Will TCP Work in mmWave 5G Cellular Networks?

Menglei Zhang, Michele Polese, Marco Mezzavilla, Jing Zhu, Sundeep Rangan, Shivendra Panwar, and Michele Zorzi

Abstract

The vast available spectrum in the millimeter-wave (mmWave) bands offers the possibility of multi-gigabit-per-second data rates for fifth generation cellular networks. However, mmWave capacity can be highly intermittent due to the vulnerability of mmWave signals to blockages and delays in directional searching. Such highly variable links present unique challenges for adaptive control mechanisms in transport layer protocols and end-to-end applications. This article considers the fundamental question of whether TCP, the most widely used transport protocol, will work in mmWave cellular systems. The article provides a comprehensive simulation study of TCP considering various factors such as the congestion control algorithm, including the recently proposed TCP BBR, edge vs. remote servers, handover and multi-connectivity, TCP packet size, and 3GPP-stack parameters. We show that the performance of TCP on mmWave links is highly dependent on different combinations of these parameters, and identify the open challenges in this area.

Introduction

End-to-end connectivity over the Internet largely relies on transport protocols that operate above the network layer. The most widely used transport protocol is the Transmission Control Protocol (TCP), designed in the 1980s [1] to offer reliable packet delivery and sending rate control to prevent congestion in the network. Reliability is accomplished with the receiver's acknowledgments (ACKs) fed back to the sender, which retransmits packets if needed, while rate control is achieved by dynamically adjusting the congestion window (i.e., the maximum amount of data that the sender can transmit without receiving ACKs). Several congestion control (CC) algorithms have been proposed in order to improve the goodput (defined as the application layer throughput) and latency of TCP over different types of networks [2].

However, the next generation of cellular networks will present new challenges for TCP,[1] specifically related to millimeter-wave (mmWave) links in the radio access network, which exhibit an erratic propagation behavior. This technology is seen as a promising enabler for the fifth generation (5G) targets of multi-gigabit-per-second data rates and ultra-low latency [3], but the end-to-end performance perceived by the user will eventually depend on the interaction with transport protocols such as TCP. Some recent studies [4, 5] have highlighted that the extreme variability of the signal quality over mmWave links yields either degraded TCP goodput and very low utilization of the resources at mmWave frequencies, or, in the presence of link-layer retransmissions, high goodput at the price of high latency. Moreover, in [4] it is shown that the bufferbloat phenomenon (i.e., the increase in latency that is caused by excessive buffering) emerges as a consequence of the presence of large buffers in the network.

Our goal is to answer the question: *Will TCP work in mmWave 5G cellular networks?* To achieve this aim, we compare the performance of different TCP CC algorithms over simulated 5G end-to-end mmWave networks considering a high-speed train and an urban macro Third Generation Partnership Project (3GPP) deployment, as further described in the next section. Our detailed simulation study demonstrates that the performance of TCP over mmWave depends critically on several aspects of the network.

1. Edge vs. Remote Server: By comparing the end-to-end performance at varying server's locations, we show that for a shorter control loop (i.e., when the server is placed at the cellular network edge, TCP can react faster to link impairments).

2. Handover and Multi-Connectivity: Due to unreliability of individual mmWave links, dense deployments of small cells with fast handover protocols are critical in maintaining stable connections and avoiding TCP timeouts.

3. CC Algorithms: With remote servers, we observe higher performance variations across different CC algorithms, while the difference is almost negligible with edge servers. Overall, TCP Bottleneck, Bandwidth, and Round-Trip (BBR) outperforms loss-based TCP in terms of both rate and latency.

4. TCP Packet Size: We quantitatively compare the benefits of transmitting larger TCP packets in Long-Term Evolution (LTE) vs. mmWave networks, and show that, given the fluctuating gigabit-per-second data rates offered at mmWave frequencies, a larger packet size provides a faster growth of the congestion window and higher achievable rate.

5. Radio Link Control (RLC) Buffer Size: We analyze TCP performance over small and large buffers. While the TCP goodput degradation caused by buffer overflow in undersized buffers is difficult to mitigate, the problem of bufferbloating (i.e., large buffer occupancy leading to delays)

The authors consider the fundamental question of whether TCP, the most widely used transport protocol, will work in mmWave cellular systems. They provide a comprehensive simulation study of TCP considering various factors such as the congestion control algorithm, including the recently proposed TCP BBR, edge vs. remote servers, handover and multi-connectivity, TCP packet size, and 3GPP-stack parameters.

[1] In this work, we focus on TCP since it is the dominant transport protocol in use today. One possible avenue for future work is to consider the User Datagram Protocol (UDP), which, however, shifts the burden of retransmissions and flow control to a higher layer, introducing similar problems as those related to TCP.

Menglei Zhang, Marco Mezzavilla, Sundeep Rangan, and Shivendra Panwar are with New York University; Michele Polese and Michele Zorzi are with the University of Padova; Jing Zhu is with Intel Corporation.

