Understanding Commercial 5G and its Implications to (Multipath) TCP

Lan Ding[a], Ye Tian[a,*], Tong Liu[a], Zhongxiang Wei[a], Xinming Zhang[a]

[a] *Anhui Key Laboratory on High-Performance Computing*
School of Computer Science and Technology, University of Science and Technology of China, Hefei, 230026, Anhui, China

Abstract

In this paper, we perform an empirical study on commercial 5G. Unlike previous works studying 5G in early deployments, we focus on fully deployed 5G networks that already serve massive numbers of users, and in particular, we perform controlled experiments to analyze TCP and multipath TCP (MPTCP) performances and behaviors over 5G. Our results show that 5G faces a dilemma in maintaining in-network buffer for TCP traffic: on one hand, TCP packets are less likely to accumulate within the network, thanks to the high data rate of 5G radio; but on the other hand, with TCP performance highly sensitive to signal quality, when signal becomes poor abruptly in mobile environment, the in-network buffer easily overflows, leading to packet loss bursts. We also confirm that employing MPTCP can aggregate capacities on 5G and 4G paths, but under the current deployment model where 5G and 4G base stations co-locate and share the same core network, MPTCP does not necessarily enhance communication reliability. In addition to active measurement experiments, we leverage a speedtest service to evaluate commercial 5G from a nationwide perspective. By exploiting over $13k$ crowdsourced speedtest results reported from 197 cities in China, we find that there exist considerable regional differences on 5G performances across the nation, and with the 4G co-locating deployment model, the 5G networks in some populous metropolitan areas have inferior performances. Our study provides insights for people to understand commercial 5G and to optimize upper-layer protocols and applications over 5G.

Keywords: 5G; TCP; Multipath TCP (MPTCP)

1. Introduction

The fifth generation mobile cellular network, namely 5G, attracts huge attentions from academics, industries, and even politics in recent years. 5G integrates a series of innovations such as new radio frequencies, new core network management, massive MIMO, channel coding, small cells, softwarization, SDN/NFV, multitenant, slicing, etc., and comparing with 4G[1], 5G can provide greater mobile bandwidth, ultra low latency, and support massive machine communications [1]. With these features, 5G is expected to drive a wide range of novel applications in entertainment, manufacturing, transportation, and public services, and is estimated to bring new businesses worth 13.2 trillion U.S. dollars to the global economy [2].

After years of research, development, and standardization, 5G has been commercial in a few countries

*Corresponding author
Email address: `yetian@ustc.edu.cn` (Ye Tian)
[1] By 4G we mean LTE.

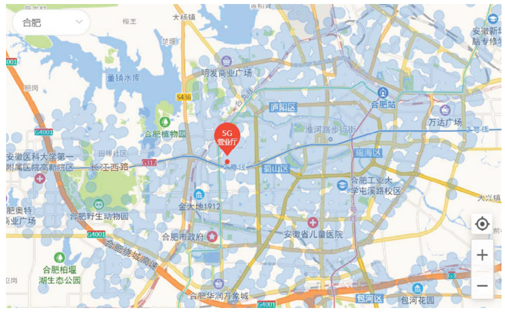

Figure 1: Mobile's 5G signal coverage in urban areas of our city on July 5, 2020, where each light blue circle represents a 5G cell. Mobile provides commercial 5G service at the 2515−2675 MHz frequency band.

around the world. Till October 2020, the three major mobile network operators, namely Mobile, Unicom, and Telecom, have together deployed over $690k$ 5G base stations that connect over $160M$ user equipments (UEs), in all 297 prefectural cites in China[2] [3]. For example, Fig. 1 presents Mobile's 5G signal coverage in our city, where each light blue circle represents a 5G cell. We can see that 5G is pervasively available in many urban areas of the city.

Understanding 5G is necessary and important to existing and emerging applications. In particular, as 5G will be a major access to the Internet in future, how it influences upper-layer protocols, such as TCP, multipath TCP (MPTCP) [4], and QUIC [5], is an important issue.

In this paper, we perform an empirical study on commercial 5G. Unlike previous works (e.g., [6] and [7]) that study 5G networks in their early deployments, which have few users and base stations, we focus on fully deployed commercial 5G networks that are deployed in large scale and already serve massive number of users. We carry out controlled measurement experiments with off-the-shelf smartphones as the UE. Since 5G is very new, there is no mature method for capturing the 5G cellular information, we develop our own method and tool on smartphones. With the methodologies, we study TCP and MPTCP performances and behaviors over 5G, under stationary as well as mobile environments. We have made the following observations.

- With stationary UE, 5G provides approximately 8 times of TCP upload and download throughputs over 4G. Thanks to the high data rate of 5G radio, TCP packets are less likely to accumulate in in-network buffer as in 3G/4G, and by eliminating bufferbloat, a TCP flow over 5G experiences much shorter RTT than in 4G. Employing MPTCP can aggregate 5G and 4G capacities, and 5G path is more preferred by MPTCP's congestion control.

- In highly mobile environment with UE moving at a speed between 76−113 km/h, a TCP flow's throughput is sensitive to the UE perceived signal. When signal strength becomes weak abruptly, which is common in mobile environment, in-network buffer quickly gets filled and overflows, leading to bursts of packet losses. For most of time, an MPTCP's 5G and 4G subflows are carried by 5G and 4G base stations co-located in the same cell; the two subflows experience throughput reductions simultaneously when the shared cell

[2]By China we mean mainland China in this paper.

Table 1. Frequencies allocated to major 5G operators in China.

Operator	Mobile	Unicom	Telecom
Frequency (MHz)	$2515 - 2675$	$3500 - 3600$	$3400 - 3500$
	$4800 - 4900$		

is faulty.

Our findings reveal that 5G faces a dilemma in maintaining in-network buffer for TCP flows: on one hand, maintaining a large in-network buffer seems no longer necessary as fewer packets are accumulated on network nodes, thanks to 5G radio's high data rate; but on the other hand, when signal strength and data rate change dramatically in mobile environment, the in-network buffer is not large enough to accommodate such variation and easily overflows, leading to bursts of packet losses.

Another implication is that although MPTCP can effectively aggregate 5G and 4G capacities, however, under the current deployment model where 5G and 4G base stations are typically co-located and share the same cell, MPTCP does not necessarily enhance wireless communication reliability, as both subflows could be influenced by the same problem, such as temporal congestion or malfunction within the shared cell.

Besides actively conducting measurement experiments, we also leverage a speedtest service to evaluate commercial 5G networks from a nationwide perspective. By exploiting over $13k$ crowdsourced speedtest results reported from 197 cities, which constitute 67% of the cities in China, we reveal that the 5G networks in different cities and provinces exhibit varying performances, but even at the 10^{th} percentile, 5G still provides a significant improvement over 4G; however, due to the 4G co-locating deployment model, the 5G networks in some populous metropolitan areas could be overloaded, and have inferior performances.

As far as we know, this is the first 5G measurement study with high-speed mobile UE, and employing MPTCP on both 5G and 4G paths. We are also the first to evaluate commercial 5G from a nationwide perspective. The remainder part of this paper is organized as follows. We introduce background and related works in Sec. 2; Sec. 3 presents our measurement methodologies; we study performances and behaviors of TCP and MPTCP over 5G with stationary and mobile UEs in Sec. 4 and Sec. 5 respectively; in Sec. 6, a nationwide evaluation on commercial 5G is presented by leveraging the crowdsourced speedtest results; finally, we conclude this paper in Sec. 7.

