain text computed by NR-Scope, and the appended C-RNTI-scrambled CRC of the DCI, separately received by NR-Scope, as previous 4G sniffers [12, 34] do for each and every received DCI. NR-Scope then obtains the C-RNTI for the UE when the TC-RNTI is promoted to the C-RNTI after MSG 4—without needing to decode the preliminary messages—and we can verify the correctness of decoded MSG 4 through CRC check. From MSG 4, we also get the CORESET position, DCI aggregation level, and the correct format of DCI that the gNB and UE will use for communication.

Decoding RACH consists of detecting the DCI that schedules

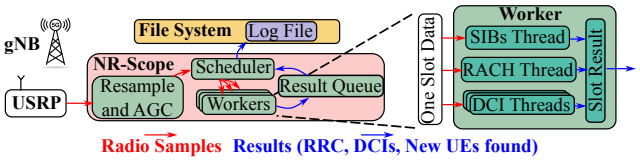

Fig. 4— The system processing pipeline of NR-Scope.

RRC Setup information and using it to decode the PDSCH that contains the actual RRC Setup message data (up to 500 bytes). Decoding one RRC Setup takes around 1 to 2 ms for signal processing, thus decoding it for every new UE is expensive as the slot only lasts 0.5 ms. According to our observation, the RRC Setup is identical among UEs, thus we can skip decoding the PDSCH and only use the DCI for new UE discovery. Unlike the SIBs, each UE only gets one RRC Setup for its initial connection, and if we miss a RACH, we can't decode the DCIs for this UE in the telemetry session. Since most of the UEs stay in the RAN for a short time (§5.3), the throughput estimation error from a missing RACH is not too dramatic.

3.2 Network Telemetry

Next, NR-Scope extracts telemetry information and estimates RAN capacity (Figure **3**). Within one TTI, we decode the DCIs in PDCCH for each UE, through which we can acquire their downlink and uplink throughput, channel condition and retransmission.

3.2.1 Telemetry Information Extraction. The goal of this step is to estimate the amount of data traffic in the data channel (PDSCH) for each UE in the RAN, and its associated physical layer parameters (bit rate and channel quality). Each UE also gets the DCI format type and aggregation level from MSG 4 [6]. Since we already decoded the required information, NR-Scope moves to the same *bandwidth part* (BWP, a fraction of the whole bandwidth) as the UE for DCI reception. With all required information known (C-RNTI, DCI format), NR-Scope receives DCI through the standard 3GPP DCI decoding process [26, 28], which yields the 30–80 bits of DCI data.[4]

3.2.2 RAN Capacity Estimation. In this step, NR-Scope uses telemetry information to calculate the capacity (bit rate) allocated to each UE and the spare RAN capacity. With the DCI and RRC messages decoded in prior steps, we calculate the *Transport Block Size* (TBS), which indicates how many bits are transmitted for the specific UE in this TTI. We calculate the TBS as in the 3GPP standard (see Appendix **A**). We record the TBS for every UE in each TTI, maintaining a sliding window to calculate the bit rate for each UE.

Due to wireless fading, the UE may fail to decode the downlink traffic, triggering re-transmission. NR-Scope captures the re-transmission information in the physical and MAC layers through tracking Hybrid Automatic Repeat Request (HARQ) information in DCIs. The gNB allocates up to 16 HARQ processes for each UE, informing the UE of each process through `harq_id` in the DCI. If the UE correctly decodes the data in one TTI and sends back an ACK, the gNB toggles the `new_data_indicator` of the DCI with the same `harq_id` to indicate new data. If the UE NACKs, the gNB uses the same `ndi` for the re-transmission. NR-Scope maintains an

[1] One system frame is 10 ms, and the system frame number (0–1024), which the UE will also need to synchronize, indexes each system frame.
[2] In 3GPP standard, *Common* configuration is for all the UEs in the RAN, while *Dedicated* configuration is for the specific UEs that receive the message.
[3] MSG 2 contains a temporary cell RNTI (TC-RNTI) and a grant for the UE to transmit MSG 3; MSG 3 contains an RRC Setup request to the gNB; the UE and gNB use the TC-RNTI for MSG 4, after which the UE is RRC connected to the gNB and the TC-RNTI becomes a C-RNTI.
[4] This includes deinterleaving and the modulation and coding scheme indicator, which are required for transport block size (TBS) calculation and our later capacity estimation. A sample of DCI with format 1-1 and its translated downlink grant for PDSCH can be found in Appendix B.

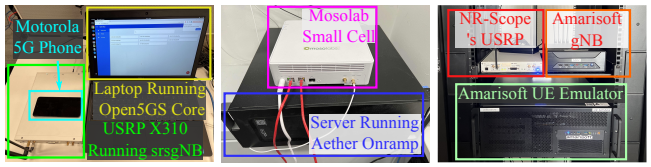

Fig. 5— Hardware used in the NR-Scope evaluation.