Digital Object Identifier: 10.1109/MCOM.2018.1701370

0163-6804/19/$25.00 © 2019 IEEE

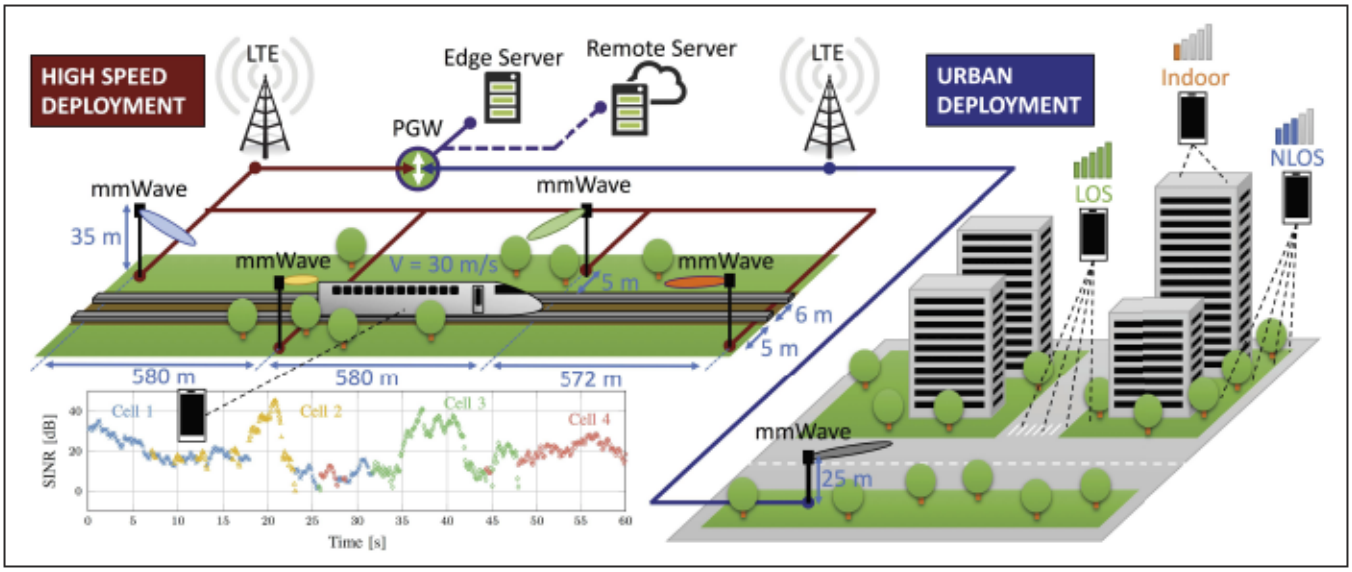

Figure 1. High-speed and urban deployment scenarios.

can be approached by appropriately designing cross-layer algorithms [6].

The rest of the article is organized as follows. We first describe the scenarios of interest in the following section. Then we list the main features of the CC algorithms considered in this study. We report the main results and observations, and draw our conclusions.

5G Deployment Scenarios

In order to assess how TCP will perform in mmWave cellular networks, we consider two of the most challenging scenarios among those specified by 3GPP in [7], that is, a high-speed train and a dense urban scenario, represented in Fig. 1. They were studied using the ns-3-based mmWave end-to-end simulation framework described in [8], which models radio access, the core network, and the 3GPP channel for the mmWave band with spatial correlation in mobility scenarios. Moreover, the protocol stack simulated by [8] also features retransmissions at both the medium access control (MAC) layer, with hybrid automatic repeat request (HARQ), and the RLC layer, using the acknowledged mode option.

High-Speed Scenario: In this scenario, shown on the left side of Fig. 1, we test the performance of TCP over a channel that varies frequently in time and under realistic mobility conditions. Multiple next generation node bases (gNBs) provide coverage to the railway, which is mostly line of sight (LoS): even if the current gNB is blocked by obstacles placed between gNBs 2 and 3, the user equipment (UE) can quickly perform a handover to another LoS gNB. The gNBs are at a height of 35 m, with an intersite distance of 580 m. The train moves at a speed of 108 km/h, and, as a result, the channel experienced by the UE varies very quickly because of severe fading and the Doppler effect, and, on a longer timescale, due to obstacles, as shown in the signal-to-interference-plus-noise ratio (SINR) plot of Fig. 1. We use the channel tracking and mobility scheme described in [9], which features fast and locally coordinated handovers for devices that are dual-connected to an mmWave gNB and a sub-6 GHz gNB (e.g., an LTE base station).

Dense Urban Scenario: In this deployment, shown on the right side of Fig. 1, we study the fairness of TCP flows over multiple UEs with different channel conditions. A single mmWave gNB placed at a height of 25 m serves a group of 10 users moving at walking speed. They are located in different positions in order to account for a mixture of channel conditions: four UEs are in LoS, thus perceiving a very high SINR, four are in non-LoS (NLoS), and the last two are inside a building, so the received power is additionally attenuated by the building penetration loss.

For both scenarios we consider two deployments of the TCP server that acts as the endpoint of the connection. The first is a traditional setup in which the server is hosted in a remote data center, with a minimum round-trip time (RTT) on the order of 40 ms, accounting for the latencies of both the core network and the public Internet. The second is a mobile edge cloud (MEC) scenario [10], in which the server is located close to the gNBs with smaller latency (on the order of 4 ms).

TCP Congestion Control Protocols

In this section, we describe the CC protocols and the TCP performance enhancement techniques considered in this article.

TCP Congestion Control Algorithms

We study four most commonly used CC algorithms.

TCP NewReno has been the default algorithm for the majority of communication systems. In the congestion avoidance phase, the congestion window `cwnd` is updated after the reception of every ACK. The update is based on the additive increase multiplicative decrease (AIMD) design: `cwnd` is increased by summing a term $\alpha/$`cwnd` for each received ACK, and divided by a factor β for each packet loss. For NewReno these parameters are fixed to $\alpha = 1$ and $\beta = 2$.

