Mobility Management for TCP in mmWave Networks

Michele Polese*, Marco Mezzavilla◇, Sundeep Rangan◇, Michele Zorzi*
*Department of Information Engineering, University of Padova, Italy, emails: {polesemi, zorzi}@dei.unipd.it
◇NYU Wireless, New York University, Brooklyn, NY, USA, emails: {mezzavilla, srangan}@nyu.edu

ABSTRACT

Communication at millimeter wave (mmWave) frequencies will likely be a cornerstone for next generation 5G cellular networks. However, providing mobility support for end-to-end applications in mmWave cellular systems is challenging due to the relatively small coverage area of individual cells, and rapid channel dynamics caused by blockage and beam-tracking. This paper presents a comprehensive performance evaluation of TCP on top of mmWave cellular systems with mobility management, detailed modeling of the channel dynamics, and end-to-end network architectures. We show how an efficient mobility management scheme in a dense network deployment can dramatically improve the performance of TCP in terms of both throughput and latency in mobile scenarios with blocking. The study also reveals that, even with fast mobility management, TCP throughput is extremely sensitive to the end-to-end delay with implications on both core network and server location.

CCS CONCEPTS

• **Networks → Transport protocols**; **Mobile networks**; *Network simulations*;

KEYWORDS

TCP; mobility; mmWave; 5G; dual connectivity; ns-3

1 INTRODUCTION

The next generation of mobile networks (5G) targets a massive performance gain with respect to 4G, with multi-gigabit-per-second throughput, ultra-low latency (below 10 ms) and ultra-reliable communications [8]. MmWave communications will likely be an enabler for 5G and the new 3GPP standard, i.e., New Radio (NR). At mmWave frequencies, indeed, there are large chunks of contiguous bandwidth available to mobile operators, which can be used to design systems capable of reaching the 5G goals [15].

mmNets'17, October 17, 2017, Snowbird, UT, USA

ACM ISBN 978-1-4503-5143-0/17/10…$15.00
https://doi.org/10.1145/3130242.3130243

Nonetheless, harnessing the massive resources available at the mmWave frequencies in end-to-end applications delivered over cellular networks faces significant challenges. The focus of this paper is on two particularly difficult problems for mmWave cellular systems, namely support for mobility and interaction with transport protocols. Individual mmWave cells will likely have relatively small coverage areas necessitating much more frequent handovers than seen in networks today. Moreover, mmWave signals suffer from blockage from many common materials and obstacles. Overcoming these blockage events requires fast switching of the serving cells or fall-back to legacy Radio Access Technologies (RATs) [11]. In addition, the fast channel dynamics and the delays associated to handover and beamtracking can have significant effects throughout the whole protocol stack. For example, initial studies of TCP over cellular systems with mmWave links [12, 22] have shown that the rapid variations in channel quality and the frequent transitions from Line of Sight (LOS) to Non Line of Sight (NLOS) can result in both bufferbloat and slow growth of the congestion window, leading to high delays and dramatic under-utilization of the channel. Extended outages can trigger TCP retransmission timeouts and connection resets. These issues are much more serious on mmWave links than on sub-6 GHz ones (e.g., LTE), because of the much higher data rate and link quality variations considered [22].

So far, however, the performance of TCP on mmWave links in highly mobile scenarios with multiple base stations has not been studied. TCP may benefit from macro diversity in a dense deployment, with multiple mmWave base stations in a small area, and from a smart mobility management that tracks the best serving base station over time. In this paper, we analyze the performance of TCP in combination with different mobility management architectures, i.e., a standalone single connectivity solution and a dual connectivity solution originally proposed in [11]. We use the end-to-end simulator introduced in [6] with the extension described in [13] to measure different metrics such as the data rate at the application layer and the latency in the Radio Access Network (RAN) in randomly generated scenarios. We provide insights on which are the parameters and the architectural choices that most affect the TCP performance in combination with mobility procedures and dense network deployment, and show that they support high throughput and low latency.