2. Background and Related Work

2.1. Background

2.1.1. 5G NR

5G's radio interface defined by 3GPP is known as New Radio (NR), and its spectrum is divided into two frequency ranges: FR1 ($\leq$6 GHz) and FR2 ($\geq$24 GHz) [8]. In particular, FR2 is referred to as the millimeter wave (mmWave) frequency. Although mmWave 5G can provide a throughput up to 20 Gbps [9], however, most of the commercial 5G networks around the world currently use FR1 for technological and economical reasons

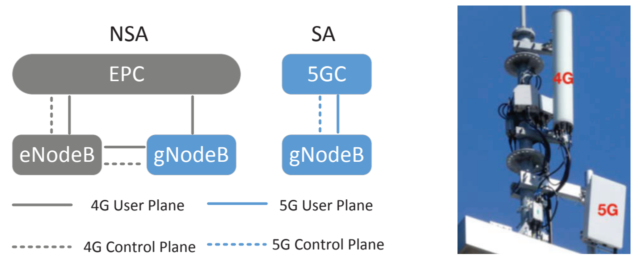

Figure 2: Left: Demonstration of NSA and SA models. Since there are many options in both models, the figure indeed presents option 3x for NSA and option 2 for SA that are adopted by Mobile as examples. Right: a 5G base station that is co-located with a 4G base station on the same tower under the NSA model.

[10]. For example, all the three major Chinese operators' commercial 5G networks currently work with the FR1 frequencies, as listed in Table 1.

2.1.2. Deployment model

3GPP defines two models for operators to transition from 4G to 5G: namely SA (Standalone Access) and NSA (Non-Standalone Access) [11]. As shown in the left figure of Fig. 2, under the NSA model, a 5G base station (gNodeB) connects to a 4G base station (eNodeB) as the master node (MN), and they connect to the same 4G core network (EPC) via optical fiber [12][13]. While under the SA model, 5G base station connects to the dedicated 5G core network (5GC). Note that in practice, an NSA 5G base station is typically co-located on the same tower with a 4G base station as shown in the right figure of Fig. 2, and they share the same cell ID.

2.2. Related Work

There is a rich literature of empirical studies on understanding cellular networks and their interactions with upper-layer protocols and applications. Tso *et al.* [14] presented a measurement study of a 3.5G (i.e., HSPA) cellular network in a wide-range of mobile environments. The authors revealed that mobility has largely negative impacts on the performance, but mobility can also improve fairness of bandwidth sharing among users. Jiang *et al.* [15] tackled the bufferbloat problem in various 3G/4G networks, and proposed a dynamic receive window adjustment algorithm to improve TCP performance over bufferbloated cellular networks. Huang *et al.* [16] studied the TCP behaviors and performances over 4G with traffics captured within a carrier's LTE network as well as controlled experiments. They observed the prevalent bufferbloat phenomena, and found that many large TCP flows under-utilize the LTE bandwidth. Li *et al.* [17] analyzed TCP behaviors and performances over 3G/4G networks with UEs moving on high-speed trains. They found that TCP performance degrades by frequent handoffs, and the spurious retransmission timeout (RTO) rate is high, due to wide RTT variations.

Recently, Multipath TCP (MPTCP) has been applied to enable smartphones to use cellular and WiFi networks simultaneously. Deng *et al.* [18] analyzed the benefits for employing MPTCP for various applications in mobile environment. Nikravesh *et al.* [19] analyzed the energy efficiency of MPTCP on mobile devices, and proposed a new mobile multipath software architecture. Coninck *et al.* [20] studied the MPTCP traffic of the iOS Siri application, and proposed a solution for enabling fast handover with low cellular usage for MPTCP under mobile environment. Abdelsalam *et al.* [21] analyzed the benefit of employing MPTCP for aggregating satellite and xDSL links, and Bujari *et al.* [22] proposed a virtualized performance enhancing proxy (vPEP) solution that makes use of MPTCP to dynamically enable satellite link as a supplementary for terrestrial link, with the purpose of optimizing web traffic performance. Besides MPTCP, Multipath-QUIC (MP-QUIC) [23] is proposed for 5G mission critical applications [24].

Since commercial 5G was just available in 2019, most studies resort to simulations for analyzing the interactions between 5G and upper-layer protocols. For example, Zhang *et al.* [25] studied TCP under an ns-3 mmWave simulation framework, and found that its performance depends on a number of TCP and cellular network parameters. To study operational 5G networks in their early deployments, Narayanan *et al.* [6] focused on the mmWave 5G networks operated by the major U.S. carriers, and analyzed the factors that impact the performance. Xu *et al.* [7] presented a comprehensive measurement study on a 5G network deployed in a limited region, and study the interactions between TCP and 5G.

Our work differs from the previous works in three aspects: First, unlike previous works that focus on 5G networks in early deployments, which have few users and limited base station deployments, we focus on fully deployed commercial 5G networks that already serve massive number of users; Second, we are the first to study 5G with high-speed mobile UE, and to employ MPTCP on both 5G and 4G paths. Finally, we provide the first evaluation on commercial 5G from a nationwide perspective.

3. Measurement Methodology

3.1. Active Measurement Setup

We first describe our methodologies for actively evaluating (MP)TCP over 5G and 4G. For carrying out measurement experiments, we have purchased 5G plans from Mobile and Unicom, and since Mobile has deployed much more base stations in our city than the other two operators combined, most of our experiments were conducted on Mobile's network. We carried out the study since May, 2020, six months after commercial 5G is available in our city. All the experiments were conducted between May and August, 2020, and had consumed a total amount of over 500 GB cellular data.

We rent a cloud server for the measurement experiments. Since inter-ISP links usually cause path inflations on China's Internet [26] and most of our experiments were conducted on Mobile's network, we rent a server from Mobile's cloud service[3], so that the UE and the server are in the same ISP. The cloud server is equipped with 4 CPU cores and has a dedicated Internet access bandwidth of 2 Gbps, which is far beyond the data rate of a 5G UE. The cloud server runs `Ubuntu` 18.04.4 with Linux kernel v5.3 that uses CUBIC [27] as the default TCP congestion control algorithm.

We use a Huawei Nova 6 smartphone as the 5G UE in our experiments. The phone runs `Android` 10 with CUBIC [27] as the default TCP congestion control algorithm. We confirm that despite 5G's high data rate,

[3] https://ecloud.10086.cn/

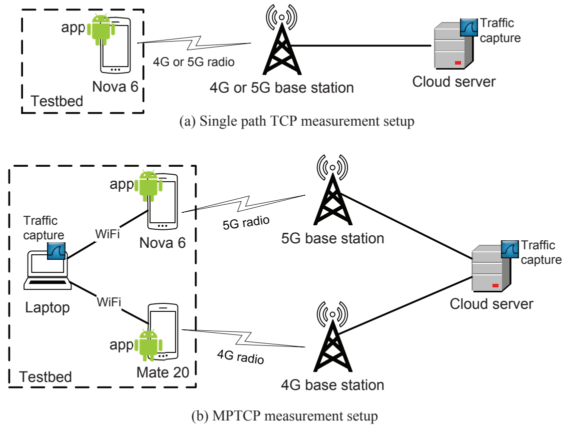

Figure 3: Measurement experiment setups for studying (a) single path TCP over 5G or 4G, and (b) MPTCP over 5G and 4G.

our UE will not be the performance bottleneck, as it is equipped with the high-end 8-core HiSilicon Kirin 990 SoC, HiSilicon Balong 5000 5G chipset, and 8 GB RAM.

For studying MPTCP on both 5G and 4G paths, which is currently not supported by any mobile devices, we use Nova 6 connecting to 5G to provide an `802.11n` WiFi hotspot at 2.4 GHz, and use a Huawei Mate 20 smartphone, which is only 4G capable, to provide another `802.11n` WiFi hotspot at 5 GHz. We then employ a Linux laptop equipped with two WiFi network interface cards to connect to the two WiFi hotspots, and transport data over an MPTCP connection composed of two subflows on 5G and 4G paths to the cloud server. We employ MPTCP v0.95.1 on both the laptop and the cloud server, and apply the "default" packet scheduler (i.e., lowest RTT first) and the olia congestion control algorithm [28].

We develop an Android smartphone app to trace cellular-related information. More specifically, the app logs 1) network type (4G or 5G) and cell ID, 2) signal strength in 5G CSI-RSRP [8] or 4G RSRP in every 200 ms, and 3) location and speed of the UE. The app also contains a TCP client that uploads to and downloads from the cloud server.