HighSpeed TCP is designed for high bandwidth-delay product (BDP) networks, in which NewReno may exhibit a very slow growth of `cwnd`. HighSpeed behaves the same as NewReno when `cwnd` is small, but when it exceeds a pre-

defined threshold, the parameters α, β become functions of `cwnd` in order to maintain a large congestion window. Moreover, the window growth of NewReno and HighSpeed depends on the ACK reception rate; thus, a shorter RTT increases the ACK frequency and further speeds up the window growth.

TCP CUBIC, instead, increases `cwnd` over time, without considering the ACK reception rate but rather capturing the absolute time since the last packet loss and using a cubical increase function for `cwnd`. It has been designed to increase the ramp-up speed of each connection while maintaining fairness with other users.

TCP BBR, recently presented by Google [11], measures bottleneck bandwidth and round-trip propagation time, or BBR, to perform congestion control. BBR to perform congestion control. It strives to match the sending rate to the estimated bottleneck bandwidth by pacing packets and setting the congestion window to `cwnd` gain × BDP, where the `cwnd` gain is a factor (= 2) that is used to balance the effects of delayed, stretched, and aggregated ACKs on bandwidth estimation.

TCP Performance Enhancement Techniques

The performance of TCP has been the object of many studies over the last decades, and, besides new CC algorithms, many other techniques have been proposed and deployed either at the endpoints of the connection (TCP sender and receiver) or inside the network.

In the case of multiple packet losses, the TCP selective acknowledgment (SACK) option [12] allows the receiver to inform the sender which packets were received successfully so that the sender can retransmit only those which were actually lost. This dramatically improves the efficiency of the TCP retransmission process.

Active queue management (AQM) schemes [13], instead, are deployed in network devices (e.g., routers, gateways, gNBs) to control the behavior of their queues and buffers. The size of these buffers plays an important role in the end-to-end performance. If the buffer is too small, many packets may be dropped when the buffer is full, according to the drop tail policy. Conversely, if the buffer is too large, bufferbloat occurs [13]. AQM techniques can be deployed at the buffers to drop packets before the queue is full so that the TCP sender can proactively react to the congestion which could arise in the near future.

Finally, there are some techniques that are typically used in combination with wireless links. The first is the usage of link-layer retransmissions between the gNB and the UE so that the losses on the channel are masked from TCP. This helps increase the goodput; however the end-to-end latency also increases, as shown in [5]. Another technique which is often used in wireless networks is proxying [2] (i.e., the connection is split into two at some level in the mobile network, e.g., at the gateway with the Internet) at the gNB), and different CC techniques are deployed over the two parts of the connection.

TCP Performance in the 3GPP Scenarios

In the following paragraphs we report the performance of the TCP CC algorithms presented previously over the 5G mmWave deployment scenarios described above, focusing on both goodput and latency. The results are averaged over multiple independent simulation runs, so the confidence intervals are small (they are, however, not shown to make the figures easier to read). In all the simulations, we use full buffer traffic with the TCP SACK option and disable the TCP delayed ACK mechanism; thus, each received packet will generate an ACK. The minimum retransmission timeout is set to 200 ms.

High-Speed Deployment Scenario

In this scenario we compare different combinations of the RLC buffer size B and the maximum segment size (MSS) P with a single TCP connection from the UE. For both the remote and edge server deployments, the RLC buffer is 10 or 100 percent of the BDP computed considering the maximum achievable data rate (3 Gb/s) and the minimum latency, that is, B equals 1.5 or 15 MB for the remote server deployment, and 0.15 or 1.5 MB for the edge server. We also consider two different MSSs, that is, a standard MSS of 1400 bytes (1.4 kB) and a large MSS of 14,000 bytes (14 kB). The goodput of saturated UDP traffic is also provided as a reference for the maximum achievable rate, as shown in Fig. 2.

Notice that, thanks to the mobility management scheme based on dual connectivity and fast secondary cell handover, and despite the high mobility of the scenario, we never observed a TCP connection reset due to an outage (i.e., even if the closest two base stations are blocked, the UE is still capable of receiving signals from other nearby gNBs). Therefore, even if blockage events are still possible, in a scenario with a dense deployment (according to 3GPP guidelines), it is possible to provide uninterrupted connectivity to the final user [14].

In the following paragraphs we provide insights on the effects of the different parameters on TCP performance over mmWave at high speed.

Impact of the Server Deployment: Loss-based TCP benefits from the shorter control loop related to an edge server deployment, as shown by comparing Figs. 2a and 2b. With the latter, indeed, the differences between the maximum goodput of the loss-based TCP versions are less marked, since the faster reaction time makes up for the differences among them. Moreover, the RTT difference between the large and small RLC buffers is lower in absolute terms (milliseconds with edge server vs. tens of milliseconds with remote server), but the ratio is approximately the same. However, for CUBIC and HighSpeed with the smallest buffer configuration, the goodput is lower with the edge than with the remote server: there is a 30 percent loss with the smallest MSS, and 50 percent with the largest one. In this case, indeed, the buffer size is very small (i.e., B = 0.15 MB), thus incurring buffer overflows,[2] which reduce the sending rate.