The rest of this paper is organized as follows. In Sec. 2 we briefly describe the state of the art of the research on TCP in mmWave networks, and in Sec. 3 we introduce the mobility management schemes adopted in this paper. Then, in Sec. 4 we report the results of the performance evaluation. Finally, in Sec. 5 we draw our conclusions and outline possible future works.

2 TCP IN MMWAVE CELLULAR NETWORKS

TCP is a transport protocol introduced in the 1980s [14] in order to guarantee reliable transmission over multiple hops in the internet, as well as flow and congestion control. The two endpoints in a TCP connection establish a control loop, where the feedback is provided by acknowledgments (ACKs) sent by the receiver, and the sender uses this information to regulate the amount of data it can transmit. The ACKs are used to signal which is the last packet received correctly and in order. Moreover, with the Selective Acknowledgment (SACK) option [20], the receiver can also inform the sender about all the packets correctly received, either in order or not, so that the sender can retransmit more efficiently only the lost ones. Several congestion control algorithms can then be used to adjust the amount of unacknowledged data that can be sent (congestion window).

Since most of the applications rely on TCP in order to transmit their data on the internet, the user experience and the performance of the end-to-end connection largely depend on the interaction between TCP and all the other layers of the protocol stack. TCP on mobile networks has been extensively studied in the last decades; for example, the papers [10, 16] report a performance evaluation on the current generation of mobile networks (LTE). However, mmWave links have a much more disruptive impact on TCP than sub-6 GHz links, because of the characteristics of the mmWave frequencies (i.e., sensitivity to blockages and higher outage probability, high variability in the channel, very high data rate). In [22], a scenario with a single user and a single base station is considered, in order to study the impact on TCP of (i) extended outages, which trigger multiple retransmission timeouts and a consequent slow growth of the congestion window, and (ii) channel fluctuations (i.e., LOS/NLOS transitions). In particular, in NLOS, most of the losses on the channel are masked by link-level retransmissions, and TCP is unaware of the variations of the channel conditions. Therefore, the congestion window keeps increasing, and the packets are buffered in different layers of the protocol stack, causing a massive increase in the end-to-end latency (i.e., the so-called bufferbloat phenomenon [7]). This is further studied in [23], where Active Queue Management (AQM) schemes are applied on a mmWave channel and a cross layer approach for downlink is proposed. In [3] a cross layer congestion control algorithm for uplink flows is considered. The authors of [17] propose a modification of the TCP sender that is able to detect whether a packet is lost because of congestion or because of a mmWave disconnection, and prevents the triggering of a retransmission timeout. Finally, the paper [12] measures the throughput of TCP over mmWave links in the presence of different link-level retransmission mechanisms, and estimates the gain of path diversity at the transport layer with Multipath TCP.

The state of the art of the research on TCP over mmWave links, however, is limited to scenarios with a single base station. In this paper we determine the performance of TCP in a more realistic scenario, with multiple base stations and different mobility procedures, and compare it against that of a scenario without mobility management.

3 MOBILITY MANAGEMENT IN MMWAVE NETWORKS

Mobility management is a critical functionality for mmWave networks. The base stations at very high frequencies will serve small cells, with an ultra-dense deployment [4] and multiple mmWave access points for each sub-6 GHz station [2]. The presence of numerous base stations with overlapping coverage areas will allow a quick reaction to the shadowing of a link by a moving obstacles, however this translates into a higher burden for the network in terms of mobility-related events[1]. These may be handovers between different mmWave base stations, or to legacy networks in case all the mmWave base stations in the area are unavailable.

