NextG-UP: A Longitudinal and Cross-Sectional Study of Uplink Performance of 5G Networks

Moinak Ghoshal[1*], Imran Khan[1*], Qiang Xu[2], Z. Jonny Kong[2], Y. Charlie Hu[2], Dimitrios Koutsonikolas[1]

[1]Institute for the Wireless Internet of Things, Northeastern University, USA, [2]Purdue University, USA

{ghoshal.m, khan.i, d.koutsonikolas}@northeastern.edu, {xu1201, kong102, ychu}@purdue.edu

ABSTRACT

5G networks are being deployed rapidly across the world and have opened door to many uplink-oriented, bandwidth-intensive applications such as Augmented Reality and Connected Autonomous Vehicles. However, the roll-out is still in the early phase and the nature of deployment also varies across different geographic regions of the world. In this demo, we present NextG-UP, an Android-based tool designed to help understand the performance and evolution of 5G networks around the world. The crowd-sourcing mobile app collects various cellular network metrics and runs a short uplink throughput/latency test.

CCS CONCEPTS

• **Networks** → **Mobile networks**; **Network measurement**; **Network performance analysis**.

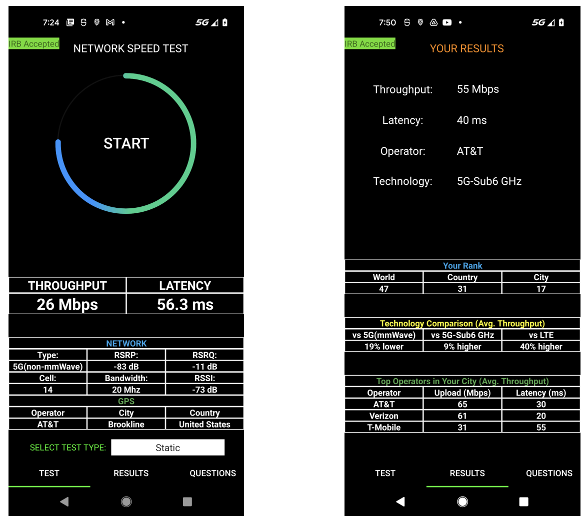

Figure 1: NextG-UP App layout.

ACM Reference Format:
Moinak Ghoshal, Imran Khan, Z. Jonny Kong, Qiang Xu, Dimitrios Koutsonikolas, Y. Charlie Hu. 2022. NextG-UP: A Longitudinal and Cross-Sectional Study of Uplink Performance of 5G Networks. In *The 28th Annual International Conference on Mobile Computing and Networking (ACM MobiCom '22), October 17–21, 2022, Sydney, NSW, Australia.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3495243.3558746

1 INTRODUCTION

The evolution of mobile networks has come a long way. The most recent generation of cellular networks, 5G, and in particular, 5G mmWave, promises unprecedented high bandwidth and ultra lower latency and holds the promise to finally support latency-critical applications (LCAs) such as Augmented Reality, Mixed Reality, and Connected Autonomous Vehicles (CAVs). Indeed, multiple mobile operators and cloud providers have formed alliances, e.g., AT&T and Microsoft, and Verizon and AWS, to showcase 5G solutions for LCAs such as AR.

However, 5G deployments are still at their infancy. A few recent measurement studies [2–4] reported that, although today's mmWave deployments indeed offer Gbps throughput and lower latency than 4G LTE, their performance is often suboptimal, coverage is sporadic, and applications cannot always take advantage of the full potential of 5G mmWave. In addition, these studies have mostly focused on measuring the 5G *downlink* performance while the uplink performance of 5G networks remains largely unknown.

In this work, we aim to fill this gap by answering the question: *How good is 5G network uplink performance, and is it sufficient to enable LCAs such as AR or CAVs?* More specifically, since most LCAs distinguish themselves from legacy apps for their heavy, bursty *uplink* data transfer, can current 5G deployments meet the uplink traffic demand of LCAs, as 5G has provisioned much higher downlink bandwidth than uplink bandwidth similarly as all its predecessors?

Answering this question systemically faces two fundamental challenges. (1) *5G deployment is rapidly evolving.* As 5G deployments are still at their infancy, operators often make

*Moinak Ghoshal and Imran Khan are co-primary authors.

ACM MobiCom '22, October 17–21, 2022, Sydney, NSW, Australia

ACM ISBN 978-1-4503-9181-8/22/10.
https://doi.org/10.1145/3495243.3558746

Table 1: List of metrics collected.

Metric	Description
GPS	User's City, Country
Network Type	5G (mmWave) /5G (sub-6 GHz)/ LTE
RSRP	Reference Signal Received Power
RSRQ	Reference Signal Received Quality
RSSI	Received Signal Strength Indicator
Cell-ID, Cell-BW	Connected cell id and frequency
Operator	User's cellular operator

infrastructural changes, which can change the network performance at a given location. Previous works performed measurements within a short time-span, ranging from a few days up to a couple of months. However, any findings from such studies might be short-lived and lead to wrong conclusions about the potential of 5G in the long-term. (2) *5G deployment exhibits geographic diversity.* Recent studies [1, 3, 4] report that 5G performance exhibits large geographical diversity. For example, in our recent work [1], we found that the 5G performance can vary substantially across operators and across cities for the same operator, as operators employ different resource allocation policies. Consequently, this question cannot be answered without a detailed, *longitudinal and cross-sectional* study of 5G uplink performance.