Impact of the Congestion Control Algorithm: The CC algorithm has a stronger impact in the remote server scenario. The best performance, in terms of goodput, is achieved by BBR with large buffer size, but it is still 400 Mb/s lower than the maximum achievable rate. Moreover, as observed in [4, 5], high goodput values also correspond to higher end-to-end latency. Howev-

The performance of TCP has been the object of many studies over the last decades, and, besides new CC algorithms, many other techniques have been proposed and deployed either at the endpoints of the connection (TCP sender and receiver) or inside the network.

[2] With large MSS, just 11 packets are enough to cause a buffer overflow.

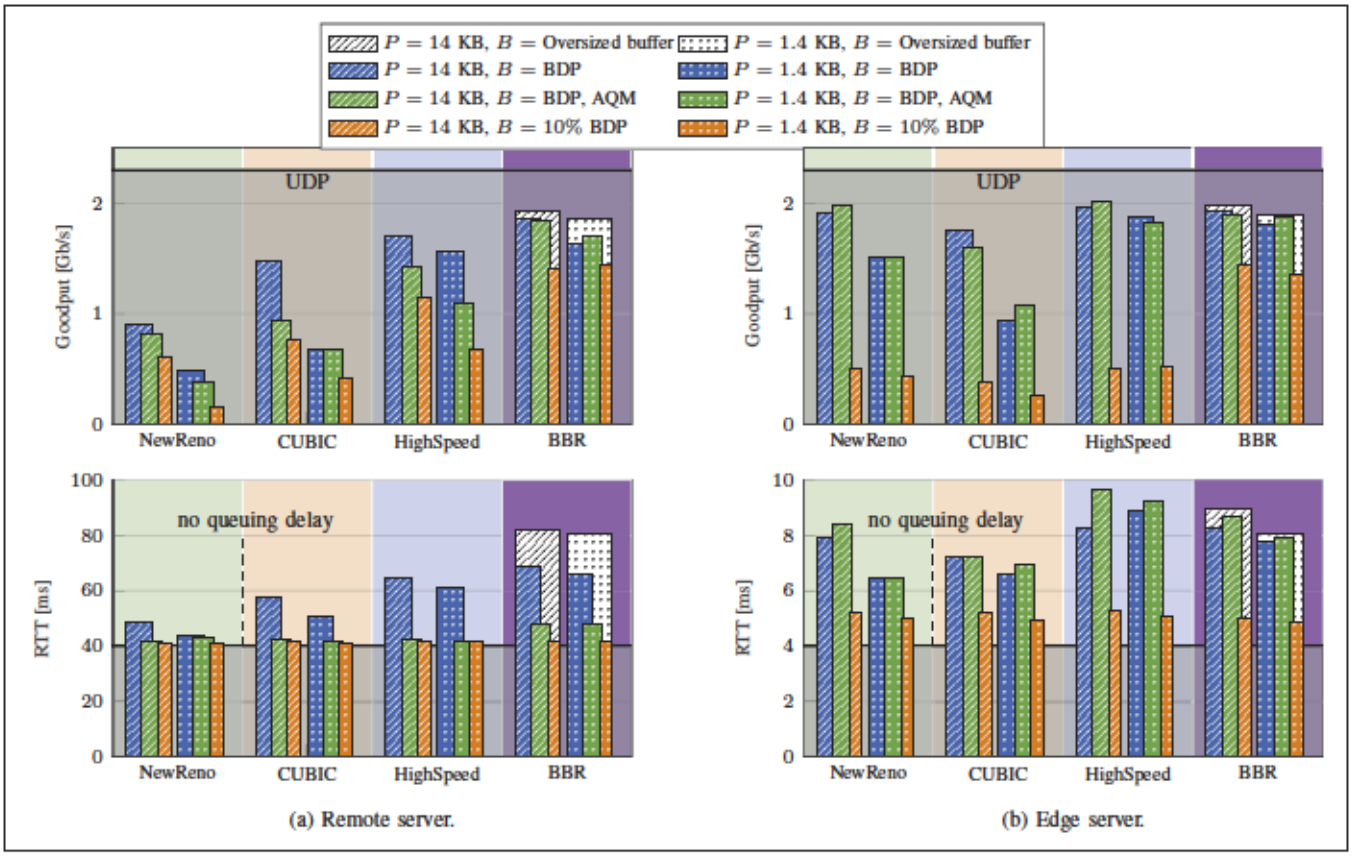

Figure 2. Goodput and RTT for the high-speed train scenario, with the remote and edge servers for different combinations of the buffer size and the MSS.

er, with small buffers, BBR produces the highest goodput (especially in the edge server scenario), with a latency comparable to loss-based TCP. BBR, indeed, regulates its sending rate to the estimated bandwidth and is not affected by packet loss (i.e., the congestion window dynamics of BBR), presented in Fig. 3a, which matches the SINR plot in Fig. 1.

However, the loss-based versions of TCP cannot adjust their congestion window fast enough to adapt to the channel variations and perform worse than BBR, especially with small buffer, as seen in Fig. 3a. Among them, TCP HighSpeed provides the highest goodput because of the aggressive window growth in the high BDP region. TCP CUBIC performs better than NewReno in the remote server case, but worse in the edge server case. This is because CUBIC's window growth is not affected by the ACK rate, and therefore is more reliable over long RTT links.