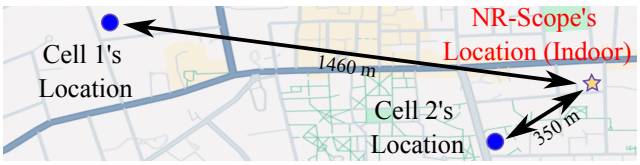

Fig. 6— Commercial 5G cells used in NR-Scope evaluation.

array for each UE to record the `ndi` from previous DCIs for each `harq_id` to detect re-transmissions.

4 Implementation

The system overview is shown in Fig. **4**, where each block inside the NR-Scope and worker block is an asynchronously running component. The USRP receives signal from the gNB and passes the radio samples to NR-Scope. NR-Scope may need to resample the samples to fit the FFT bins onto the subcarriers[5], use *automatic gain control* (AGC) for better signal strength, and passes the slot data to the scheduler. The scheduler copies the data and its state (known UE list, cell's configurations) to an idle worker. For each slot data, the worker spawns SIBs thread, RACH thread and DCI threads for SIBs decoding, UE discovery and DCIs extraction, and then put the slot result into the result queue. The scheduler tries to collect results from the result queue, writes the results to a log file, and updates its state (RRC messages from the cell and new UEs).

We implement NR-Scope's telemetry functionalities in around 4,000 lines of C++ code (excluding reused code) processing radio signals received by a USRP. We reuse the physical layer signal processing modules from an open-source 5G library srsRAN [27], *i.e.*, a wireless channel estimator, a demodulator, and a frame synchronizer. srsRAN only supports processing for FDD base station with 15 kHz subcarrier spacing, thus we modify its low level processing code extensively, such as Fourier transform scheduling and phase compensation, to support different subcarrier spacing, higher bandwidth and TDD base station processing. The major computational cost comes from the FFT of each slot, demodulation and CRC check for all the detected UEs in the RAN. To catch up with the TTI, NR-Scope employs multiple threads in DCI extraction, UE list is sharded among threads, and the final results are gathered from the threads. Furthermore, the worker pool design enables asynchronous, on-demand slot data processing, further lowering the host CPU requirement if we don't need the real-time DCI output (§5.3).

5 Evaluation

We first introduce overall methodology (§5.1), then evaluate NR-Scope's telemetry in DCI decoding, PRB detection and flow throughput estimation accuracy (§5.2), and processing time and coverage (§5.3), finally explore its use cases (§5.4).

5.1 Methodology

We evaluate NR-Scope in the following two real 5G Standalone networks, as shown in Fig. **5** and 6. In srsRAN, Mosolab and T-Mobile networks, Motorola *Moto G* 5G phones serve as the clients, while in Amarisoft's network, one UE emulator is used.

[srsRAN/Open5GS]: In this open-source 5G Standalone network, we run Open5GS on a laptop, and an srsRAN gNB [27] with a X310 USRP. The gNB runs on 5G band n41 in TDD mode, with downlink center frequency 2524.95 MHz, 30 kHz SCS, and bandwidth 20 MHz.

[Mosolabs/Aether]: In this Private 5G Standalone network, we configure an Aether Onramp 5G software defined core network [25] on a server machine, and a Sercomm Mosolabs small cell [24]. The gNB runs in the CBRS band (n48) in TDD mode, controlled by a spectrum access server (SAS) [13], on center frequency 3561.6 MHz, 30 kHz SCS, and bandwidth 20 MHz.

[Amari Callbox]: In this industrial cellular test equipment, we use its build-in core and gNB [1], which transmits the wireless signals with its build-in SDR. The gNB runs in band n78 in TDD, on center frequency 3489.42 MHz, 30 kHz SCS, and bandwidth 20 MHz.

[T-Mobile Cells]: There are 2 T-Mobile commercial cells near our lab, with cell 1 running in band n25 (FDD, 15 kHz SCS and 10 MHz bandwidth, 1989.85 MHz center frequency) and cell 2 in n71 (FDD, 15kHz SCS and 15 MHz bandwidth, 622.85 MHz center frequency). Both commercial cells use BWP 1 for communication, while all the others above use BWP 0. The distances between NR-Scope and the two cells are shown in Fig. **6**,[6] showing the practical effective range.

5.2 RAN Telemetry

We evaluate NR-Scope telemetry time series data, DCI miss rates, PRB decoding accuracy, and RAN throughput estimation accuracy.

5.2.1 Decoding accuracy. We evaluate telemetry accuracy both via DCI miss rate and PRB decoding accuracy. For this experiment, we use the [srsRAN/Open5GS] fully open-source network. This enables us to collect detailed physical layer *ground truth* for all UEs from srsRAN's log, in terms of TTI index, DCI content and downlink grants (frequency and time resource allocation, and TBS). We match the number of DCIs captured by NR-Scope and srsRAN's log using the timestamp and the TTI index, through which we calculate a DCI decoding *miss rate*. We repeat each evaluation ten times; each observation period lasts approximately ten minutes.