To address these challenges, we designed a crowd-sourced Android app, NextG-UP, to collect cellular uplink bandwidth and latency data from users around the world over a period of 12 months. We plan to make the collected dataset publicly available. In this demo, we will showcase the NextG-UP app and present the app's workflow, the type of data it collects, and the results that will be displayed to the end-user.

The ultimate goals of developing such a 5G uplink performance monitoring tool are four-fold: (1) It will expose user-experienced 5G performance to 5G users around the world; (2) It will track and expose the evolution of the "digital-divide" as the society enters the 5G-era; (3) It will help mobile network operators to improve the deployment process and online operations of 5G networks; and (4) It will help mobile app developers to optimize and co-evolve the functional performance of apps with 5G performance.

2 DESIGN OVERVIEW

Functionality. The NextG-UP app has two main functionalities. It collects various cellular network metrics and measures uplink TCP throughput and RTT. To collect the network metrics, the application requires the users to grant permission to access certain data on the phone (TELEPHONY, GPS, etc.). A detailed list of the collected metrics is shown in Table 1.

Cloud servers. To measure uplink throughput and RTT, we have setup testing servers in two different zones: US and Europe. In US, we further divide the zones in two sub-zones: East Coast and West Coast. In each coast, we are running two different types of servers. One is an AWS Wavelength server specially designed for edge computing that provides ultra-low latency service for Verizon users located in the US, and the other one is an AWS Cloud server. The AWS cloud servers are deployed in North Virginia and Oregon whereas the AWS Wavelength servers are deployed in Boston and San Francisco. For the Europe zone, we deployed an AWS cloud server in Frankfurt.

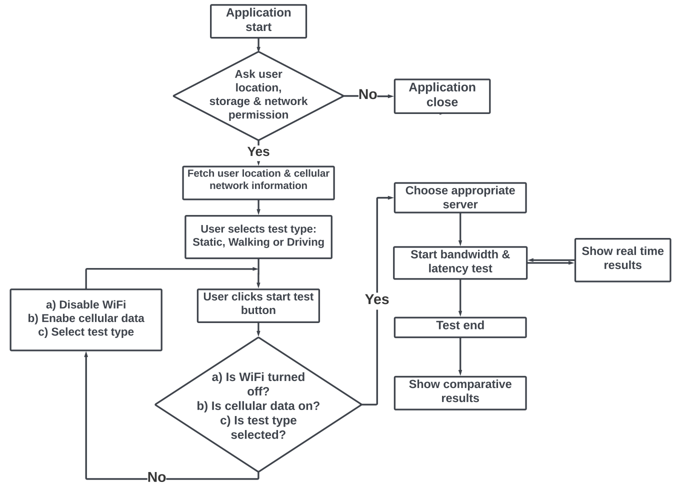

Figure 2: Workflow of the NextG-UP application.

App design. The workflow of the NextG-UP application is shown in Fig. 2. The app collects the user location in the background and chooses the nearest server. The user is then prompted to select between three test types (static, walking, or driving). Since the app aims at measuring cellular network performance only, there is a check in the background to notify users if they have their WiFi turned on. Once these checks are completed, the app measures the uplink TCP throughput using nuttcp-8.1.4. The throughput test runs for 10 s. Then the RTT test begins using the ping utility and sends 11 ICMP packets 200 ms apart from each other. During the tests, the throughput and RTT values are displayed on the screen in real-time. After the tests are completed, a summary of the results is sent back to a server, which provides feedback (user's rank in the city, top operators in the user's city, etc.) to the user about their test performance (Fig. 1). The app is light-weight. The app image is 6.5 MB in size and uses less than 250 MB of memory while running.

IRB approval. This study is IRB approved.

3 DEMONSTRATION

We have the AWS testing servers up and running and we will use a smartphone to demonstrate the app. Conference attendees will also be asked to download the app from Google play-store on their Android smartphones and run it. They will be able to see the results on their smartphones and submit any feedback or questions. We do not require any additional infrastructure from the conference.

REFERENCES

[1] Phuc Dinh, Moinak Ghoshal, Dimitrios Koutsonikolas, and Joerg Widmer. 2022. Demystifying Resource Allocation Policies inOperational 5G mmWave Networks. In *In Proc. of IEEE WoWMoM*.

[2] 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 *Proc. of ACM WWW*.

[3] Arvind Narayanan, Xumiao Zhang, Ruiyang Zhu, Ahmad Hassan, Shuowei Jin, Xiao Zhu, Denis Rybkin, Dustin Zhang, Michael Yang, Z. Morley Mao, Feng Qian, and Zhi-Li Zhang. 2021. A Variegated Look at 5G in the Wild: Performance, Power, and QoE Implications. In *Proc. of ACM SIGCOMM*.

[4] Muhammad Iqbal Rochman, Vanlin Sathya, Norlen Nunez, Damian Fernandez, Monisha Ghosh, Ahmed S. Ibrahim, and William Payne. 2022. A Comparison Study of Cellular Deployments in Chicago and Miami Using Apps on Smartphones. In *Proc. of ACM WiNTECH*.