For determining the network type, we find that the standard Android API `getDataNetworkType()` is unable to identify 5G, as it returns LTE all the time despite that the UE has already connected to 5G. We also test the method in [6], but could not find the `nrState` and `isNrAvailable` fields in the string converted from the `ServiceState` object. Fortunately, we find that an API provided by Huawei EMUI, namely `getHwNetworkType`(), can tell the network type, i.e., 4G or 5G. After identifying the network type, we convert the `CellInfo` object into `CellInfoNr` for 5G or `CellInfoLte` for 4G, and call `getCsiRsrp()` or `getRsrp()` to obtain the signal strength under different networks respectively.

To sum up, Fig. 3 presents our setups for studying TCP and MPTCP over 5G and 4G. In particular, to study single-path TCP, we only use Nova 6 as in Fig. 3(a), which runs our app to transfer TCP traffic and logs 5G and 4G cellular information. For studying MPTCP, our testbed is composed of two smartphones and a laptop computer as shown in Fig. 3(b). Finally, the testbed can be either indoor or outdoor, and static or mobile (i.e., on a moving car) in different measurement experiments.

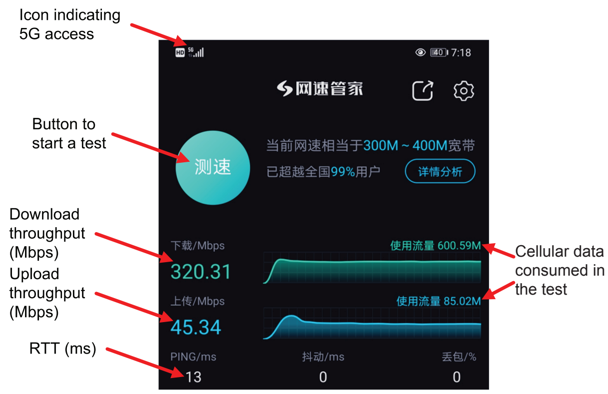

Figure 4: Speedtest app interface and an example of test result.

3.2. Passive Measurement Leveraging Speedtest

In addition to active measurement experiments, we also leverage a speedtest service to evaluate the commercial 5G networks with a passive approach. The service is provided at `5g.speedtest.cn`, and its infrastructure contains thousands of server nodes deployed in all the provinces and ISPs in mainland China. A speedtest works as follows: when a user launches the speedtest app on his 5G smartphone, the client automatically selects a proximate server node; the client first measures the latency to the node with ping, then it downloads from and uploads to the node with four parallel HTTP flows for measuring the throughputs. Fig. 4 demonstrates a speedtest result displayed by the app, which is composed of the download and upload throughputs, and the RTT between the UE and the selected server node. Clearly, if we can collect speedtest results reported by massive number of users all over the country, we are able to evaluate 5G from a nationwide perspective.

Although the `5g.speedtest.cn` website displays only ten recent speedtest results with highest download throughputs, however, we find that the website's raw source code actually contains dozens of results[4], which are updated frequently. We perform a few tests using our 5G UE under strong and weak signals (thus have high and low throughputs), and find that all the results appear on the website later. We conclude that the website encodes all or most of the recent test results.

To harvest the crowdsourced speedtest results on `5g.speedtest.cn`, we develop a crawler to constantly crawl the website every five minutes for 42 days from June 3 to July 15, 2020. We have collected 13,483 distinct speedtest results reported from 197 cities, which constitute 67% of the cities, in all the 31 provinces in China. The crowdsourced speedtest results enable us to perform a nationwide evaluation on commercial 5G.

4. Evaluation with Stationary UE

In this section, we study and evaluate the commercial 5G networks with stationary UE.

Table 2. RTTs between UE and Baidu frontend servers over Mobile and Unicom's 5G and 4G networks.

	Mobile (ms)	Unicom (ms)
5G	17.97 ± 2.80	10.07 ± 1.91
4G	28.75 ± 6.11	21.07 ± 2.01

4.1. Latency

We first compare end-to-end latencies of 5G and 4G. To this end, we connect the UE to Mobile's 5G network as in Fig. 3(a), and ping from our measurement app to the frontend server of `baidu.com` 50 times. We then turn off the 5G radio interface to make the UE connect to 4G, and repeat the pings. During the experiment, the UE is stationary and stays in the same cell that has both 5G and 4G accesses all the time. We also perform the same measurement with Unicom.

Table 2 presents the RTTs between our UE and Baidu's frontend servers over Mobile and Unicom's 5G and 4G networks. We exclude the first ping, as it could be extraordinarily long due to name resolution and radio resource control (RRC) [15][16]. We can see that for both operators, 5G reduces the latency dozens of milliseconds over 4G. Since in this experiment, an operator's 5G and 4G base stations share the same cell, the part of the network paths traversed by the ICMP packets outside the cellular networks should be same. Therefore, we conclude that the latency reduction comes from the flatten network architecture of 5G.

4.2. TCP Performance and Behavior

We then focus on TCP, which is the dominant connection-oriented transport-layer protocol, and examine how a TCP flow behaves over 5G. The experiment setup is shown in Fig. 3(a) with outdoor UE.

We run a TCP server at the cloud server, and develop a TCP client within our measurement app on the UE to connect to the server. In the downloading test, the server sends randomly generated data to the client as fast as it is allowed by TCP congestion control. Note that we transfer random data instead of downloading a file, as the latter approach might be bottlenecked by the cloud server's disk I/O. In the uploading test, the client also sends random data to the server as fast as it could. We keep the UE stationary and make sure that it stays in the same cell that provides 5G and 4G accesses.

Unlike [6], which employs up to 16 HTTP connections in their study, we use only one single TCP flow in the experiment. This is because we would like to investigate how 5G impacts TCP behaviors, therefore we use one TCP flow to avoid interference of other TCP flows from our UE that compete the channel between the UE and base station.

4.2.1. Download and upload throughput

We perform ten experiments, each lasting 60 seconds, with stationary UE in 5G and 4G networks. We use `pypcapkit` [29] and `tcptrace` [30] to process the traffics captured at the cloud server, and list the UE's averaged upload and download throughputs on Mobile's 5G and 4G networks in Table 3.

[4]The `5g.speedtest.cn` website uses CSS to control display and hide of HTML elements.

Table 3. TCP upload and download throughputs, data-in-flight and RTTs during bulk transfer over 5G and 4G networks.

	5G	4G
Upload throughput (Mbps)	96.23 ± 4.43	12.19 ± 0.11
Download throughput (Mbps)	335.23 ± 25.18	43.89 ± 3.15
Data-in-flight (kB)	1579.62 ± 304.93	1789.72 ± 208.44
RTT during bulk transfer (ms)	40.61 ± 5.02	333.44 ± 39.17

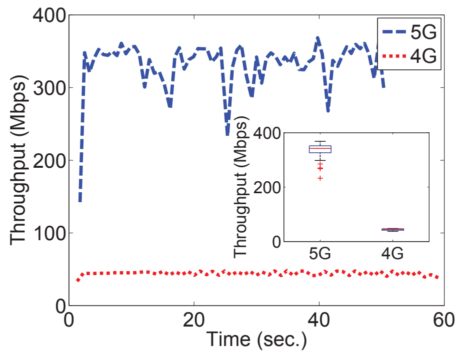

Figure 5: Comparison of download throughputs over 5G and 4G during a representative 60-second interval.

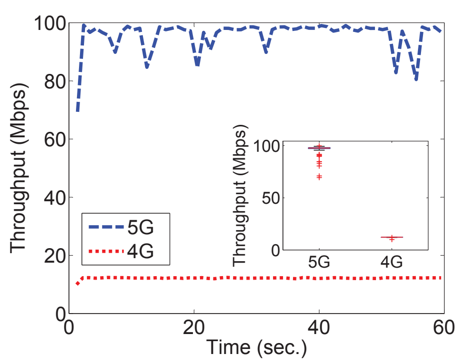

Figure 6: Comparison of upload throughputs over 5G and 4G during a representative 60-second interval.

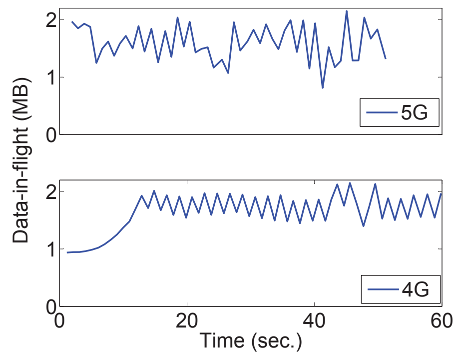

Figure 7: Comparison of in-flight data of download TCP flows in bulk transfer over 5G and 4G during a representative 60-second interval.