Impact of the MSS: The MSS does not affect the performance of BBR, which probes the bandwidth with a different mechanism, whereas, for loss-based TCP, the impact of the MSS on the goodput is remarkable.[3] The standard MSS of P = 1.4 kB exhibits much worse performance compared to a larger MSS of P = 14 kB. This happens because, in congestion avoidance, the congestion window increases by MSS bytes every RTT, if all the packets are received correctly and delayed ACK is not used, so the smaller the MSS, the slower the window growth. Hence, the MSS dictates the congestion window's growth, which is particularly critical in mmWave networks for two main reasons:

- The mmWave peak capacity is at least one order of magnitude higher than in LTE, so the congestion window will take a much longer time to reach the achievable link rate. In this case, we can gain in performance by simply using a larger MSS, as depicted in Fig. 2.
- In addition, the channel fluctuations in the mmWave band will result in frequent quality drops, thus often requiring the congestion window to quickly ramp up to the link capacity to avoid underutilizing the channel.

Large MSS — mmWave vs. LTE: Aimed at better illustrating why larger packets are particularly important in 5G mmWave networks, we also provide a performance comparison against LTE in the same scenario,[4] and report in Table 1 and Fig. 3 detailed results focusing on the impact of the TCP MSS on the congestion window growth and, consequently on the goodput of the system. Only a single user is placed in the high-speed train scenario; thus, the drops in the congestion window are due to the worsening of the channel quality and not to contention with other flows. Figure 3 shows that the loss-based TCP congestion window with a small MSS grows very slowly in congestion avoidance, and consequently loss-based TCP does not fully exploit the available bandwidth during the intervals in which the received signal has a very high SINR (i.e., at t = 20 s and t = 40 s, as shown in Fig. 1). The large MSS helps speed up the congestion window's growth, which translates into higher goodput. Conversely, the goodput degradation associated with small packets is less relevant in LTE networks, given that the goodput is limited by the available bandwidth and not

[3] Typically, TCP segments are mapped to multiple MAC/PHY data units, which complicates the dependence between a larger value of the TCP MSS and the correspondingly higher packet error probability over the wireless link. This nontrivial relationship, which would deserve a study by itself, has been properly captured in our numerical results.

[4] For the LTE setup, the small buffer represents 50 percent of the BDP (i.e., 0.08 and 0.2 MB for edge and remote server, respectively), because a 10 percent BDP buffer would be too small to protect from random fluctuations of the channel.

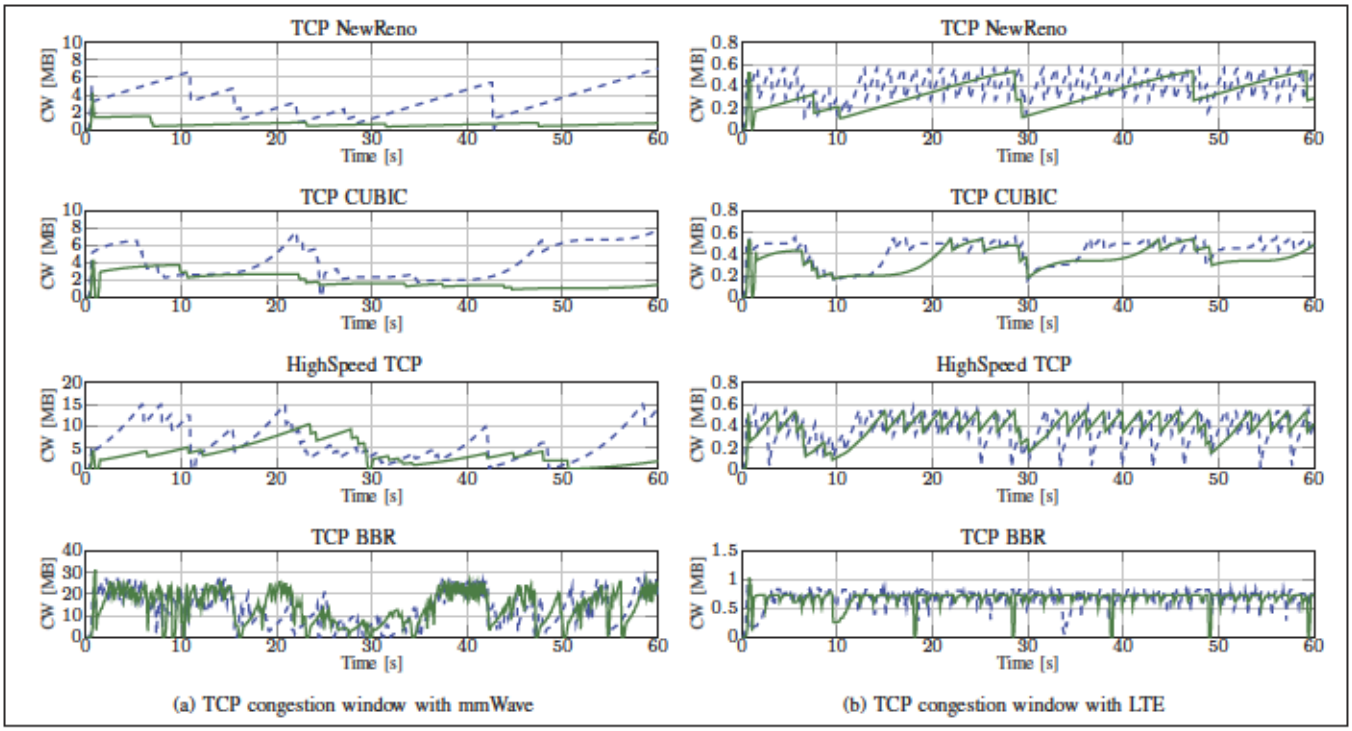

Figure 3. Congestion window evolution over time for different CC algorithms. The scenario is configured with remote servers and small RLC buffers: a) TCP congestion window with mmWave; b) TCP congestion window with LTE.

by the congestion window increase rate. These trends are reflected in Table 1. Among all loss-based TCP versions, only HighSpeed increases its congestion window fast enough even when transferring small packets. As a consequence, the goodput gain obtained with large MSS values is much smaller.