Firstly, we test the DCI miss rate for each UE, we connect UEs into the srsgNB, and in the meantime use NR-Scope to decode the DCIs. Then we match the DCIs decoded by NR-Scope and DCIs in srsgNB's log based on the timestamp and TTI indexes to calculate the miss rate for each UE. We perform the same evaluation with Amarisoft network, with up to 64 UEs. NR-Scope achieves a very low DCI miss rate (Fig. **7**), it detects the vast majority of downlink and uplink DCIs, with miss rates of 0.33% and 0.28% in srsRAN and 0.93% and 0.31% in Amarisoft network: two 9's of reliability.

Secondly, we test REs decoding correctness in the same settings. We compare the downlink and uplink grants decoded by NR-Scope

[5]The resampling process is only needed for USRP with twinRX daughterboard, we use CBX-120 daughterboard for evaluations in all cells other than T-Mobile cells and mixed of CBX-120 and TwinRX for the T-Mobile cells evaluation.

[6]The cell towers' location information is provided by https://www.cellmapper.net/.

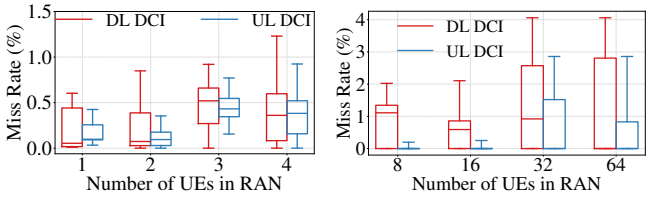

(a) srsRAN with phones as UEs. **(b)** Amarisoft gNB and UEs.

Fig. 7— *DCI miss rate* across 10 minutes experiments.

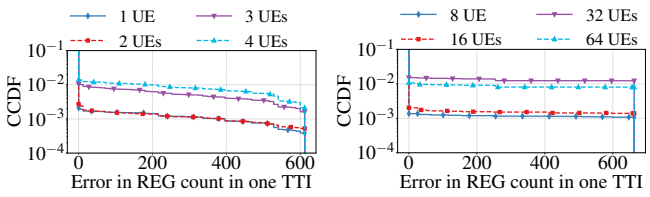

(a) srsRAN with phones as UEs **(b)** Amarisoft gNB and UEs.

Fig. 8— *RE decoding errors* across 10 minutes experiments.

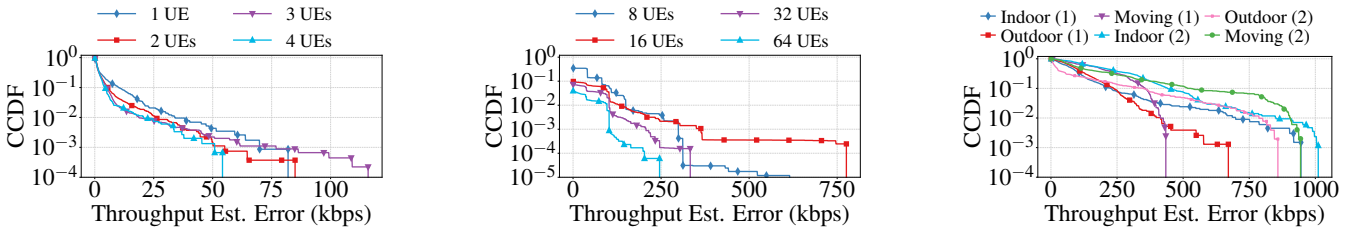

(a) UEs in Mosolab cell. **(b)** UEs in Amarisoft cell. **(c)** UE in T-Mobile Cells: UE status (cell).

Fig. 9— *Accuracy of NR-Scope* throughput estimation (ground truth: tcpdump for Mosolab and T-Mobile, log for Amarisoft cell.)

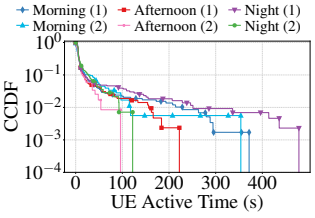

Fig. 10— UEs' active time in T-Mobile cells (Time (cell).)

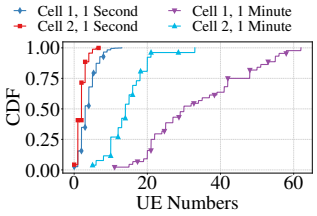

Fig. 11— Number of active UEs per second or minute.

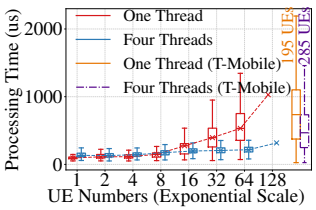

Fig. 12— Processing time with one or four threads.