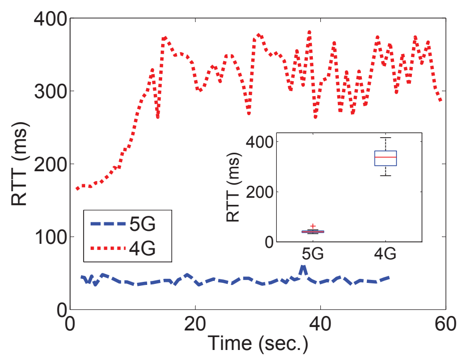

Figure 8: Comparison of RTTs of download TCP flows in bulk transfer over 5G and 4G during a representative 60-second interval.

We also select representative 60-second experiments over 5G and 4G, and plot the download throughputs in Fig. 5. From Table 3 and Fig. 5, one can see that 5G achieves an averaged TCP throughput exceeding 330 Mbps, which is nearly 8 times of the one 4G can provide. We recognize that the download throughput is dwarfed by the ones reported in [6] and [7]. There are two reasons for such a difference: First, the 5G networks studied in [6] employ mmWave, which by nature can achieve a much higher data rate comparing with the FR1 frequencies used by Mobile and other Chinese operators in this study. Second, the 5G networks studied in [6] and [7] are in their early deployments with few users, however, we carry out the experiment after 5G is fully commercial for over six months in our city, so our UE actually shares the base station with many other UEs in the same cell.

The upload throughputs of 5G and 4G for the same 60-second experiments are presented in Fig. 6. Again, Table 3 and Fig. 6 indicate that 5G brings a significant improvement, as the throughput is close to 8 times of 4G. In addition, Table 3, Fig. 5, and Fig. 6 confirm the previous observation that 5G has a larger performance variation than 4G [6].

4.2.2. RTT and in-flight data

We list in Table 3 the averaged amount of in-flight data and RTT during bulk transfer summarized from the ten experiments, and in Fig. 7 and Fig. 8, we compare the in-flight data and RTTs of the download TCP flows in the representative 60-second experiments over 5G and 4G. The two figures and Table 3 show that the

TCP flows in both networks have over 1.5 MB of data in flight, which is quite large. However, by examining the RTTs in Fig. 8, we find that the large amounts of in-flight data are caused by different reasons: In the 5G network, the RTT is low (around 40 ms) and close to the RTT measured with ping, which indicates that the queueing delay during bulk transfer is small. Since packets are not severely queued at any intermediate nodes along the path, the large amount of in-flight data is indeed being transmitted on wired and wireless links along the path. On the other hand, in the 4G network, the RTT is high (around 333 ms) and far beyond the ping RTT. Since the Internet part of the paths in both tests should be same, the large RTT indicates that the in-flight data is actually buffered within the 4G network, and the queueing delay caused by buffering constitutes most of the RTT.

To further understand the relation between RTT and in-flight data, we focus on the first 15 seconds after the download starts. As shown in Fig. 7 and Fig. 8, in the 4G network, it takes over 15 seconds for both RTT and in-flight data to grow, this is because as the download starts, data that has arrived at the 4G base station but has yet been transmitted to the UE gradually builds up a large queue. Note that such a *bufferbloat* phenomena has been long observed in 3G and 4G networks [15][16], as in 3G/4G, there exists a big mismatch between the wired and wireless data rates, and with 4G radio data rate far lower than the wired one, many packets are buffered on 4G base stations waiting to be transmitted. However, since 5G NR can provide a data rate that is close to or even higher than the wired network, packets will not be heavily buffered within the 5G base station before being transmitted to the UE.

The elimination of the long-existing bufferbloat phenomena in cellular networks suggests a subtle benefit of 5G, that is, short-lived flows will not be influenced by the large queue formed by long-lived and high-throughput flows when they coexist. For example, suppose a user streams an HD video and browses websites simultaneously over 4G or 5G. In 4G, the user will experience a longer page loading time (PLT) than usual, as the web browsing flows will have long queueing delays caused by the video streaming flow. However, PLT should be shorter in 5G with short RTTs for all the flows in the network.

4.3. MPTCP Performance and Behavior

Recently, MPTCP has been applied in mobile wireless networks for aggregating cellular and WiFi bandwidths [18]. As 5G will co-exist with 4G in near future, one potential application of MPTCP is to aggregate 5G and 4G capacities for aggregating bandwidths on both links or for enhancing communication reliability.

In this experiment, we investigate how MPTCP can aggregate 5G and 4G bandwidths. As shown in Fig. 3(b), we use a 5G and a 4G smartphones to provide two WiFi hotspots, and use an MPTCP-capable laptop computer to connect to the two WiFi hotspots for running MPTCP over 5G and 4G. All the devices are outdoor and stationary. Since the 5G download throughput far exceeds the typical `802.11n` data rate in real life, which is between 40−50 MB/s [31], to avoid bottleneck caused by WiFi, we examine the MPTCP upload throughout rather than the download one in this experiment. We capture traffics on both ends of the MPTCP connection, and use `pypcapkit` [29] and `mptcptrace` [32] for analyzing captured traffics.

We repeat the experiment five times, and summarize the results in Fig. 9 by presenting boxplots of the MPTCP connection's upload throughput, as well as the throughputs of the two MPTCP subflows on the 5G

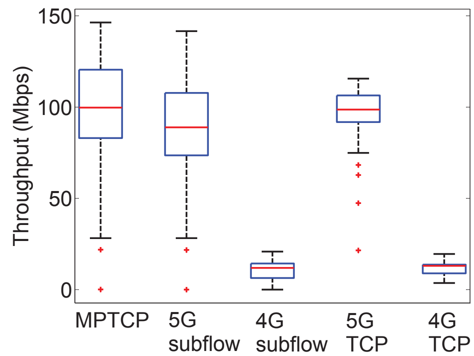

Figure 9: Comparison of throughputs of MPTCP, its 5G and 4G subflows, and 5G and 4G single path TCP flows.

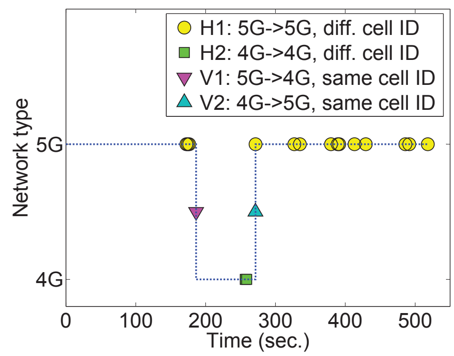

Figure 10: Four types of handoffs experienced by UE accessing to Mobile during a bicycle trip.

and 4G paths. For comparison, we also measure the upload throughput of single path TCP with the laptop connecting to only one WiFi hotspot provided by the 5G or 4G UE, and present the results. From the figure we can see that, MPTCP can effectively aggregate 5G and 4G capacities, and the MPTCP throughput is higher than either 5G or 4G single path TCP. Moreover, MPTCP utilizes 4G more poorly than 5G, as the ratio between the throughputs of the 4G subflow and the 4G single path TCP, which is 87.5%, is considerably lower than the ratio of 92.2% on 5G. We explain the observations with MPTCP's coupled congestion control, in which an MPTCP subflow expands its congestion window (cwnd) no more aggressively than a regular TCP flow; moreover, the olia algorithm favors the subflow with shorter RTT [28], therefore the 5G subflow, which has a much shorter RTT, has a higher bandwidth utilization than the 4G subflow.

5. Evaluation with Mobile UE

In this section, we study and evaluate the commercial 5G networks with UE in mobile environments.