Several mobility management schemes for 5G networks have been proposed in the literature, either with single connectivity [21] or with multi connectivity [11, 18]. With the first approach, the mobile terminal is connected to a single base station at a time, as in legacy cellular networks (not including systems supporting carrier aggregation), and has to perform a traditional Hard Handover (HH) in order to adapt the serving base station or to switch from the mmWave RAN to a sub-6 GHz one. The second architecture, instead, allows the User Equipment (UE) to connect to multiple base stations at any given time, in order to perform coordinated transmissions [9] or to quickly switch or handover between them [18]. Multi connectivity can also be exploited in a multi RAT scenario, where the user is connected to both a sub-6 GHz base station (e.g., an LTE base station) and a mmWave one. This is the architecture proposed in [11], where the LTE base station acts as a mobility anchor for the dual-connected user and enables fast handovers between mmWave cells, without the need to report or to cooperate with the core network. Moreover, the packets are first routed to the LTE

[1] With this term we refer to all the events in which the user changes the serving base station.

base station, which processes them up to the PDCP layer and then are forwarded to the remote mmWave base stations via the X2 interface.

TCP may benefit from a timely adaptation of the serving base station, however, as shown in [10], the handover procedure may negatively impact both the TCP throughput and the end-to-end latency. Given the need for frequent handovers in mmWave networks, it is important to design mobility procedures that aim at minimizing the packet losses during handovers, as well as the interruption time caused by the switch between base stations or across different RATs. Moreover, when designing core network and backhaul architectures for mmWave cellular networks, it is important to minimize the end-to-end latency. MmWave cellular networks are indeed an example of high Bandwidth-Delay Product (BDP) networks, because of the large amount of bandwidth (and thus high data rate) available at the physical layer (in the order of Gbit/s), but the currently available TCP congestion control algorithms in high BDP scenarios offer a decreasing throughput as the end-to-end latency increases [19].

In this paper, we consider three different mobility management schemes. The baseline is a basic strategy, in which the UE connects to the mmWave base station with the highest Signal to Noise Ratio (SNR) and as it moves it does not update the serving access point (*no handover*). We look at the no handover case to demonstrate the value of dense deployments and macro diversity. Moreover, this baseline is equivalent to the deployment considered in [22]. We then consider a single connectivity approach, in which each mmWave base station is directly connected to the core network, and a HH is required to update the serving mmWave base station or fall back to the legacy LTE RAT (*Single Connectivity with HH*). The last scheme is the dual connectivity architecture proposed in [11] (*Dual Connectivity (DC)*).

4 PERFORMANCE EVALUATION

4.1 Simulator and Scenarios

For the performance evaluation we run a simulation campaign using the end-to-end mmWave module for ns–3 described in [6] with the dual connectivity extension presented in [13]. The simulated 3GPP-like protocol stack and the different nodes are presented in Fig. 1. An example of scenario is shown in Fig. 2. There are three mmWave and one LTE base stations, and N_{obs} obstacles of different size in the area between the base stations and the user. The obstacles are placed in the scenario randomly in each simulation run, in order to capture different possible propagation environments. They model buildings, trees, or other people and they force a NLOS condition when interposed between the user and a base station. Moreover, the link is considered to be in outage if the SNR is below a threshold Ω. The channel model is the one proposed by 3GPP in [1], with the spatial consistency option, so that as the user moves the channel matrix is updated in a correlated way. The user moves at speed v along the horizontal path from $x = 50$ m to $x = 150$ m, and then turns back and repeats the path multiple times, so that it is possible to measure the performance of TCP in steady-state. The application layer simulates a file transfer with full buffer.

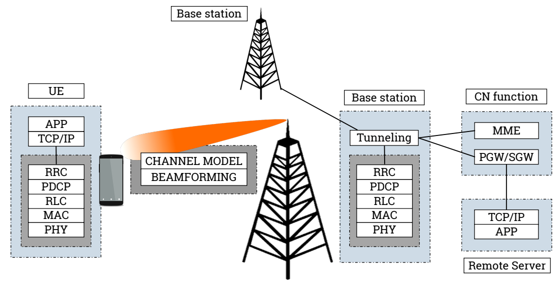

Figure 1: End-to-end protocol stack considered in the performance evaluation.