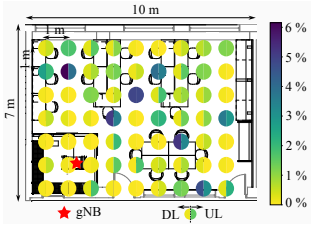

Fig. 13— DCI miss rate across the floor (64 UEs).

and the corresponding grants in the srsgNB log. After we match the DCIs, we compare the decoded number of REGs with the ground truth within each TTI (0.5 ms with 30 kHz subcarrier spacing) and calculate the difference. Fig. **8** shows that NR-Scope achieves an average of 0.77 REG estimation errors per TTI, and most of the time (> 99%), the REG estimation error is zero. The number of REGs allocated for each UE can be up to several hundreds, and so errors of this level will not materially affect the bit rate estimation, as we show next. This demonstrates that NR-Scope can decode physical layer PRB information with high accuracy and granularity.

5.2.2 Throughput estimation accuracy. Here we evaluate the downlink throughput estimation of NR-Scope, in the [Mosolabs/Aether], [Amari Callbox], and [T-Mobile Cells] network. In the commercial small and T-Mobile cell we do not have access to the physical layer ground truth as we did previously, and so for *ground truth* we instead run tcpdump [29] on the phone to capture network packets, calculate the bit rate and compare the results. We connect different numbers of UE into the small cell, which use the data to watching videos or downloading files, and at the same time we use NR-Scope to decode their DCIs. We test NR-Scope under different UE usage scenarios, including static, moving and blocked. In the Amarisoft cell, we extract the gNB's log as the ground truth.

NR-Scope achieves a highly accurate bandwidth usage estimation: Fig. **9(a)** shows the results: NR-Scope estimates the UEs' throughput with 75th percentile errors of 2.33 Kbits/second across static, blocked, and moving scenarios, respectively[7]. In the Amarisoft cell, the 95th percentile error is 35.856 Kbits/second. In the two commercial [T-Mobile Cells], NR-Scope's median throughput estimation errors are 42.56 Kbits/second. The average downlink bit rate for all UEs is 3.35 Mbit/second, 5.73 Mbit/second, and 4.88 Mbit/second in Mosolab, Amarisoft, and T-Mobile cells in this evaluation, so the majority bit-rate estimation errors are under 0.9%.

5.3 Micro-Benchmarks

5.3.1 T-Mobile Cell UE and RAN Behavior. We provide a measurement of UE active time in the [T-Mobile Cells] in Fig. **10**. For each time of day, we measure each cell for 10 min. with three repetitions. We detect 400 to 600 distinct UEs in cell 1 in different time of data and 100 to 200 distinct UEs in cell 2. According to the data, 90 percent of UEs stay in the RAN for less than 35 seconds, showing an unique "come-and-go" cellular network pattern. Furthermore, the number of UEs that the gNB schedules per second and minute is shown in Fig. **11** – less than 60 UE most of one minute period.

5.3.2 Processing time. We demonstrate that NR-Scope can operate in scale. We evaluate the signal processing time and DCI translation time of NR-Scope in one or four DCI threads with different number of UEs in [Amari Callbox] (20MHz bandwidth), and one of [T-Mobile Cells] (10MHz bandwidth). The computation consists two parts: signal processing and DCI decoding for each known UE. The signal processing has a complexity of $O(nlogn)$ for FFT and demodulation, where n is the one slot sample number, growing with bandwidth. The DCI decoding has a growing trend of $O(m)$, where m is the

[7]Figures for all three UE conditions is shown in Appendix C

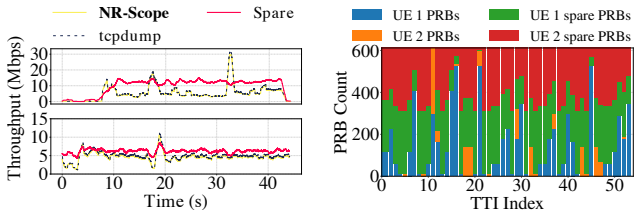

(a) Bit rate estimation for UE 1 (up) and UE 2 (down).

(b) Used REs and fair share REs estimated by NR-Scope for two UEs.

Fig. 14— *NR-Scope spare capacity estimation:* We calculate the bit rate and split the unused REs equally among UEs.

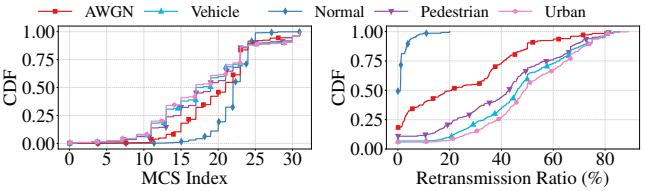

Fig. 15— *NR-Scope telemetry for MCS and retransmission detection:* The UEs are emulated with different channels.

number of UEs, forming a $O(nlogn + m)$ complexity. Fig. **12** shows the linear trend with UE numbers. For the T-Mobile cell with more UEs, NR-Scope can keep up with more workers (§4).