Large packets introduce an additional benefit: due to (1) a reduced TCP/IP header overhead and (2) a reduced number of TCP ACKs, there will be more available downlink/uplink resources, resulting in higher goodput values.

This solution may not be practical in an end-to-end network in which the maximum transmission unit (MTU) is not entirely in the control of the mobile network provider and is typically dictated by the adoption of Ethernet links (i.e., an MTU of 1500 bytes). In contrast, in MEC scenarios, in which the whole network is deployed by a single operator, it is possible to support a large MSS thanks to Ethernet jumboframes [15].

Impact of the Buffer Size and AQM: The buffer size is also critical for the performance of TCP. As shown in Fig. 2, large buffers generally yield higher goodput, because the probability of buffer overflow is smaller, and they offer more effective protection against rapid and temporary variations of the mmWave channel quality. However, when a large buffer is coupled with loss-based TCP, the latency inevitably increases. Conversely, smaller buffers provide lower latency at the price of lower goodput.

For loss-based TCP, an intermediate solution is provided by applying AQM to the largest buffer, especially in the remote server scenario. Controlled delay management (CoDel) is used as the default AQM in our simulation because of its simple configuration. It controls the buffering latency by dropping packets when the minimum queuing delay within an interval, starting from 100 ms, is larger than 5 ms. CoDel is picked as an example to show the trade-off between latency and goodput by using AQM. Our goal for this article is not to select the best AQM scheme or optimize AQM, which in itself is a very interesting topic and could be considered for future research. As shown in Fig. 2a, the goodput with the AQM option is larger than that with the smallest buffer, and in some cases (i.e., for the smallest packet size) is comparable to that of the BDP buffer without AQM, which in general yields the highest goodput. However, the latency is equivalent to the one associated with the small buffer, which is the lowest. In the edge server scenario the TCP control loop is short (the RTT is 4 ms), and the reaction to congestion is quick. Hence, its performance is indeed equivalent to having BDP buffers without AQM.

		Remote server		Edge server	
	Buffer size	BDP	10% BDP	BDP	10% BDP
TCP NewReno	LTE	1.06	1.17	0.80	0.65
	mmWave	1.81	3.96	1.27	1.15
TCP CUBIC	LTE	1.06	1.15	1.03	0.89
	mmWave	2.2	1.83	1.89	1.44
HighSpeed TCP	LTE	1.08	0.9	0.94	0.95
	mmWave	1.09	1.69	1.05	0.98
TCP BBR	LTE	1.00	0.96	1.02	0.82
	mmWave	1.14	0.97	1.06	1.06

Table 1. Ratio between the goodput achieved with P = 14 kB and with P = 1.4 kB for different configurations of the simulated scenario.

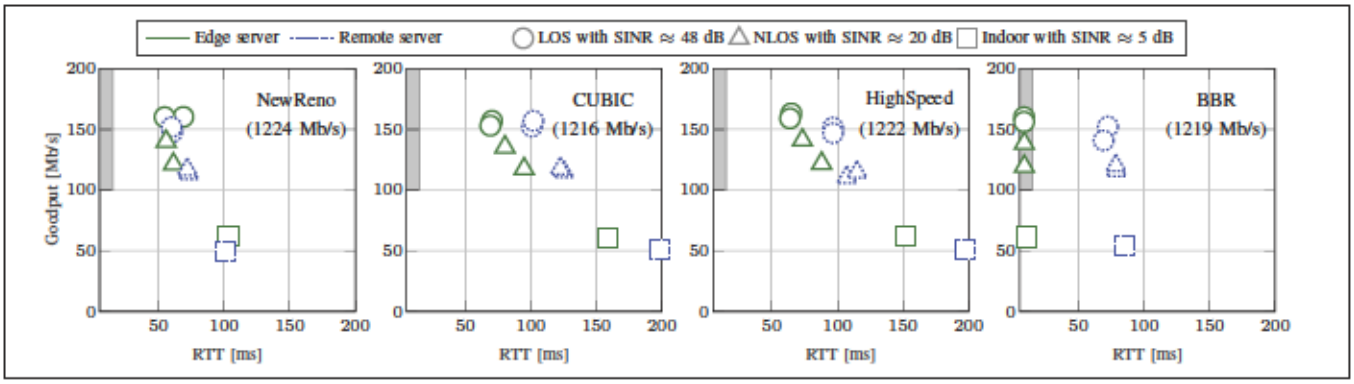

Figure 4. Goodput vs RTT for ten UEs in the Urban Scenario, for different choices of the CC algorithm.

	Loss-based	MSS impacts goodput	Summary	Considerations over 5G
TCP NewReno	Yes	Yes	Remote server: lowest goodput	Need to move servers to the edge
TCP CUBIC	Yes	Yes	Edge server: lowest goodput	Need to increase MSS
HighSpeed TCP	Yes	Only remote server	Big buffer: high goodput and high latency	Need to mitigate latency with AQM
TCP BBR	No	No	Big buffer: high goodput and high latency *Small buffer: small rate reduction and low latency*	Small buffer is preferred Performs well over mixed UE conditions

Table 2. Results of the CC algorithms over 5G deployments.

