5G Tracker – A Crowdsourced Platform to Enable Research Using Commercial 5G Services

Arvind Narayanan, Eman Ramadan, Jacob Quant, Peiqi Ji, Feng Qian, Zhi-Li Zhang
Department of Computer Science & Engineering, University of Minnesota, U.S.A.

ABSTRACT

While 5G has offered many opportunities for research, the majority of studies have been conducted with constrained experimental settings or done privately by 5G operators. Even a year after the launch of commercial 5G networks, research over commercial 5G has been limited due to the lack of publicly available tools and datasets. In this paper, we propose 5G Tracker – a crowdsourced platform intended to aid researchers in collecting and leveraging large-scale 5G datasets. This platform includes an Android app that records passive and active measurements tailored to 5G networks and research. We have been using 5G Tracker for over 8 months, during which time we have collected over 4 million data points, consuming over 50 TB of cellular data across multiple 5G carriers in the U.S. Our experience shows that 5G performance is affected by several user-side contextual factors besides location such as user mobility level, orientation, weather, location dynamics (*e.g.,* moving vehicles), and environmental features such as pillars, foliage, and buildings. This is partly because mmWave signals (considered key to mainstream 5G) are known to be highly sensitive to obstructions and user mobility. These observations highlight the need to move towards building *context-aware* 5G performance prediction models that can provide guidance for decisions at various layers such as preemptive handoff, multi-path scheduling, tower placement, and *"5G-aware"* application development. Finally, we showcase the utility of our platform by building a first of kind, interactive 5G coverage mapping application as a case study driven by the data we collected, which is publicly available at: `https://5gophers.umn.edu`.

CCS CONCEPTS

• Information systems → Crowdsourcing; **• General and reference → Measurement**; **• Networks → Wireless access networks**; **Mobile networks**;

KEYWORDS

5G, 5GTracker, 5G Tracker, 5Gophers, measurements, crowd-sourcing

ACM Reference Format:
Arvind Narayanan, Eman Ramadan, Jacob Quant, Peiqi Ji, Feng Qian, Zhi-Li Zhang. 2020. 5G Tracker – A Crowdsourced Platform to Enable Research Using Commercial 5G Services. In *ACM Special Interest Group on Data Communication (SIGCOMM '20 Demos and Posters), August 10–14, 2020, Virtual Event, USA.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3405837.3411394

SIGCOMM '20 Demos and Posters, August 10–14, 2020, Virtual Event, USA

ACM ISBN 978-1-4503-8048-5/20/08…$15.00
https://doi.org/10.1145/3405837.3411394

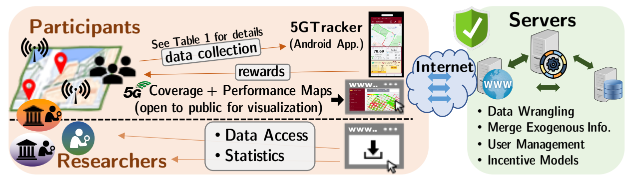

Figure 1: Overview of 5G Tracker platform.

1 INTRODUCTION & MOTIVATION

With its touted high speed and low latency, 5th generation (5G) wireless technology is envisaged to usher in a new smart and connected world. Within a year of its commercial debut in 2019, more than 125 carriers have deployed 5G services worldwide [3]. But even in the midst of such fast growth and excitement, multiple factors make research of commercial 5G networks challenging. Carriers' choice of which radio frequencies to use (*e.g.,* whether they support low-band, mid-band, and/or high-band/mmWave 5G [2, 6]) drastically affect their networks' coverage and performance. The absence of openly available tools to collect 5G data and the sparsity of most carriers' commercial 5G deployments pose obstacles to researchers relying on data collected from commercial 5G services. 5G operators collect several key performance indications (KPIs) from their customers that are vital for network planning, resource allocation, deployment and operations. However, such data are typically not released to the public due to concerns regarding competitive advantage and customer privacy. Some researchers might resort to building their own measurement platforms, but this can be a prohibitively costly undertaking for many research problems due to the effort required to develop and validate the system and the need to obtain 5G devices with subscriptions.

In the past, Mobile Crowd Sensing (MCS) [5, 7] – a popular crowd sourcing paradigm – has shown the potential to enable large-scale data collection by incentivizing users to install a mobile app in their smartphone and to collect and upload data of interest. However, at this time, none exist that are tailored for 5G networks and research.