5.3.3 Telemetry Coverage. Here we evaluate the performance of NR-Scope across the floor. We set up the USRP in different locations to measure 64 UEs in Amarisoft network to acquire DCI miss rate, shown in Fig. **13**. Mostly, DCI miss rate is near zero, and the miss rate slightly increases when the signal quality is bad. Coverage is jointly affected by the gNB's transmitting power (§5.2.1) and the USRP's receiving signal quality. Operational cells have a higher transmission power for better coverage, and our evaluation in the [T-Mobile Cells] with distances 1460 and 350 meters (§5.2.2) demonstrates NR-Scope's practicality in commercial cells. Users can find a location with good signal quality and stay there for all the telemetry.

5.4 Use Cases

5.4.1 RAN Spare Capacity Estimation. We acquire the whole bandwidth of the cell through RRC Setup (§3). In each TTI, we can split unused REs evenly across UEs and recalculate these REs to yield a fair-share spare capacity attributable to each UE. During this demonstration, we connect 2 UEs to the Mosolab network, then use both NR-Scope and tcpdump [29] on each phone to calculate each UE's downlink throughput. As shown in Fig. **14**, the capacity estimation for each UE is highly accurate in time (the curve tracks just under ground truth). Turning to Fig. **14(a)** we see NR-Scope calculates the spare bit rate for each UE through the fair share of spare REs. Although the number of fair share REs are the same for each UE, the calculated spare bit rates are different because two UEs have different modulation and coding rates in the same TTI. This feedback could be delivered to app servers without involving the RAN, which as the bottleneck usually incurs large latency.

5.4.2 MCS and Retransmission Behavior. Here we connect 64 UEs to the Amarisoft RAN, using the channel simulator to simulate, AWGN, Vehicle, Pedestrian and Urban channels, according to 3GPP standards. We can observe UE's MCS index and retransmission behavior in different channel conditions in Fig. **15**. The MCS indexes in different channel conditions show that gNB tends to use higher MCS index with a lower retransmission ratio in better channel conditions (normal and AWGN). With channel condition feedback from NR-Scope, service provider can adjust sending strategy accordingly. The coefficient of determinations between NR-Scope's result and the ground truth are 0.9970 and 0.9862 for MCS and retransmission.

6 Discussion and Potential Applications

Privacy. Cellular telemetry is broadcasted without encryption, so it is accessible to anyone with easily-obtainable equipment, and is used here for the sole purpose of communications network design. Telemetry data itself (C-RNTI) does not identify persons, and the ability to combine the data with another dataset to identify an individual is so attenuated as to be generally impracticable. The relevant IRB has waived the associated master study.

Congestion control. The UE can instruct NR-Scope to send channel feedback to a sender, using NR-Scope as a service. NR-Scope's feedback is faster than half an RTT, as it can shortcut the full round trip path with only the upstream RAN to sender sub-path [30].

Security. The RRC messages and the resource allocation patterns that NR-Scope reveals can aid security assessments of the RAN, particularly to identify surveillance equipment and RAN vendors [23].

Internet Measurement. NR-Scope can be used to perform out-of-loop measurement of cellular networks for better understanding of RAN behavior (Fig. **10**) and design of RAN-aware applications.

RAN Aware Design for Closed RAN. NR-Scope provides fine-grained RAN monitoring even for a closed RAN, enabling unified spectrum-aware designs for both open and closed RANs.

7 Limitations and Future Work

On-device Processing. NR-Scope receives raw radio signals, thus can only be implemented on a phone with a customized firmware.

UCI Decoding. NR-Scope NR-Scope now decodes the DCIs in the downlink channel for uplink scheduling result. UCI in the uplink channel, containing the scheduling request and *Channel Quality Indicator* (CQI) (Fig. **1**), could be useful for uplink data scheduling analysis and is considered as a future work of NR-Scope.

Post-Processing Library. NR-Scope will leverage multiple USRPs to decode multiple cells, with the resulting data streams analyzed for carrier aggregation and UE handover events, fusing these into an aggregate data stream. We leave this as our future work.

8 Conclusion

This work has presented the first open-source 5G Standalone RAN passive network telemetry tool. NR-Scope does not require the cooperation of mobile network operators, or phone manufacturers, nor does it require modifications to the UE.

Acknowledgements

This work is supported by the National Science Foundation under Grant Nos. AST-2232457, CNS-2223556 and ITE-2326928.