BBR tries to solve this problem without modifying the buffers in the routers by maintaining a congestion window equal to twice the BDP regardless of packet loss, as shown by Fig. 3. As a consequence, latency is only doubled in large buffers, and the goodput is slightly reduced in small buffers. These behaviors are also observed in the oversized and 10 percent BDP buffer cases of Fig. 2.

Urban Deployment Scenario

In this scenario we consider 10 UEs attached to a single mmWave gNB. In particular, we position four UEs in LoS conditions, four in NLoS and two inside a building. The average SINR for each channel condition is provided in Fig. 4. Notice that, with low blockage density and walking speed, the channel condition is relatively stable over time. For each UE pair, one is connected to an edge server, and the other is connected to a remote server. In this way, it is possible to test the performance of TCP over a mixture of different conditions. The gNB uses a round-robin (RR) scheduler, so the resource management at the base station does not have an impact on the fairness among different flows. All the UEs use the same TCP version. We consider a standard MSS of 1400 bytes and an RLC buffer size of 1.5 MB for each UE.

Figure 4 shows the average cell goodput (labeled in parentheses) and the goodput-latency trade-off for each type of user, separately, and for each CC algorithm, in order to evaluate the fairness and the overall performance of different TCP versions with respect to different user channel conditions.

All CC algorithms achieve the same average cell goodput, and similar goodput per UE. However, the RTT varies a lot among the CC algorithms. The reason is that all UEs use the same buffer size regardless of their channel conditions and network latency. As a consequence, the RLC buffer size may be large for some UEs, such as those at the edge. Therefore, the CC algorithms that adopt a more aggressive window growth policy, such as CUBIC and HighSpeed, yield much higher latency. For the loss-based TCP, NLoS and indoor UEs suffer from higher latency: given the same buffer size and backhaul data rate, a reduced bottleneck bandwidth results in an increased queueing delay in the buffers, until TCP settles to a steady state phase. BBR, instead, limits the congestion window to twice the estimated BDP, and results in a maximum latency of 2 × minimum RTT. We also draw a gray area in the plot representing a typical 5G application requirement (i.e., goodput greater than 100 Mb/s and delay lower than 10 ms). In this scenario, among all CC algorithms, only BBR meets this requirement for the UEs connected to an edge server, and only under good channel conditions.

Conclusions

The massive but intermittent capacity available at mmWave frequencies introduces new challenges for all layers of the protocol stack, including TCP, the most widely used transport protocol. The interplay between congestion control algorithms and mmWave channel quality fluctuations makes the topic particularly complex, and represents the key driver behind this work. We have carried out a thorough simulation campaign, based on ns-3, across 3GPP-inspired scenarios, whose results are summarized in Table 2. The main findings and some relevant research questions are listed as follows:

- TCP benefits from a shorter control loop, where the server is placed at the cellular network edge and can react faster to link impairments. *Should we (re)consider splitting TCP at some point?*
- Moreover, when the RTT is high, loss-based TCP underutilizes the mmWave capacity, while those based on congestion (e.g., BBR) show improved performance by estimating

the bandwidth of mmWave links. This means that new approaches based on *more refined abstractions of the end-to-end network* can be studied for highly variable and high-data-rate mmWave links.

- Multi-connectivity and smart handovers, supported by advanced beamtracking and beamswitching techniques, will result in robust TCP connections. *How densely should we deploy mmWave cells? How do we support backhaul for densely deployed mmWave cells?*
- We show very clearly how loss-based TCP over mmWave bands can greatly benefit from using larger datagrams. *Has the time come to break the legacy MTU value* by natively supporting larger packets in a wider set of networks?
- Finally, it is well known that buffer size must scale proportionally to BDP to achieve maximum TCP goodput. However, it is very challenging to properly dimension the buffers for mmWave links, given the rapid bandwidth variations between LoS and NLoS conditions, and to protect from link losses without introducing bufferbloat. Given the low latency requirement and massive available bandwidth, *is it beneficial to trade bandwidth for lower latency,* for example, by running BBR over small RLC buffer configurations?

We believe that these insights will stimulate further exploration of this important topic, which is essential to fully exploit the true potential of mmWave communications. Moreover, the observations provided by this initial simulation-based study can be used to guide the design of experimental activities, which are necessary to further validate the challenges that mmWave links pose to TCP, and to test novel techniques to improve the end-to-end user experience in mmWave cellular networks.

Acknowledgments

This work was supported in part by NSF grants 1302336, 1564142, and 1547332, NIST 70NANB17H166, SRC, and the industrial affiliate members of NYU WIRELESS.