2 SYSTEM DESIGN OF 5G TRACKER

In this paper, we present 5G Tracker – an MCS-based system that we have developed to enable researchers to use 5G data collected from commercially operational 5G services. Fig. 1 shows the overview of 5G Tracker platform. **Android app.** Participants worldwide can install 5G Tracker – an Android mobile app. The app records passive measurements (*e.g.,* user's location, network state information, active radio type, *etc.*) by accessing relevant Android APIs (or by parsing raw string representations of objects when APIs are unavailable) sampled periodically at a configurable interval (1 sec. by default). Passive measurement collection does not require any user interaction, running completely in the background.

This research was in part supported by NSF under Grants CNS-1917424, CNS-1903880, CNS-1915122, CNS-1618339, CNS-1617729, CNS-1814322, CNS-1831140, CNS-1836772, and CNS-1901103.

Table 1: Partial list of fields captured by 5G Tracker.

Type	Fields Captured
Static Info.	Phone Make & Model, Carrier Name, OS, iPerf3.7 Config, Ping Config
N/W State	active network type (WiFi / cellular), IP address of network interfaces
Phone State	Timestamp, CPU & RAM utilization, battery & CPU temperature
Cellular State	Active Radio Type (5G-NR/4G/3G), Cell ID, Physical Cell ID (PCI), NR Bands, ARFCN, Tracking Area Code (TAC), Mobile Network Code (MNC), Total RX/TX Bytes, Signal Strength Info. (RSSI, RSRP, SINR, RSRQ, SSRSRP, SSRSRQ, SSSINR, CSIRSRP, CSIRSRQ, CSISINR)
WiFi State	SSID, Signal Strength, Link Bandwidth, Channel Information
User State	Geo-coordinates, moving speed, compass direction, user activity
Ping Logs	Round Trip Time (RTT)
iPerf 3.7 Logs	Uplink/Downlink Speed Per Second, Total RX/TX Bytes, TCP Retransmits, TCP Send Window, RTT

Upon the participants' direction, the data collected gets uploaded to our server using a secure connection. In return for their efforts, we incentivize participants with gift cards. Our platform also facilitates active measurement collection by integrating two popular tools directly into the app: `ping` and `iPerf3.7`. These tools enable interested users (*e.g.* researchers) to perform controlled measurements repeatedly. Table 1 lists a subset of fields recorded by the app. Additionally, we also allow users to report (and confirm) important information such as presence/installation of 5G base stations at a location and other ground truth contextual information (*e.g.* whether the user is driving/walking, weather conditions, etc.). For ethical reasons, our platform never collects any personally identifiable information (PII)[1] of the participants. Besides granting OS-level permissions, our mobile app does not require *root* privileges to run.

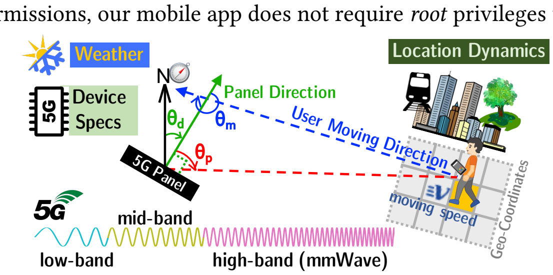

Figure 2: Contextual Information affecting 5G Performance.

Server-End. 5G Tracker's backend server cluster serves multiple purposes: receiving, storing, processing, and presenting data submitted by users. It provides an HTTPS-based API for registering mobile devices, submitting collected data, updating incentive models[2], as well as for user account management. An important step performed at the server is to incorporate relevant exogenous information such as spatial relationship between the user and tower (distance, orientation angles, *etc.*). Fig. 2 illustrates a range of user-side contextual features (*e.g.*, weather, location dynamics, *etc.*) that we have experienced to affect 5G performance. Without doubt there are likely several carrier-side factors that affect 5G performance. However, in the past, similar MCS platforms [4, 7] have helped the research community without requiring collaboration with carriers. Server-side factors such as location of server, protocols used, *etc.* can also affect user-perceived 5G performance. However such factors are generally controllable by researchers. To illustrate how 5G Tracker can be used, next we present a case study to generate *data-driven* 5G coverage and performance maps.

[1] Email addresses are used exclusively for authentication and incentive delivery. No form of user identifier is included in the publicly available dataset.
[2] Our platform allows researchers (under collaboration status) to specify their own incentive model driven by their research needs.

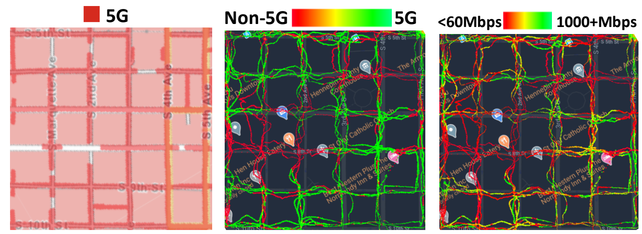

(a) By carrier [1]. **(b) % of 5G samples.** **(c) Throughput map.**

Figure 3: 5G Coverage and Throughput Maps of Downtown region in a large U.S. city.