References

[1] [n. d.]. AMARI Callbox Mini. ([n. d.]). https://www.amarisoft.com/test-and-measurement/device-testing/device-products/amari-callbox-mini

[2] [n. d.]. Mobile data traffic forecast – Mobility Report. ([n. d.]). https://www.ericsson.com/en/reports-and-papers/mobility-report/dataforecasts/mobile-traffic-forecast

[3] [n. d.]. Qualcomm® eXtensible Diagnostic Monitor. ([n. d.]). https://www.qualcomm.com/support/software-tools/tools.qualcomm-extensible-diagnostic-monitor.feb127fd-592d-4c60-acad-c7a0a54a99ea#overview

[4] [n. d.]. Wireless Network Monitoring Intelligence & Insights. ([n. d.]). https://thinkrf.com/

[5] 3GPP. 2008. *Requirements for further advancements for Evolved Universal Terrestrial Radio Access (E-UTRA) (LTE-Advanced).* Technical Report (TR) 36.913. 3rd Generation Partnership Project (3GPP). http://www.3gpp.org/DynaReport/36913.htm

[6] 3GPP. 2022. *Release 17 Description; Summary of Rel-17 Work Items.* Technical Report (TR) 21.917. 3rd Generation Partnership Project (3GPP). http://www.3gpp.org/DynaReport/21917.htm

[7] Venkat Arun and Hari Balakrishnan. 2018. Copa: Practical Delay-Based Congestion Control for the Internet. In *Proceedings of the Applied Networking Research Workshop*. ACM, Montreal QC Canada, 19–19. https://doi.org/10.1145/3232755.3232783

[8] Niklas Blum, Serge Lachapelle, and Harald Alvestrand. 2021. WebRTC: real-time communication for the open web platform. *Commun. ACM* 64, 8 (Aug. 2021), 50–54. https://doi.org/10.1145/3453182

[9] Nicola Bui and Joerg Widmer. 2016. OWL: a reliable online watcher for LTE control channel measurements. In *Proceedings of the 5th Workshop on All Things Cellular: Operations, Applications and Challenges (ATC '16)*. Association for Computing Machinery, New York, NY, USA, 25–30. https://doi.org/10.1145/2980055.2980057

[10] Neal Cardwell, Yuchung Cheng, C. Stephen Gunn, Soheil Hassas Yeganeh, and Van Jacobson. 2016. BBR: Congestion-Based Congestion Control: Measuring bottleneck bandwidth and round-trip propagation time. *Queue* 14, 5 (Oct. 2016), 20–53. https://doi.org/10.1145/3012426.3022184

[11] Yongzhou Chen, Ruihao Yao, Haitham Hassanieh, and Radhika Mittal. 2023. {Channel-Aware} 5G {RAN} Slicing with Customizable Schedulers. 1767–1782. https://www.usenix.org/conference/nsdi23/presentation/chen-yongzhou

[12] Robert Falkenberg and Christian Wietfeld. 2019. FALCON: An Accurate Real-Time Monitor for Client-Based Mobile Network Data Analytics. In *2019 IEEE Global Communications Conference (GLOBECOM)*. Institute of Electrical and Electronics Engineers, Piscataway, NJ, 1–7. https://doi.org/10.1109/GLOBECOM38437.2019.9014096

[13] FCC. 2023. 3.5 GHz band overview. (2023). https://www.fcc.gov/wireless/bureau-divisions/mobility-division/35-ghz-band/35-ghz-band-overview[fcc.gov]

[14] Sadjad Fouladi, John Emmons, Emre Orbay, Catherine Wu, Riad S Wahby, and Keith Winstein. 2018. Salsify: Low-Latency Network Video through Tighter Integration between a Video Codec and a Transport Protocol. In *Proc. of the USENIX Symposium on Networked Systems Design and Implementation (NSDI)*. USENIX Association.

[15] Prateesh Goyal, Mohammad Alizadeh, and Hari Balakrishnan. [n. d.]. ABC: A Simple Explicit Congestion Controller for Wireless Networks. ([n. d.]).

[16] Sangtae Ha, Injong Rhee, and Lisong Xu. 2008. CUBIC: a new TCP-friendly high-speed TCP variant. *ACM SIGOPS Operating Systems Review* 42, 5 (July 2008), 64–74. https://doi.org/10.1145/1400097.1400105

[17] Tuan Dinh Hoang, CheolJun Park, Mincheol Son, Taekkyung Oh, Sangwook Bae, Junho Ahn, BeomSeok Oh, and Yongdae Kim. 2023. LTESniffer: An Open-source LTE Downlink/Uplink Eavesdropper. In *Proceedings of the 16th ACM Conference on Security and Privacy in Wireless and Mobile Networks*. Association for Computing Machinery, New York, NY, USA, 43–48. https://doi.org/10.1145/3558482.3590196

[18] Keysight. 2023. Keysight WaveJudge. (2023). https://www.keysight.com/us/en/products/wireless-analyzers/wavejudge-wireless-analyzer-solutions.html

[19] Swarun Kumar, Ezzeldin Hamed, Dina Katabi, and Li Erran Li. 2014. LTE Radio Analytics Made Easy and Accessible. *SIGCOMM Comput. Commun. Rev.* 44, 4 (Aug. 2014), 211–222. https://doi.org/10.1145/2740070.2626320

[20] Yuanjie Li, Chunyi Peng, Zengwen Yuan, Jiayao Li, Haotian Deng, and Tao Wang. 2016. Mobileinsight: extracting and analyzing cellular network information on smartphones. In *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking (MobiCom)*. Association for Computing Machinery, New York, NY, USA, 202–215. https://doi.org/10.1145/2973750.2973751

[21] Feng Lu, Hao Du, Ankur Jain, Geoffrey M. Voelker, Alex C. Snoeren, and Andreas Terzis. 2015. CQIC: Revisiting Cross-Layer Congestion Control for Cellular Networks. In *ACM HotMobile*.