5.1. Handoff

We first examine UE handoffs under 5G. Currently, a mobile UE in a cellular network experiences two kinds of handoffs: horizontal handoff and vertical handoff. A horizontal handoff happens when the UE switches from one cell to another under the same network type, caused by reasons such as mobility; and a vertical handoff means that the UE changes its network type (e.g., from 4G to 5G). As one can see in Fig. 1, since the operators

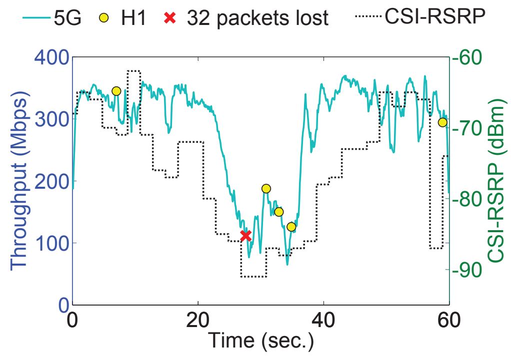

Figure 11: Download throughput, signal strength, $H1$ handoffs, and pakcet loss burst of 5G UE under mobile environment.

are unable to provide 5G signal coverage in all the areas, it is inevitable for a mobile UE to be demoted to 4G when moving to some areas where the 5G signal is weak or unavailable, and be promoted back to 5G after the signal is available and becomes strong again. As a matter of fact, both vertical and horizontal handoffs are frequently experienced by our UE in the mobile environment, because the cells are small.

As an example, in Fig. 10, we present the handoffs that our UE, which accesses to Mobile, has experienced during a 9-minute bicycle trip in our campus (recall that the measurement app records network type and cell ID every 200 ms). From the figure we can see that there are four types of handoffs: the horizontal handoffs between different cells of 5G or 4G, which we refer to as $H1$ and $H2$ handoffs; and the vertical handoffs for demoting UE from 5G to 4G and promoting from 4G to 5G in the same cell, which are referred to as $V1$ and $V2$ handoffs respectively. To confirm that the four types of handoffs are primitive (i.e., a handoff is not composed of other handoffs), we continuously record UE's network type and cell ID (rather than on 200 ms intervals) when the UE is performing the four types of handoffs, and find they are indivisible. We also analyze the handoffs with the UE accessing to Unicom, and have the same result.

Note our observation of primitive 5G-to-5G horizontal handoff is very different from [6], which reports that such a handoff (i.e., the $H1$ handoff) is accomplished with a sequence of other horizontal and vertical handoffs. The difference suggests that commercial 5G networks are diverse and still evolve rapidly around the world.

5.2. TCP Performance and Behavior

We then evaluate TCP over 5G under a mobile environment. The experiment setup is shown in Fig. 3(a) with the UE in a high-speed moving car. More specifically, we select a highway passing our city and drive a car along it back and forth. By referring our city's 5G signal coverage map as in Fig. 1 and detecting by ourself, we find that most part of the selected route is covered by Mobile's 5G signal, but there are still some sections where 5G is unavailable. In our experiment, we use the 5G UE to download from the cloud server with a single TCP flow, and the UE also records signal strength, car speed, and UE handoffs. The experiment lasts 62.7 minutes, during which the UE connects to 5G for overall 48.6 minutes. In the experiment, we have recorded 443, 94, 38, and 38 $H1$, $H2$, $V1$, and $V2$ handoffs respectively. The car speed is between $76-113$ km/h.

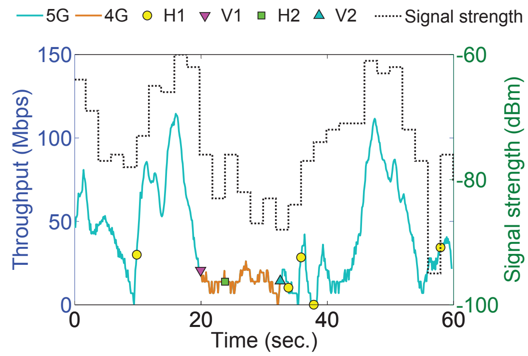

Figure 12: Download throughput, signal strength (CSI-RSRP for 5G and RSRP for 4G), and four types of handoffs of 5G UE under mobile environment.

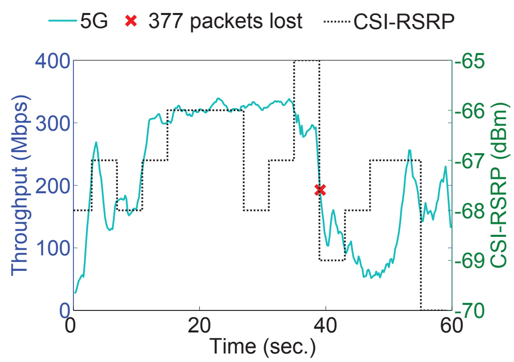

Figure 13: Download throughput, signal strength, and packet loss burst of 5G UE under mobile environment.

5.2.1. Throughput and impacting factor

We study the UE's download throughput and investigate the impacting factors. Our first observation is that under the highly mobile environment, the UE's download throughput varies significantly. For example, in Fig. 11, Fig. 12, and Fig. 13, each presenting the UE's download throughputs in a 60-second interval, the throughput varies from over 300 Mbps to less than 100 Mbps within a few seconds. Clearly, such a large throughput variation will confuse the bandwidth-sensitive applications, and lead to inconsistent user experiences.

Unlike the previous study [6], which identifies the frequent UE handoff as the major cause for throughput reductions, we find that the $H1$ handoff, i.e., the switching from one 5G cell to another, does not cause throughput drops. For example, on Fig. 11 we can find $H1$ handoffs happening when the download throughput is maintained at above 300 Mbps; and in Fig. 12, we also find that $H1$ handoffs occur while the download throughput rapidly increases. In addition, no TCP packets are lost on UE handoffs. From these observations, we confirm that horizontal handoff hardly impacts TCP performance over 5G.

After carefully examining the experiment data, we find that the UE's download throughput is sensitive to the signal strength perceived by the UE. In particular, when the signal becomes too weak, the throughput drops significantly, and the throughput increases when the signal becomes strong. For example, in Fig. 11 and Fig. 12, when the signal's CSI-RSRP drops below -80 dBm, the throughput decreases over 200 Mbps and 100 Mbps respectively. Moreover, the download throughput varies in a way that is correlated with the signal strength, and the Pearson's correlation coefficients between the throughput and the signal strength in Fig. 11

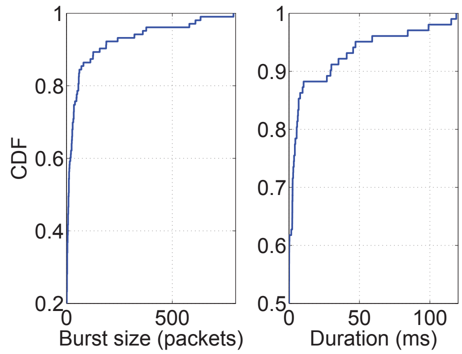

Figure 14: Distributions of sizes (in # of packets) and durations of packet loss bursts.

and Fig. 12 are 0.7506 and 0.7410 respectively. Although we do not observe such a strong correlation in Fig. 13, however, one can see that the throughput drops severely when the signal suddenly becomes weak at the 39^{th} second.

As for the vertical handoff, after analyzing the 38 pairs of $V1$ and $V2$ handoffs recorded in the experiment, we find that typically, a $V1$ handoff happens after the 5G signal has already become weak and the throughput has already dropped. For example, in Fig. 12, the period when the UE switches to 4G nearly overlaps with the period that the signal strength is below -80 dBm.

5.2.2. Retransmission and packet loss

We then focus on retransmissions during the data transfer. We have identified a total number of 8,139 retransmitted TCP packets in the captured trace, constituting less than 0.04% of the total traffic. All the retransmissions happen when the UE connects to 5G.

After analyzing the traces, we find that all the retransmissions are caused by retransmission timeout (RTO) at the sender. Note that an RTO in TCP can be either a normal RTO, which is caused by a really lost packet, or a spurious RTO, where the packet is not really lost but just delayed. We apply the method in [33] to differentiate the RTOs, and find that only 724 RTOs are spurious, while the remaining 7,415 RTOs, constituting 91.10% of the retransmissions, are caused by really lost packets.