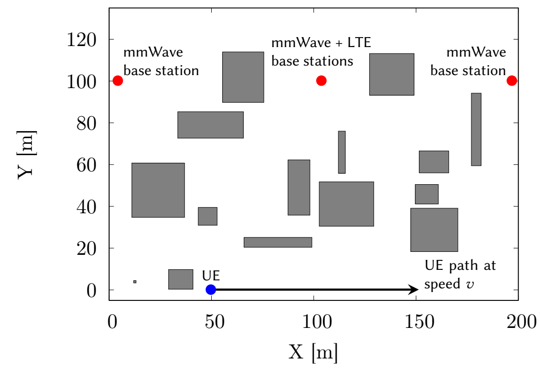

Figure 2: Example of simulation scenario. The grey rectangles are 15 randomly deployed non-overlapping obstacles (e.g., cars, buildings, people, trees).

Parameter	Value
mmWave carrier frequency	28 GHz
mmWave bandwidth	1 GHz
LTE carrier frequency (DL)	2.1 GHz
LTE bandwidth	20 MHz
3GPP Channel Scenario	Urban Micro
mmWave outage threshold Ω	−5 dB
mmWave max PHY rate	3.2 Gbit/s
X2 link latency D_{X2}	1 ms
S1 link latency D_{S1}	1 ms
PGW to remote server latency D_{RS}	[0, 10, 20] ms
RLC buffer size B_{RLC}	1 MB
RLC AM reordering timer	1 ms
S1-MME link latency D_{MME}	10 ms
UE speed v	5 m/s
Number of obstacles N_{obs}	[5, 15]
TCP Maximum Segment Size	1400 byte

Table 1: Simulation parameters

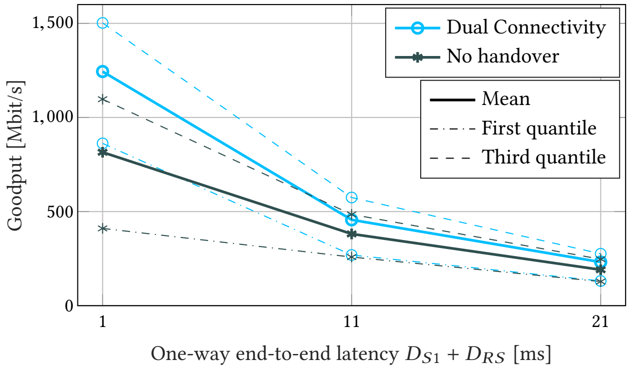

Figure 3: Goodput for a scenario with $N_{\text{obs}} = 15$ random obstacles, with and without handovers.

We compare three different values for the latency D_{RS} of the wired link between the Packet Gateway (PGW) and the application server, in both a Mobile Edge Cloud (MEC) scenario with an Edge Server (ES) (i.e., $D_{RS} = 0$ ms since the server is deployed in the core network) and a scenario with a Remote Server (RS) (i.e., $D_{RS} = 10$ or 20 ms). HighSpeed TCP is used for congestion control, since it was designed for high bandwidth-delay product networks [5]. The main parameters of the simulations are summarized in Table 1.

The metrics we consider are the application layer throughput (i.e., goodput) and the one-way RAN latency (measured from the time at which a packet enters the PDCP layer to when it is received at the UE side at the same layer). The second metric accounts also for the retransmissions at the MAC (with Hybrid Automatic Repeat reQuest (HARQ)) and the Radio Link Control (RLC) layers, and for the additional X2 latency introduced by the forwarding of packets between base stations. This happens during mobility-related events with both the single connectivity and the DC architectures, and for each packet with the latter, when the UE uses the mmWave RAT.

4.2 Goodput

In this section, we compare the performance of the different mobility management architectures. RLC Acknowledged Mode (AM) is used to perform additional retransmissions at the RLC layer, in order to minimize the packet losses over the wireless link and during handover procedures.