[22] Norbert Ludant, Pieter Robyns, and Guevara Noubir. 2023. From 5G Sniffing to Harvesting Leakages of Privacy-Preserving Messengers. In *2023 IEEE Symposium on Security and Privacy (SP)*. 3146–3161. https://doi.org/10.1109/SP46215.2023.10179353 ISSN: 2375-1207.

[23] Alexander Marder, Jon Larrea, Kc Claffy, Eric Kline, Kyle Jamieson, Bradley Huffaker, Lincoln Thurlow, and Matthew Luckie. 2024. REVEAL: Real-time Evaluation and Verification of External Adversarial Links. In *IEEE Military Communications Conference (MILCOM)*. IEEE.

[24] Mosolab. 2023. Mosolab Canopy Small Cell. (2023). https://opennetworking.org/products/moso-canopy-5g-indoor-small-cell/

[25] ONF. 2023. ONF Aether Onramp Project. (2023). https://opennetworking.org/aether/

[26] ShareTechNote. 2023. PDCCH Transmission Process, ShareTechnote. (2023). https://www.sharetechnote.com/html/5G/5G_PDCCH.html#PDCCH_Transport_Process

[27] Software Radio System. 2023. srsRAN Project: Open Source RAN. (2023). https://docs.srsran.com/projects/project/en/latest/tutorials/source/cotsUE/source/index.html

[28] Kazuki Takeda, Huilin Xu, Taehyoung Kim, Karol Schober, and Xingqin Lin. 2020. Understanding the Heart of the 5G Air Interface: An Overview of Physical Downlink Control Channel for 5G New Radio. *IEEE Communications Standards Magazine* 4, 3 (2020), 22–29. https://doi.org/10.1109/MCOMSTD.001.1900048

[29] Android Tcpdump. 2023. Official site for the tcpdump binary for Android devices. (2023). https://www.androidtcpdump.com/

[30] Haoran Wan and Kyle Jamieson. 2024. Evolving Mobile Cloud Gaming with 5G Standalone Network Telemetry. (Feb. 2024). https://doi.org/10.48550/arXiv.2402.04454 arXiv:2402.04454 [cs].

[31] Keith Winstein, Anirudh Sivaraman, and Hari Balakrishnan. 2013. Stochastic Forecasts Achieve High Throughput and Low Delay over Cellular Networks. In *Proc. of the USENIX Symposium on Networked Systems Design and Implementation (NSDI)*.

[32] Xiufeng Xie, Xinyu Zhang, Swarun Kumar, and Li Erran Li. 2015. piStream: Physical Layer Informed Adaptive Video Streaming over LTE. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking (MobiCom '15)*. Association for Computing Machinery, New York, NY, USA, 413–425. https://doi.org/10.1145/2789168.2790118

[33] Xiufeng Xie, Xinyu Zhang, and Shilin Zhu. 2017. Accelerating Mobile Web Loading Using Cellular Link Information. In *Proceedings of the 15th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys '17)*. Association for Computing Machinery, New York, NY, USA, 427–439. https://doi.org/10.1145/3081333.3081367

[34] Yaxiong Xie and Kyle Jamieson. 2022. NG-Scope: Fine-Grained Telemetry for NextG Cellular Networks. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 6, 1 (Feb. 2022), 1–26. https://doi.org/10.1145/3508032

[35] Yaxiong Xie, Fan Yi, and Kyle Jamieson. 2020. PBE-CC: Congestion Control via Endpoint-Centric, Physical-Layer Bandwidth Measurements. In *Proc. of the Conference of the ACM Special Interest Group on Data Communication (SIGCOMM)*. Association for Computing Machinery, New York, NY, USA, 451–464. https://doi.org/10.1145/3387514.3405880

A Transport Block Size Calculation

The following development restates the 3GPP standard TBS calculation [6] for convenience here. Firstly, we calculate the intermediate variable N_{RE}:

$$N_{RE}^{'} = N_{sc}^{RB} \times N_{symb}^{sh} - N_{DMRS}^{PRB} - N_{oh}^{PRB}, \quad (1)$$

$$N_{RE} = min(156, N_{RE}^{'}) * n_{PRB}. \quad (2)$$

N_{sc}^{RB} is the number of subcarriers per resource block and it's 12. N_{symb}^{sh} is the number of time domain OFDM symbols allocated through DCI, which can be found in t_alloc in the DCI grant in Appendix B. N_{DMRS}^{PRB} is the number of resource elements for DMRS per PRB and N_{oh}^{PRB} is the overhead of PDSCH, which are both defined in RRC messages. n_{PRB} is the frequency domain resource block allocated by DCI, which can be found in f_alloc in the DCI grant in Appendix B. N_{RE} represents the effective resource elements (1 OFDM symbol × 1 subcarrier) allocated for the UE in the DCI.