Since most of the RTOs are caused by packet losses, we investigate the reasons behind. To this end, we group the RTOs into clusters with a threshold of 50 ms, that is, we group two consecutive RTOs into the same cluster if they happen within 50 ms. As a result, we have grouped the 7,411 normal RTOs into 103 clusters, and their total duration is only 1.96 seconds. Fig. 14 presents the distributions of the sizes and durations of the clusters. From the figure we can see that packets are lost in bursts: typically, a burst lasts only a few milliseconds, but contains tens to hundreds of lost packets, and in many cases, packets with continuous sequence numbers are lost.

To understand the packet loss bursts, we refer to the signal strength. We find that many bursts happen when the signal perceived by the UE is poor or becomes weak abruptly. For example in Fig. 11, a burst containing 32 lost packets happens when the signal's CSI-RSRP is -91 dBm, which is ranked only as "fair" for mobile signal strength; and Fig. 13 contains a burst with 377 lost packets, which happens at the moment when

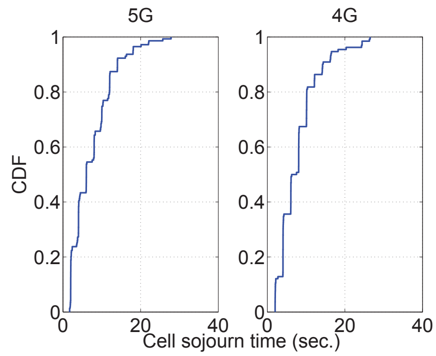

Figure 15: Cell sojourn time for mobile UEs in 5G and 4G networks.

the signal's CSI-RSRP suddenly drops from −65 dBm to −69 dBm.

Note that our observation here is very different from the previous study on 4G [17], which reports that in a highly mobile environment (UE is on a high-speed train moving at a speed over 300 km/h), most of the RTOs are spurious because of RTO estimation inaccuracy. However, in our study, majority of the RTOs are caused by really lost packets. The reasons are twofold: First, spurious RTOs are caused by large and varying-sized in-network buffer in 4G networks [17], but as we have seen in Sec. 4, 5G generally does not form large in-network buffer. Second, when the signal perceived by the UE becomes poor and can no longer sustain high data rate, the in-network buffer quickly gets filled and overflows, leading to packet loss bursts.

As discussed in previous works [15][16], 3G/4G cellular networks prefer to maintain large in-network buffers to absorb sudden losses in capacities, but large amount of buffered TCP packets lead to bufferbloat and result in issues such as inaccurate RTT estimation and spurious retransmission. Nevertheless, our findings in this and previous sections reveal that 5G faces a new dilemma for buffering TCP packets within the network: on one hand, for stationary UE, maintaining a large in-network buffer seems no longer necessary as fewer packets are accumulated in the buffer, thanks to 5G NR's high data rate; but on the other hand, when signal strength and data rate change dramatically in mobile environment, the in-network buffer is not large enough to accommodate such variation and easily overflows.

5.3. MPTCP Performance and Behavior

One major motivation for introducing MPTCP to mobile environment is to enhance communication reliability. For example, previous studies confirm that by employing MPTCP on both WiFi and cellular networks, QoE of interactive applications such as Siri can be greatly improved [20].

In this section, we examine whether such benefit exists by applying MPTCP over 5G and 4G under a mobile environment. As shown in Fig. 3(b), we set up two WiFi hotspts using 5G and 4G UEs, and upload from the MPTCP-capable laptop computer to the MPTCP-capable cloud server. All the devices were on a car moving on the same route as in the previous section, and we record 5G and 4G cellular information such as cell ID and signal strength on both UEs.

Since 5G UE connects to 4G when 5G signal is weak or unavailable, we focus our analysis on an 18.2-minute period, during which both 5G and 4G are available and each network carries an MPTCP subflow.

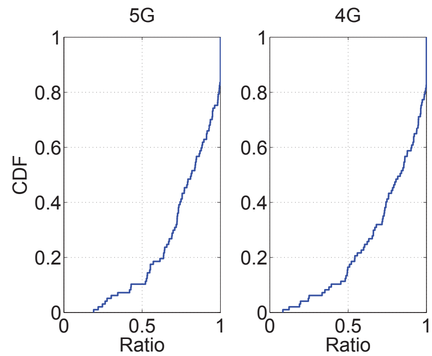

Figure 16: Ratios between t_o and $(t_3 - t_1)$ (left), and t_o and $(t_4 - t_2)$ (right).

We first look at the 5G and 4G base stations our UEs connect to. During the 18.2-minute period, the 5G and 4G UEs have connected to 144 and 132 distinct base stations respectively. In Fig. 15, we present the durations when our UE stays in a cell, for all the 5G and 4G cells. We can see that the distribution curves in two figures are stair like, with a stair width of about 2 seconds. This observation suggests that Mobile makes handoff decisions every 2 seconds, for UEs in both 5G and 4G networks. We also find that a UE stays in a cell for a very short time, with the median cell sojourn time of 6.0 seconds in 5G and 6.9 seconds in 4G; in addition, for considerable number of cells, the UEs stay for no more than 2 seconds, which is the minimum cell sojourn time given the 2-second handoff decision granularity. One concern is that 2 seconds may be too long for very fast moving UE, such as one on a high-speed train [17]. Finally, we note that the two distributions in Fig. 15 are very close, suggesting that the 5G and 4G base stations have similar signal coverages.

We then focus on the 5G and 4G cells that share the same cell ID. We have identified 98 such cells, in which our 5G and 4G UEs stay for 14.2 and 14.0 minutes respectively during the 18.2-minute period. For each shared cell, we compute the duration when both UEs stay in the cell as

$$t_o = \min(t_3, t_4) - \max(t_1, t_2)$$

where t_1 and t_2 are the moments of the $V1$ handoffs that the 5G and 4G UEs enter into the cell, and t_3 and t_4 are the moments that they leave the cell. We then compute the ratio between t_o and $(t_3 - t_1)$ for 5G, and the ratio between t_o and $(t_4 - t_2)$ for 4G, which are the ratios of cell sojourn time when both UEs are in the same cell. We present the distributions of the ratios in Fig. 16. From the figure we can see that, for most of the time, an MPTCP connection's 4G and 5G subflows are carried by base stations in the same cell. Recall that under the NSA model, the co-locating 5G and 4G base stations share the same core network within the cell, our observation suggests that when the shared network facility are faulty, MPTCP over 5G and 4G can not improve the communication reliability despite that an additional subflow on a different cellular network path is employed.

To testify our argument, we examine the throughputs achieved by the MPTCP connection's subflows on 5G and 4G paths, and we present the throughputs in an exemplary 60-second period in Fig. 17. From the figure we can see that between the 4^{th} and 12^{th} second, both subflows' throughputs drop to zero, and further investigation shows that during the interval, both our 5G and 4G UEs are in the cell with ID `96665900`.

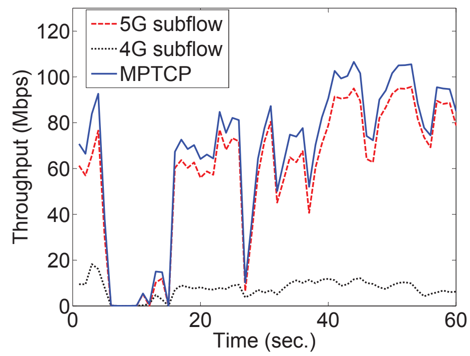

Figure 17: Throughputs of 5G and 4G subflows as well as the aggregated MPTCP throughput during an exemplary 60-second period.

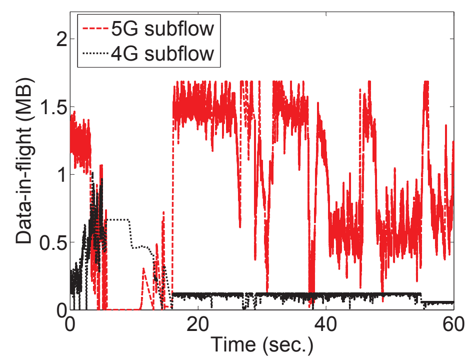

Figure 18: In-flight data of 5G and 4G subflows during an exemplary 60-second period.