References

[1] J. Postel, "Transmission Control Protocol," IETF RFC 793, Sep. 1981.
[2] K. Liu and J. Y. Lee, "On Improving TCP Performance over Mobile Data Networks," *IEEE Trans. Mobile Comp.*, vol. 15, no. 10, Oct. 2016, pp. 2522–36.
[3] M. Xiao *et al.*, "Millimeter Wave Communications for Future Mobile Networks," *IEEE JSAC*, vol. 35, no. 9, Sept. 2017, pp. 1909–35.
[4] M. Zhang *et al.*, "Transport Layer Performance in 5G mmWave Cellular," *Proc. IEEE INFOCOM Wksps.*, Apr. 2016, pp. 730–35.
[5] M. Polese, R. Jana, and M. Zorzi, "TCP and MP-TCP in 5G mmWave Networks," *IEEE Internet Computing*, vol. 21, no. 5, Sept. 2017, pp. 12–19.
[6] M. Zhang *et al.*, "TCP Dynamics over mmWave Links," *Proc. IEEE 18th Int'l. Wksp. Signal Processing Advances in Wireless Commun.*, July 2017.
[7] 3GPP, "TR 38.913, Study on Scenarios and Requirements for Next Generation Access Technologies, V14.1.0," 2017.
[8] M. Mezzavilla *et al.*, "End-to-End Simulation of 5G mmWave Networks," *IEEE Commun. Surveys & Tutorials*, vol. 20, no. 3, 3rd qtr. 2018, pp. 2237–63.
[9] M. Polese *et al.*, "Improved Handover Through Dual Connectivity in 5G mmWave Mobile Networks," *IEEE JSAC*, vol. 35, no. 9, Sept. 2017, pp. 2069–84.
[10] P. Mach and Z. Becvar, "Mobile Edge Computing: A Survey on Architecture and Computation Offloading," *IEEE Commun. Surveys & Tutorials*, vol. 19, no. 3, 3rd qtr., 2017, pp. 1628–56.
[11] N. Cardwell *et al.*, "BBR: Congestion-Based Congestion Control," *Queue*, vol. 14, no. 5, 2016, p. 50.
[12] M. Mathis *et al.*, "TCP Selective Acknowledgment Options," IETF RFC 2018, Oct. 1996.
[13] Y. Gong *et al.*, "Fighting the Bufferbloat: On the Coexistence of AQM and Low Priority Congestion Control," *Computer Networks*, vol. 65, June 2014, pp. 255–67.
[14] M. Polese *et al.*, "Mobility Management for TCP in mmWave Networks," *Proc. 1st ACM Wksp. Millimeter-Wave Networks and Sensing Systems 2017*, ser. mmNets '17, 2017, pp. 11–16.
[15] M. Mezzavilla *et al.*, "Evaluation of Jumboframes Feasibility in LTE Access Networks," *Proc. IEEE ICC*, June 2013, pp. 5964–68.

Biographies

MENGLEI ZHANG received his B.S. degree in electrical engineering from Nanjing University of Science and Technology, China, in 2013, and his M.S. degree in electrical engineering in 2015 from New York University (NYU) Tandon School of Engineering, Brooklyn, where he is currently working toward a Ph.D. degree in electrical engineering with Prof. Rangan. His research interests include wireless communications, channel modeling, congestion control, and system level simulation.

MICHELE POLESE [S'17] received his B.Sc. (2014) and M.Sc. (2016) in telecommunication engineering from the University of Padova, Italy. Since October 2016 he has been a Ph.D. student in the Department of Information Engineering at the University of Padova. He visited NYU and AT&T Labs, Bedminster, New Jersey. His research interests focus on the analysis and development of protocols and architectures for 5G mmWave networks.

MARCO MEZZAVILLA is a research scientist at the NYU Tandon School of Engineering. He received his B.Sc. (2007), M.Sc. (2010) in telecommunications engineering, and Ph.D. (2013) in information engineering from the University of Padova, Italy. He is serving as a reviewer for IEEE conferences, journals, and magazines. His research interests include design and validation of communication protocols and applications of 4G/5G wireless technologies, multimedia traffic optimization, radio resource management, spectrum sharing, convex optimization, cognitive networks, and experimental analysis.

JING ZHU [M'04, SM'12] received B.S. and M.S. degrees from Tsinghua University, China, in 1999 and 2001, respectively, and a Ph.D. in 2004 from the University of Washington, Seattle, all in electrical engineering. He is currently a principal engineer at Intel Corporation. His main research interests are system design, performance optimization, and applications for heterogeneous wireless networks, including 4G/5G cellular systems, high-density wireless LANs, and mobile ad hoc networks.

SUNDEEP RANGAN [F'15] is a professor of electrical and computer engineering at NYU and associate director of NYU WIRELESS. He received his Ph.D. from the University of California, Berkeley in electrical engineering. In 2000, he co-founded (with four others) Flarion Technologies, a spinoff of Bell Labs that developed Flash OFDM, the first cellular OFDM data system. It was acquired by Qualcomm in 2006, where he was a director of engineering prior to joining NYU in 2010.

SHIVENDRA PANWAR [S'82, M'85, SM'00, F'11] is a professor of electrical and computer engineering at NYU. He is the director of the New York State Center for Advanced Technology in Telecommunications (CATT), a member of NYU WIRELESS, and the faculty director of the NY City Media Lab. He was co-awarded the IEEE Communication Society's Leonard G. Abraham Prize, and is a co-author of the *IEEE Multimedia Communications* and ICC Best Papers for 2011 and 2016, respectively.

MICHELE ZORZI [F'07] is with the Information Engineering Department of the University of Padova. His present research interests focus on various aspects of wireless communications. He was Editor-in-Chief of *IEEE Wireless Communications* from 2003 to 2005, *IEEE Transactions on Communications* from 2008 to 2011, and, at present, *IEEE Transactions on Cognitive Communications and Networking*. He served as a Member-at-Large of the ComSoc Board of Governors from 2009 to 2011, and as Director of Education and Training from 2014 to 2015.

The observations provided by this initial simulation-based study can be used to guide the design of experimental activities, which are necessary to further validate the challenges that mmWave links pose to TCP, and to test novel techniques to improve the end-to-end user experience in mmWave cellular networks.