The second step is to calculate N_{info}: $N_{info} = N_{RE} \times R \times Q_m \times v$, where R is the code rate and Q_m is the modulation order, which are delivered through the DCI's MSC value and the UE checks the predefined tables with it [6]. And v is number of layers, transferred in the MSG 4's element "$pdsch-ServingCellConfig$": "$maxMIMO-Layers$".

Final step is to calculate TBS, with some value calculation:

- If $N_{info} \leq 3824$, we have $N_{info}^{'} = 2^n \times round(\frac{N_{info}-24}{2^n})$, where $n = \lfloor log_2(N_{info} - 24) \rfloor - 5$. Then, if $R \leq 1/4$, we have $TBS = 8C\lceil \frac{N_{info}^{'}+24C}{8C} \rceil - 24$, where $C = \lceil \frac{N_{info}^{'}+24}{3814} \rceil$. If $R > 1/4$, we have $TBS = 8\lceil \frac{N_{info}^{'}+24}{8} \rceil - 24$ if $N_{info}^{'} < 8424$, and $TBS = 8C\lceil \frac{N_{info}^{'}+24C}{8C} \rceil - 24$ if $N_{info}^{'} \geq 8424$, where $C = \lceil \frac{N_{info}^{'}+24}{8424} \rceil$.
- If $N_{info} > 3824$, we have $N_{info}^{'} = max(24, 2^n \lfloor \frac{N_{info}}{2^n} \rfloor)$, where $n = max(3, \lfloor log_2(N_{info}) \rfloor - 6)$. Then we check table 5.1.3.2-2 in [6] to determine the TBS.

B DCI and Grant

A downlink DCI looks like:

```
DCI:
    c-rnti=0x4296, dci=1_1, ss=ue, L=0, cce=7, f_alloc=0x33,
    t_alloc=0x0, mcs=27, ndi=0, rv=0, harq_id=11, dai=2,
    tpc=1, harq_feedback=2, ports=7, srs_request=0, dmrs_id=0
```

And its translated grant for PDSCH is like:

```
Grant:
    rnti=0x4296, rnti_type=C-RNTI, k=0, mapping=A, t_alloc=2:12,
    f_alloc=0:2, nof_dmrs_cdm_grps=2, beta_dmrs=1.412538,
    nof_layers=2, n_scid=0,
CW0:
    mod=256QAM, nof_layers=2, mcs=27, tbs=3240, R=0.926,
    rv=0, ndi=0, nof_re=432, nof_bits=3456
SCH:
    mcs_table=256qam, xoverhead=0
```

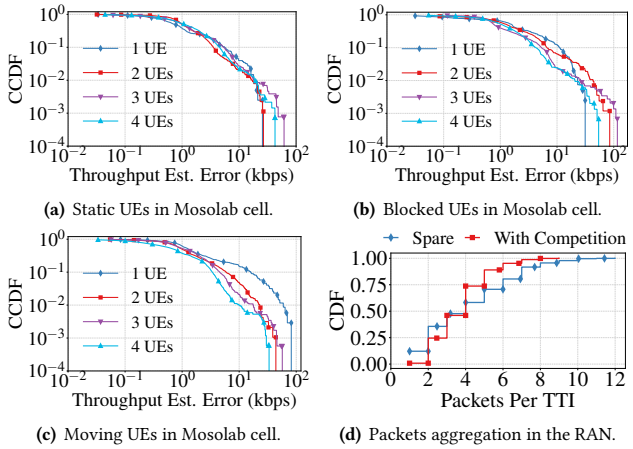

(a) Static UEs in Mosolab cell. **(b)** Blocked UEs in Mosolab cell.

(c) Moving UEs in Mosolab cell. **(d)** Packets aggregation in the RAN.

Fig. 16— *Accuracy of NR-Scope* throughput estimation under different UE status (a-c). Packet aggregation pattern (d).

C Detailed Figure for Fig. 16

Here we show the more detailed result for UEs in different scenarios in Fig. **16**.

D Packet Aggregation

NR-Scope can be used to analyze the packet aggregation behavior of RAN. We compare the TBS in each TTI and the receiving packet size to get packets per TTI measurement, as shown in Fig. **16(d)**. Packets aggregated in one TTI have nearly the same arrival time in physical layer, schemes using inter packet arrived time as bandwidth indicator may fail but can benefit from NR-Scope's precise RAN capacity estimation.

E Open Source Tool Release

There is an open source release of the NR-Scope tool on Github. The software release is accompanied by a README file with detailed installation and running instructions, as well as an end-user agreement requiring attestation of responsible use of the tool.