We further analyze in-flight data on 5G and 4G subflows for the 60-second period in Fig. 18. We can see that MPTCP prefers the 5G subflow most of the time because of its shorter RTT. But between the 4^{th} and 12^{th} second, when packets are continuously lost on the 5G subflow, the MPTCP connection resorts to the 4G subflow and cause many packets to buffer at the 4G base station within a few seconds.

Our observation suggests that unlike previous studies that use MPTCP over uncorrelated paths (e.g., terrestrial and satellite paths in [21] and [22]) to enhance communication reliability, under the NSA 5G, employing MPTCP over 5G and 4G paths does not necessarily improve communication reliability, as very likely, the 5G and 4G base stations carrying subflows of an MPTCP connection are in the same cell and impacted by the same problem such as congestion or malfunction within the cell.

6. A Nationwide Evaluation

In this section, we leverage the speedtest service at `5g.speedtest.cn`, and analyze performance of the commercial 5G networks of China from a nationwide perspective.

6.1. Performance over the Nation

As described in Sec. 3.2, `5g.speedtest.cn` displays the speedtest results from users over the country. By constantly crawling the website, we have collected 13,483 crowdsourced speedtest results reported in June and July, 2020 from 197 cities in all the 31 provinces in China. Although `5g.speedtest.cn` does not reveal on

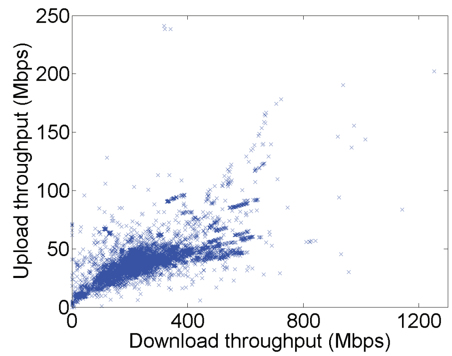

Figure 19: Download and upload throughputs of all the speedtest results over the nation in scatter plot.

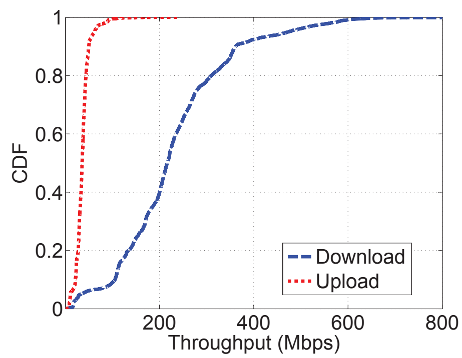

Figure 20: Distributions of 5G download and upload throughputs in speedtest results.

which operator's 5G network a test is performed, however, we believe that this does not influence our analysis, as our active measurement studies on Mobile and Unicom suggest that the major operators in China adopt similar technologies to build their 5G networks, and have similar characteristics.

In Fig. 19, we present the download and upload throughputs of the speedtest results. One can see that the throughputs measured over the country vary significantly: the download throughput ranges from 1.58 Mbps to 1253.61 Mbps, and the upload throughput varies from 0.21 Mbps to 241.19 Mbps. Another observation is that statistically, the download and upload throughputs are linearly correlated, with a Pearson's correlation coefficient of 0.7016.

In Fig. 20, we plot the distributions of the download and upload throughputs of the 13,483 speedtest results. From the figure we can see that 80% of the tests report download and upload throughputs within 300 Mbps and 45 Mbps respectively. On the other hand, even at the 10^{th} percentile, which corresponds to a download throughput of 105.66 Mbps and an upload throughput of 21.84 Mbps, 5G can still provide a significant improvement over 4G, as the latter's download throughput hardly exceeds 50 Mbps in practice [34].

6.2. Cities and Provinces

We then focus on performances of the 5G networks in individual cities. Fig. 21 presents the four cities with the highest 5G download throughputs, and we also present the nationwide median download throughput, which is 216.54 Mbps, on the figure for comparison. We find that all the four cities are small or medium in size, and have throughputs much higher than the nationwide median. Moreover, the speedtest results show that many

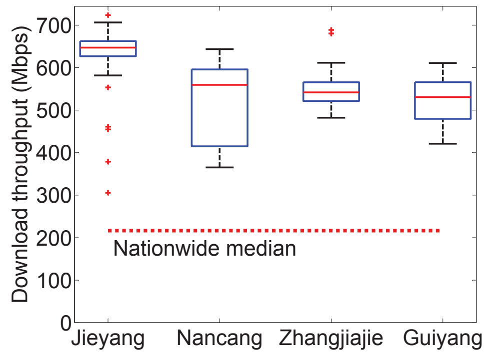

Figure 21: 5G download throughputs of the top four cities of Jieyang, Nancang, Zhangjiajie, and Guiyang.

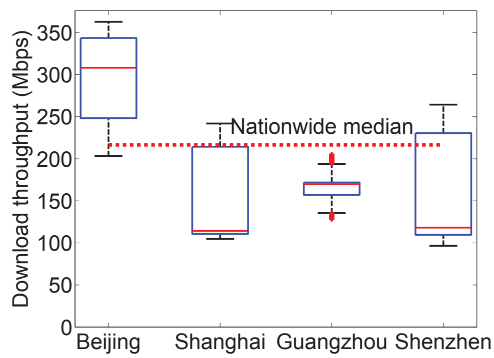

Figure 22: 5G download throughputs of the four largest metropolises of Beijing, Shanghai, Guangzhou, and Shenzhen.

small-medium cities have relatively higher throughputs.

On the other hand, in Fig. 22, we present the download throughputs of the four largest metropolises in China: Beijing, Shanghai, Guangzhou, and Shenzhen. We find that only Beijing has a throughput higher than the nationwide median, while all the other three cities have much lower throughputs.

We explain the variation of the 5G network performances among different cities as the following. In the large populous metropolises, the 5G service is available for several months, and people in these cities, who are generally younger and more wealthy, are more willing to switch to 5G by purchasing 5G plans and 5G smartphones; on the other hand, in small and medium-sized cities, 5G is available for a shorter time, and people in these cities are less eager to switch to 5G. As a result, the 5G base stations in the populous metropolises are connected by more UEs, and provide relatively lower throughputs comparing with many small cities across the nation.

In Fig. 23, we present the averaged 5G download throughputs for all the 31 provinces in mainland China. Most provinces have a throughput between 200−400 Mbps, but some industrialized and populous eastern provinces, such as Guangdong and Jiangsu, have mean throughputs below 200 Mbps, while some rural western provinces like Qinghai, Guizhou, and Xinjiang have averaged throughputs above 400 Mbps. The regional difference suggests that even with the FR1 frequencies, there are a lot of spaces for the operators to improve their 5G networks. In particular, recall that under the NSA model, the 5G base stations are co-located with the 4G ones, thus the density of the 5G base stations are indeed restricted by the 4G deployment. In future,

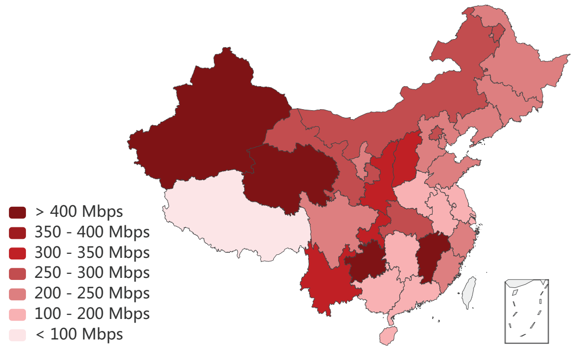

Figure 23: Averaged 5G download throughput of different provinces in China.

it is expected that the ultra-dense small cell deployment [35] would be employed to improve the capacities of the 5G networks in populous metropolitan areas.

7. Conclusion

In this work, we performed one of the first measurement studies on commercial 5G. By conducting controlled experiments with off-the-shelf smartphone, we evaluated performances of 5G networks, and analyzed behaviors of TCP and MPTCP over 5G. Our findings show that in addition to providing a much higher throughput than 4G, the commercial 5G has unique characteristics including shorter latency, large performance variation, and bursts of packet losses in mobile environment. Our findings suggest that commercial 5G faces a dilemma in maintaining in-network buffer for TCP, and due to the 4G co-locating NSA model, applying MPTCP over 5G and 4G does not necessarily improve communication realisability. We also leveraged a speedtest service to perform a nationwide evaluation on the commercial 5G in China. By leveraging crowdsourced speedtest results, we find that there exist considerable regional differences on 5G performances across the nation, and populous metropolitan areas have relatively inferior performances. Our findings provide valuable insights for people to improve upper-layer protocols and applications over 5G.