The first key finding is related to the gain in goodput with mobility management schemes (in particular with DC) vs. the no handover approach, as shown in Fig. 3 for different values of the fixed network delay $D_{S1} + D_{RS}$ and the scenario with $N_{\text{obs}} = 15$ obstacles. We report the average value of the goodput over multiple independent runs of the simulation, with different randomly generated scenarios (solid line), as well as the first and third quantiles (respectively, the dash-dotted and dashed lines). It can be seen that the DC option always outperforms the no handover option, with a gain up to 50% (400 Mbit/s) for the average goodput, demonstrating that access to multiple, densely deployed base stations is essential to maintain high throughputs in mmWave deployments. Individual mmWave links are highly susceptible to blocking and handover (either between mmWave cells or via fallback to LTE cells) is necessary for macro-diversity.

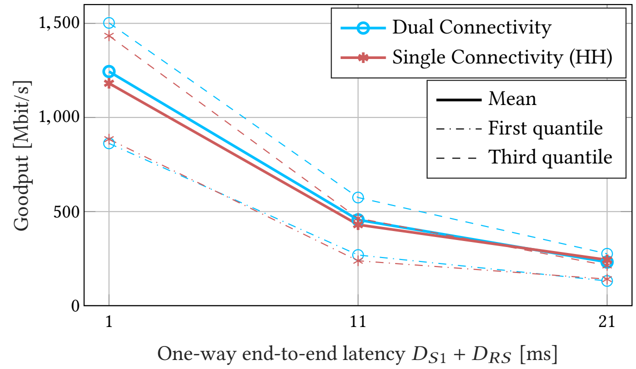

(a) $N_{\text{obs}} = 15$ **random obstacles.**

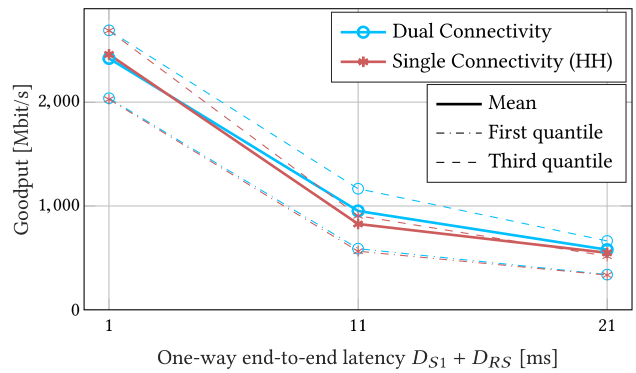

(b) $N_{\text{obs}} = 5$ **random obstacles.**

Figure 4: Goodput for a scenario with a single connectivity or dual connectivity mobility management scheme for a different number of obstacles N_{obs}.

The second observation is that the end-to-end network latency has a very significant impact on the goodput, which decreases as $D_{S1} + D_{RS}$ increases. With an ES ($D_{RS} = 0$), the goodput is limited by the physical layer data rate provided by the mmWave link, while with both RS options ($D_{RS} \geq 10$ ms) the goodput is limited by the congestion window increase rate. This observation suggests that state-of-the-art TCP mechanisms are unable to ramp up to the available throughput in mmWave channels with high variability unless network delays are very low. Core network optimization and content placement will thus likely be key to obtain the full throughput in a mmWave setting.

For a comparison between DC and single connectivity, Fig. 4a and 4b show the average goodput and the first and third quantiles for both architectures with a different number of obstacles N_{obs}. We see that DC offers a modest gain over

HH at larger network delays. While this gain is relatively small, it should be stated that the goodput here is measured on average for the entire runs where handover events are relatively infrequent. Thus, the difference in average throughput is not large. We will see in the next section that the more significant gain is in latency. In general, the dual connectivity option manages to complete the handovers between mmWave base stations or the switches across RATs in a shorter time, with fewer packet losses, therefore it sustains a generally higher goodput. However, the single connectivity solution manages to reach a better performance when there is a short interval of time with the channel in LOS condition and the user does not change the serving base station. In this case, indeed, the overall latency of the single connectivity option is smaller than that of the dual connectivity deployment[2], therefore the congestion window grows more quickly. In the scenario with $N_{\text{obs}} = 5$ and the ES, we observed that, if the same latency is considered in the fixed part of the network, then the solution with dual connectivity gains on average 400 Mbit/s (20%) with respect to the single connectivity architecture.