3 A CASE STUDY: 5G COVERAGE MAPS

Even as we witness rapid 5G deployments worldwide, carriers, especially those that provide mmWave-based 5G and claim ultra-fast speeds, have largely struggled to provide meaningful coverage maps. We know that mmWave radio signals are highly directional, require line-of-sight (LoS) for maintaining 5G connectivity, and have limited range [6]. Particularly, they are sensitive to obstructions and user mobility thus requiring carriers to densely deploy 5G base stations and adopt complex beam tracking algorithms to ensure seamless 5G connectivity [2]. Fig. 3a shows a coverage map of a mmWave-based 5G network in the downtown area of a city in the U.S. under study. According to the carrier, it appears that the entire region has 5G coverage. Using our platform with `iPerf3`-mode enabled, we walk and drive through several areas in this region following random trajectories, visiting each location multiple times to improve confidence. We then split the map into grids where each grid represents an area of 2m×2m. For each grid we compute two measures: (1) how many samples out of the total samples were connected to the 5G network and (2) the mean throughput perceived in each grid. We see that this carrier's coverage map glosses over many details: (1) in many geographic pockets, we did not get 5G connectivity (see Fig. 3b for our *measured* 5G *coverage* area), (2) even for areas where we were connected to 5G, throughput performance varied widely from 0 Mbps to >1 Gbps (see Fig. 3c for our *measured* 5G *throughput* map). Although not shown, the throughput map further varies when filtered by criteria such as the moving direction/orientation and speed of the user (see Fig. 2), highlighting the sensitivity of mmWave 5G's performance to several factors. We have been using 5G Tracker for over 8 months and have collected more than 4 million data points. Using this dataset, we have released these first of a kind 5G coverage and performance maps out for public consumption at the URL: `https://5gophers.umn.edu`.

4 CONCLUSION

We present 5G Tracker – a crowdsourced platform that helps collect passive and active measurements from operational 5G carriers. It is sufficiently non-invasive for users to adopt yet flexible enough to enable researchers to conduct rich, controlled, repeatable experiments, thus allowing them to focus on research rather than developing tools. We also present a case study that visualizes 5G coverage and performance using the dataset. We believe combining measurements with user-side contextual information (beyond location) will aid in building robust 5G performance prediction models. At the same time, the data collected over time will also be vital in understanding the evolution of commercial 5G performance.

REFERENCES

[1] 2020. 5G Coverage Map: This is 5G Built Right | Verizon. (2020). Retrieved June 2020 from https://www.verizon.com/5g/coverage-map/

[2] 2020. Keysight Technologies | Engineering the 5G World Just Got Easier. (2020). Retrieved June 2020 from https://www.keysight.com/us/en/assets/7119-1223/ebooks/Engineering-the-5G-World.pdf

[3] 2020. Ookla 5G Map - Tracking 5G Network Rollouts Around the World. (2020). https://www.speedtest.net/ookla-5g-map

[4] 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 âĂŹ16)*. Association for Computing Machinery, New York, NY, USA, 202âĂŞ215. https://doi.org/10.1145/2973750.2973751

[5] Yi-Hao Lin, Jyh-Cheng Chen, Chih-Yu Lin, Bo-Yue Su, and Pei-Yu Lee. 2018. SensingGO: Toward Mobile/Cellular Data Measurement with Social and Rewarding Activities. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking (MobiCom âĂŹ18)*. Association for Computing Machinery, New York, NY, USA, 765âĂŞ767. https://doi.org/10.1145/3241539.3267733

[6] Arvind Narayanan, Eman Ramadan, Jason Carpenter, Qingxu Liu, Yu Liu, Feng Qian, and Zhi-Li Zhang. 2020. A First Look at Commercial 5G Performance on Smartphones. In *Proceedings of The Web Conference 2020 (WWW âĂŹ20)*. Association for Computing Machinery, New York, NY, USA, 894âĂŞ905. https://doi.org/10.1145/3366423.3380169

[7] Ashkan Nikravesh, Hongyi Yao, Shichang Xu, David Choffnes, and Z. Morley Mao. 2015. Mobilyzer: An Open Platform for Controllable Mobile Network Measurements. In *Proceedings of the 13th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys âĂŹ15)*. Association for Computing Machinery, New York, NY, USA, 389âĂŞ404. https://doi.org/10.1145/2742647.2742670