Acknowledgements

This work was supported in part by the National Natural Science Foundation of China under Grants 61672486 and 61672485, and in part by the Strategic Priority Research Program of the Chinese Academy of Sciences under Grant XDC02070300.

References

[1] Minimum requirements related to technical performance for IMT-2020 radio interface(s) (Nov. 2017).

[2] The 5G economy: How 5G will contribute to the global economy (Nov. 2019).
URL https://www.qualcomm.com/media/documents/files/ihs-5g-economic-impact-study-2019.pdf

[3] China MIIT: 690 thousand 5G base stations deployed, connecting over 160 million UEs (in Chinese) (Oct. 2020).
URL https://finance.sina.com.cn/tech/2020-10-23/doc-iiznezxr7663660.shtml

[4] C. Paasch, O. Bonaventure, Multipath TCP, ACM Queue 12 (2) (2014) 1–12.

[5] A. Langley, A. Riddoch, A. Wilk, A. Vicente, et al., The QUIC transport protocol: Design and Internet-scale deployment, in: Proc. of SIGCOMM, Los Angeles, CA, USA, 2017.

[6] A. Narayanan, E. Ramadan, J. Carpenter, Q. Liu, Y. Liu, F. Qian, Z.-L. Zhang, A first look at commercial 5G performance on smartphones, in: Proc. of WWW, Taipei, Taiwan, 2020.

[7] D. Xu, A. Zhou, X. Zhang, G. Wang, X. Liu, C. An, Y. Shi, L. Liu, H. Ma, Understanding operational 5G: A first measurement study on its coverage, performance and energy consumption, in: Proc. of SIGCOMM, Virtual Event, NY, USA, 2020.

[8] 3gpp specification series: 38series, accessed: Jul. 5, 2020.
URL https://www.3gpp.org/DynaReport/38-series.htm

[9] Huawei announces peak cell throughput exceeds 20 Gbit/s in C-band test of China's 5G technology R&D trial (Sep. 2017).
URL https://www.huawei.com/en/news/2017/9/Huawei-5G-20Gbps

[10] List of 5G NR networks, accessed: Jul. 5, 2020.
URL https://en.wikipedia.org/wiki/List_of_5G_NR_networks

[11] H. Ekström, Non-standalone and standalone: two standards-based paths to 5G (Jul. 2019).
URL https://www.ericsson.com/en/blog/2019/7/standalone-and-non-standalone-5g-nr-two-5g-tracks

[12] Z. Zhu, W. Lu, L. Zhang, N. Ansari, Dynamic service provisioning in elastic optical networks with hybrid single-/multi-path routing, J. Lightw. Technol. 31 (2013) 15–22.

[13] L. Gong, X. Zhou, X. Liu, W. Zhao, W. Lu, Z. Zhu, Efficient resource allocation for all-optical multicasting over spectrum-sliced elastic optical networks, J. Opt. Commun. Netw. 5 (2013) 836–847.

[14] F. P. Tso, J. Teng, W. Jia, D. Xuan, Mobility: A double-edged sword for HSPA networks: A large-scale test on Hong Kong mobile HSPA networks, IEEE Trans. on Parallel & Distributed Systems 23 (10) (2012) 1895–1907.

[15] H. Jiang, Y. Wang, K. Lee, I. Rhee, Tackling bufferbloat in 3G/4G networks, in: Proc. of IMC, Boston, MA, USA, 2012.

[16] J. Huang, F. Qian, Y. Guo, Y. Zhou, Q. Xu, Z. M. Mao, S. Sen, O. Spatscheck, An in-depth study of lte: Effect of network protocol and application behavior on performance, in: Proc. of SIGCOMM, Hong Kong, China, 2013.

[17] L. Li, K. Xu, D. Wang, C. Peng, K. Zheng, R. Mijumbi, Q. Xiao, A longitudinal measurement study of TCP performance and behavior in 3G/4G networks over high speed rails, IEEE/ACM Trans. on Networking 25 (4) (2017) 2195–2208.

[18] S. Deng, R. Netravali, A. Sivaraman, H. Balakrishnan, WiFi, LTE, or both? measuring multi-homed wireless internet performance, in: Proc. of IMC, Vancouver, BC, Canada, 2014.

[19] A. Nikravesh, Y. Guo, F. Qian, Z. M. Mao, S. Sen, An in-depth understanding of multipath TCP on mobile devices: Measurement and system design, in: Proc. of MobiCom, New York, NY, USA, 2016.

[20] Q. D. Coninck, O. Bonaventure, Tuning multipath tcp for interactive applications on smartphones, in: Proc. of IFIP Networking, Zurich, Switzerland, 2018.

[21] A. Abdelsalam, M. Luglio, C. Roseti, F. Zampognaro, Analysis of bandwidth aggregation techniques for combined use of satellite and xDSL broadband links, Intl. J. of Satellite Commun. & Netw. 37 (2) (2019) 76–90.

[22] A. Bujari, M. Luglio, C. E. Palazzi, M. Quadrini, C. Roseti, F. Zampognaro, A virtual PEP for web optimization over a hybrid backhaul, IEEE Commun. Mag. 58 (10) (2020) 42–48.

[23] Q. D. Coninck, O. Bonaventure, Multipath QUIC: Design and evaluation, in: Proc. of CoNEXT, Incheon, Republic of Korea, 2017.

[24] R. S. Mogensen, C. Markmoller, T. K. Madsen, T. Kolding, G. Pocovi, M. Lauridsen, Selective redundant mp-quic for 5g mission critical wireless applications, in: Proc. of IEEE VTC2019-Spring, Kuala Lumpur, Malaysia, 2019.

[25] M. Zhang, M. Polese, M. Mezzavilla, J. Zhu, S. Rangan, S. Panwar, M. Zorzi, Will TCP work in mmWave 5G cellular networks?, IEEE Communication Magazine 57 (1) (2019) 65–71.

[26] Y. Tian, R. Dey, Y. Liu, K. W. Ross, Topology mapping and geolocating for China's Internet, IEEE Trans. on Parallel & Distributed Systems 24 (9) (2013) 1908–1917.

[27] S. Ha, I. Rhee, L. Xu, CUBIC: a new TCP-friendly high-speed TCP variant, ACM SIGOPS Operating Systems Review 42 (5) (2008) 64–74.

[28] R. Khalili, N. Gast, M. Popovic, J.-Y. L. Boudec, MPTCP is not pareto-optimal: performance issues and a possible solution, IEEE/ACM Trans. on Networking 21 (5) (2013) 1651–1665.

[29] pypcapkit, accessed: Jul. 5, 2020.
URL `https://pypi.org/project/pypcapkit/`

[30] S. Ostermann, tcptrace official homepage, accessed: Jul. 5, 2020.
URL `http://www.tcptrace.org/`

[31] What is the actual real-life speed of wireless networks?, accessed: Apr. 24, 2021.
URL `https://www.speedguide.net/faq/what-is-the-actual-real-life-speed-of-wireless-374`

[32] B. Hesmans, O. Bonaventure, Tracing multipath tcp connections, in: Proc. of SIGCOMM, Chicago, IL, USA, 2014.

[33] F. Vacirca, T. Ziegler, E. Hasenleithner, An algorithm to detect TCP spurious timeouts and its application to operational UMTS/GPRS networks, Computer Networks 50 (16) (2006) 2981–3001.

[34] 4G upload and download througputs of three major mobile operators (in Chinese) (Nov. 2018).
URL `https://new.qq.com/omn/20181118/20181118A0BXMH.html`

[35] F. Tseng, L. Chou, H. Chao, J. Wang, Ultra-dense small cell planning using cognitive radio network toward 5G, IEEE Wireless Communication 22 (6) (2015) 76–83.