Finally, the number of obstacles N_{obs} plays a major role in the achievable goodput, which is up to 2 times higher with 5 obstacles than with 15. In the first case, indeed, there is a higher probability of having a LOS channel, thus a higher data rate available at the physical layer.

4.3 Latency

Fig. 5 reports the boxplots for the RAN latency of successfully received packets at the PDCP layer, for different mobility management schemes and different values of N_{obs}. It can be immediately seen that adapting the serving base station to the best one available not only increases the goodput, but also reduces the latency. The handover procedures may occasionally introduce additional latency because of the handover interruption time (i.e., the interval from the detachment from the source base station and the connection to the target one), but they are necessary to track the best serving base station and thus increase the probability of being connected with a LOS link. Therefore, the packet transmissions benefit from the higher available data rate from the lower number of HARQ and RLC retransmissions. Moreover, thanks to a dense deployment and to the handover or switch procedures, it is possible to avoid outages and most of the LOS to NLOS transitions that cause the buffering (and thus latency) at the RLC layer that was measured in [22] in combination with TCP as the transport layer, thus containing the bufferbloat issue.

[2]At least with the core network architecture considered in this paper and described in Sec. 3. It is due to the forwarding latency on the X2 link from the PDCP layer in the LTE base station to the mmWave base station.

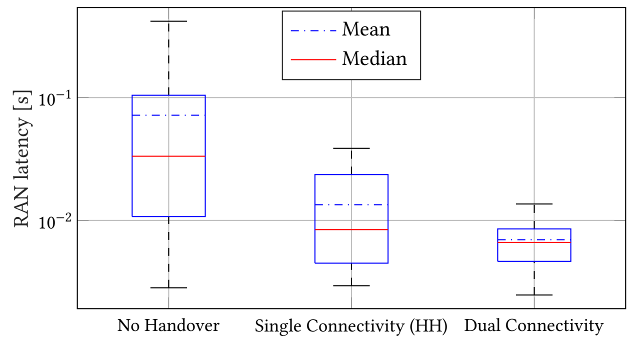

(a) $N_{\text{obs}} = 15$ random obstacles.

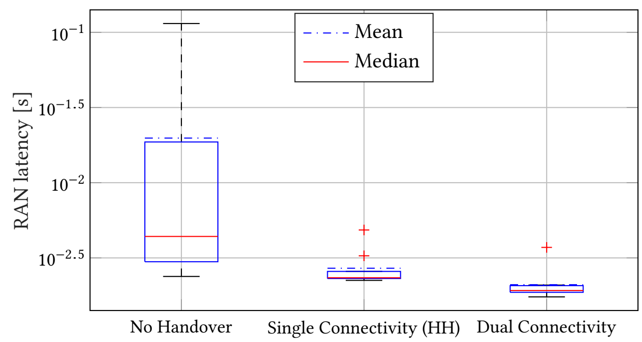

(b) $N_{\text{obs}} = 5$ random obstacles.

Figure 5: RAN one-way latency for the three different mobility management schemes, with a different number of obstacles N_{obs}. Notice that the y-axis is in logarithmic scale.

Finally, if we consider the two architectures in which the handovers are allowed, the one with dual connectivity manages to keep the latency at a minimum, and with a smaller variability as shown by the boxplots, thanks to the faster handover or RAT switch procedures [11].

4.4 RLC AM and RLC UM

In the previous sections, we considered the Acknowledged Mode of RLC, since it is usually combined with TCP, while the Unacknowledged Mode (UM) is used with best effort protocols, since it does not provide retransmissions. However, thanks to the lack of RLC layer retransmissions and the need for packet reordering at the receiver, the UM reduces the latency, and has a smaller impact on the X2 links during the handover and switch events, since with RLC AM both the transmitted but not acknowledged and the not-yet transmitted packets are forwarded from the source to the target base station, while with RLC UM only the latter are forwarded. Fig. 6 shows the goodput (solid bars) and the latency (dotted bars) for the Edge Server scenario, i.e., the one in which the TCP control loop is as short as possible. It can be seen that, as expected, RLC AM yields a higher goodput at the price of an increase in the RAN latency. Moreover, the drop in goodput of RLC UM is more noticeable with the DC architecture, since

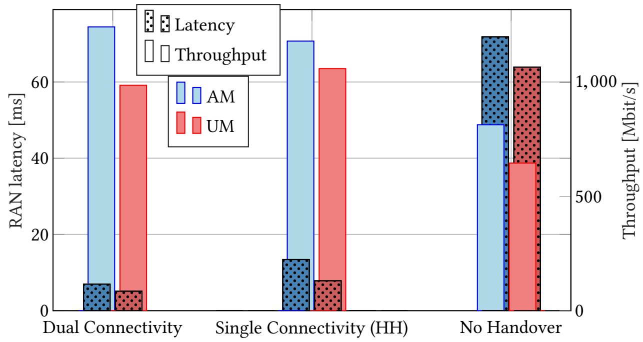

Figure 6: RAN latency and goodput for different mobility management schemes, with RLC AM and UM and $N_{obs} = 15$ obstacles.

the additional forwarding on the X2 interface decreases the reactivity of the TCP control loop. However, DC with RLC AM maintains in this scenario a slightly smaller latency than the single connectivity architecture with RLC UM, and also a larger throughput, showing that with the proper mobility management architecture it is preferable to use RLC AM with respect to UM.

5 CONCLUSIONS

In this paper, we showed the impact of mobility management architectures and macro diversity on the performance of TCP on mmWave cellular networks. The mmWave channel is intermittent because of blockages, and its performance is highly variable in LOS and NLOS. A fast adaptation of the serving base station becomes then fundamental, and TCP benefits from a mobility management scheme and a dense deployment in terms of both throughput and latency. In a dense scenario, where the probability of being connected to a LOS base station is higher, the buffering delay in the RAN and the mmWave outages have a reduced effect on the RAN latency. Moreover, the end-to-end latency has a significant impact on the goodput that TCP offers to the application layer, because of the unresponsiveness of the TCP control loop for high BDP networks.

As part of our future work we will extend our study to larger simulation scenarios and to a real testbed. Moreover, we will both design network architectures that can reduce the end-to-end latency of a connection, and improve the responsiveness of TCP by adopting a performance enhancing proxy in the core network.

REFERENCES

[1] 3GPP. 2017. TR 38.900, Study on channel model for frequency spectrum above 6 GHz, V14.2.0. (2017).

[2] 3GPP. 2017. TR 38.913, Study on Scenarios and Requirements for Next Generation Access Technologies, V14.1.0. (2017).

[3] T. Azzino, M. Drago, M. Polese, A. Zanella, and M. Zorzi. 2017. X-TCP: A Cross Layer Approach for TCP Uplink Flows in mmWave Networks. In *16th Annual Mediterranean Ad Hoc Networking Workshop (Med-Hoc-Net'17)*.

[4] N. Bhushan, J. Li, D. Malladi, R. Gilmore, D. Brenner, A. Damnjanovic, R. T. Sukhavasi, C. Patel, and S. Geirhofer. 2014. Network densification: the dominant theme for wireless evolution into 5G. *IEEE Communications Magazine* 52, 2 (Feb. 2014), 82–89.

[5] S. Floyd. 2003. *HighSpeed TCP for large congestion windows.* RFC 3649.

[6] R. Ford, M. Zhang, S. Dutta, M. Mezzavilla, S. Rangan, and M. Zorzi. 2016. A framework for end-to-end evaluation of 5G mmwave cellular networks in ns-3. In *Proceedings of the Workshop on ns-3*. ACM, 85–92.

[7] J. Gettys and K. Nichols. 2011. Bufferbloat: Dark Buffers in the Internet. *ACM Queue* 9, 11, Article 40 (Nov. 2011), 15 pages.

[8] M. Iwamura. 2015. NGMN View on 5G Architecture. In *2015 IEEE 81st Vehicular Technology Conference (VTC Spring)*.

[9] D. Maamari, N. Devroye, and D. Tuninetti. 2016. Coverage in mmWave Cellular Networks With Base Station Co-Operation. *IEEE Transactions on Wireless Communications* 15, 4 (Apr. 2016), 2981–2994.

[10] B. Nguyen, A. Banerjee, V. Gopalakrishnan, S. Kasera, S. Lee, A. Shaikh, and J. Van der Merwe. 2014. Towards Understanding TCP Performance on LTE/EPC Mobile Networks. In *Proceedings of the 4th Workshop on All Things Cellular: Operations, Applications, and Challenges*. ACM, New York, NY, USA, 41–46.

[11] M. Polese, M. Giordani, M. Mezzavilla, S. Rangan, and M. Zorzi. 2017. Improved Handover Through Dual Connectivity in 5G mmWave Mobile Networks. *IEEE JSAC Special Issue on Millimeter Wave Communications for Future Mobile Networks* 35, 9 (September 2017).

[12] M. Polese, R. Jana, and M. Zorzi. 2017. TCP in 5G mmWave Networks: Link Level Retransmissions and MP-TCP. In *2017 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE.

[13] M. Polese, M. Mezzavilla, and M. Zorzi. 2016. Performance Comparison of Dual Connectivity and Hard Handover for LTE-5G Tight Integration. In *Proceedings of the 9th EAI International Conference on Simulation Tools and Techniques (SIMUTOOLS'16)*. 118–123.

[14] J. Postel. 1981. Transmission Control Protocol. RFC 793. (Sept. 1981). Available on-line at `http://www.faqs.org/rfcs/rfc793.html`.

[15] S. Rangan, T. S. Rappaport, and E. Erkip. 2014. Millimeter-Wave Cellular Wireless Networks: Potentials and Challenges. *Proc. IEEE* 102, 3 (Mar. 2014), 366–385.

[16] Remi Robert, Eneko Atxutegi, Ake Arvidsson, Fidel Liberal, Anna Brunstrom, and Karl-Johan Grinnemo. 2016. Behaviour of common TCP variants over LTE. In *Global Communications Conference (GLOBECOM), 2016 IEEE*. IEEE, 1–7.

[17] H. Sato and M. Bandai. 2017. A Transport Protocol for Millimeter-Wave Links. *IEICE Communications Express* (May 2017).

[18] F. B. Tesema, A. Awada, I. Viering, M. Simsek, and G. P. Fettweis. Mobility Modeling and Performance Evaluation of Multi-Connectivity in 5G Intra-Frequency Networks. In *2015 IEEE Globecom Workshops*.

[19] G. Vardoyan, N. S. V. Rao, and D. Towsley. 2016. Models of TCP in high-BDP environments and their experimental validation. In *2016 IEEE 24th International Conference on Network Protocols (ICNP)*.

[20] L. Wang, M. Allman, E. Blanton, and K. Fall. 2012. *A Conservative Loss Recovery Algorithm Based on Selective Acknowledgment (SACK) for TCP.* RFC 6675.

[21] V. Yazici, U. C. Kozat, and M. O. Sunay. 2014. A new control plane for 5G network architecture with a case study on unified handoff, mobility, and routing management. *IEEE Communications Magazine* 52, 11 (Nov. 2014), 76–85.

[22] M. Zhang, M. Mezzavilla, R. Ford, S. Rangan, S. S. Panwar, E. Mellios, D. Kong, A. R. Nix, and M. Zorzi. 2016. Transport Layer Performance in 5G mmWave Cellular. In *2016 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*.

[23] M. Zhang, M. Mezzavilla, J. Zhu, S. Rangan, and S. S. Panwar. 2016. The Bufferbloat Problem over Intermittent Multi-Gbps mmWave Links. arXiv:1611.02117 [cs.NI]. (Nov. 2016).