TCP and MP-TCP in 5G mmWave networks

Michele Polese*, Rittwik Jana^, Michele Zorzi*
*Department of Information Engineering, University of Padova, Italy
Email: {polesemi, zorzi}@dei.unipd.it
^AT&T Labs-Research, Bedminster NJ, USA
Email: rjana@research.att.com

Abstract

Future 5G networks will likely include mmWave radio access communication links, because of their potential multi-gigabit-per-second capacity. However, these frequencies are characterized by very dynamic channel conditions which lead to wide fluctuations in the received signal quality. This article explains how the end-to-end user experience in mobile mmWave networks could be affected by a sub-optimal interaction between the most widely used transport protocol, TCP, and mmWave links. It also provides insights on the throughput-latency trade-off when Multipath TCP (MP-TCP) is used judiciously across various links (e.g., LTE and mmWave).

Keywords

5G, TCP, Multipath TCP, mmWave, congestion control.

I. INTRODUCTION

The next generation of mobile networks (5G) will be standardized before 2020 in order to address the growth of mobile internet traffic. According to the Next Generation Mobile Network (NGMN) alliance, 5G networks shall provide [1] (i) a user bitrate of at least 50 Mbit/s at cell edges, with Gbit/s peaks in the most favorable conditions; (ii) ultra-low end-to-end latency (i.e., possibly below 10 ms); (iii) very high reliability and availability of communications; and (iv) support for low-power massive Machine Type Communications (MTC).

MmWave communications will play a major role in meeting the throughput target of 5G networks. At frequencies above 10 GHz, indeed, there is a high availability of contiguous spectrum that can be allocated to cellular networks. However, mmWaves present numerous challenges and issues that must be addressed in order to make this technology market-ready. In fact, these frequencies suffer from high isotropic pathloss and blockage by most solid materials, such as for example buildings, cars, and also the human body [2], that may result in service unavailability (outage). The high antenna gain provided by using multiple antenna techniques (i.e., beamforming and massive MIMO) can make up for the pathloss, but how to provide a reliable service in the presence of frequent blockages is still an open question. In particular, a sudden transition of a mmWave link from a Line-of-Sight (LOS) to a Non-Line-of-Sight (NLOS) channel state generates wide fluctuations in the SINR (in the order of 30 dB, according to [2]), and therefore the capacity offered changes drastically, thereby compromising the end-to-end user experience.

These extreme propagation conditions demand a new design of the physical and Medium Access Control (MAC) layers, but also have an impact on the interplay with the higher layers of the protocol stack. The most widely used reliable data transport protocol is TCP, which was however designed in a different historical context, and according to different needs and constraints. TCP considers packet losses as an implicit notification of congestion on the link, and therefore reduces the sending rate trying to relieve it. However, the lossy nature of mmWave links may trigger the TCP congestion control mechanisms even if there is no actual congestion. This may yield a sub-optimal end-to-end performance and waste the great potential of mmWave links.

In this article, we provide an overview of the impact of mmWave communications on the performance of TCP, outlining which are the main novelties that the usage of a mmWave Radio Access Network (RAN) introduces in end-to-end connections at the transport layer, and how the adoption of the recent Multipath TCP (MP-TCP) option can improve the performance of TCP on these kinds of links. Finally, we suggest possible future research directions. A more technical discussion with additional results can be found in [3].

II. RECENT ADVANCES IN TCP

TCP was designed in the 1980s as a connection-oriented and reliable protocol that provides end-to-end connectivity over multiple hops and congestion control (CC). Reliability is enabled by a retransmission mechanism, based on the acknowledgments received by the TCP transmitter. Congestion control is implemented by different algorithms that increase and/or decrease the maximum amount of unacknowledged data that the sender is allowed to transmit (*congestion window*), reacting to network events such as packet losses. There have been several evolutions of the original congestion control algorithms: the latest Request For Comments (RFC) describing them is [4], and the survey in [5] lists 13 TCP variants implemented in the Linux kernel. In particular, the latest versions of the major Operating Systems use TCP CUBIC as the default [6]. This TCP flavor was designed in order to react more promptly to packet losses, and to restore the connection to the highest available rate in a shorter time than legacy designs based on TCP New Reno. In the latter, after a packet loss, the sender decreases its congestion window by half, and increases it by one packet in each RTT, thus requiring a very long time to fill a high capacity link. With TCP CUBIC, instead, the growth is independent of the RTT, which makes it possible to reach a higher throughput more quickly than in New Reno, while behaving fairly towards flows using other versions of TCP [6].

Another trend in modern cellular networking is to exploit the presence of multiple network interfaces in mobile devices. A typical example is the usage of MP-TCP for vertical and seamless handovers between 3G/4G cellular networks and Wi-Fi hotspots [7]. MP-TCP is an extension of TCP that enables multipath transport, i.e., the application communicates with a traditional TCP socket, which transparently handles multiple subflows on different interfaces, such as, for example, Wi-Fi, the cellular network, and Ethernet. MP-TCP is currently under discussion in the IETF [8], [9], and is designed around three main goals [10]. First, it should improve the throughput, in the sense that it should perform at least as well as a single path TCP (SP-TCP) flow on the best path available. Second, it should not use more resources than standard TCP flows. Finally, it should steer more packets towards less congested paths.

The core of MP-TCP is its congestion control algorithm, that manages the congestion window of each different subflow in a coupled or uncoupled manner. With an uncoupled congestion control, each subflow is independent, i.e., the congestion window is updated separately in each path. With a coupled approach, instead, the congestion window of the different subflows is increased and decreased in a coordinated manner, considering the congestion of all the available subflows. By coupling the different subflows, the authors of [10] claim that it is possible to reach the above MP-TCP design goals. The first coupled CC proposed in [10] is however criticized in [11] and in [12], because it (i) transmits too much traffic on congested paths and (ii) is unfriendly with respect to SP-TCP. The authors of [11] present the Opportunistic Linked Increases Algorithm (OLIA), which is designed to overcome these two issues, but according to [12] presents non-responsiveness problems with respect to congestion changes in the subflows. The Balanced Linked Adaptation algorithm (BALIA) [12] addresses both the problems of the original CC and those of OLIA. In particular, the parameters of the protocol are derived through a theoretical analysis of the performance of multipath congestion control algorithms. However, these schemes are based on the legacy design of Reno and New Reno congestion control algorithms, and thus share with them the drawbacks described in the previous paragraph.

In this article, we discuss the main issues that arise when TCP and MP-TCP are deployed in networks with mmWave links, showing how state-of-the-art transport protocols interact with this technology.

III. TCP PERFORMANCE ON MMWAVE NETWORKS

A. The mmWave Channel

Since TCP is the most widely used transport protocol, it is important to understand the interactions that exist between mobile networks (i.e., wireless channels) and the TCP performance. In wireless networks the loss of a packet is not necessarily caused by congestion, but may instead be due to a sudden (and possibly only temporary) drop in signal quality. In [14], [15], the authors study the behavior of TCP in relation to a complex mobile network such as LTE, showing (i) that as the distance between the User Equipment (UE) and the evolved Node Base (eNB) increases the TCP throughput degrades and (ii) how TCP is affected by network events such as handovers.

MmWave networks are expected to reach an order of magnitude higher throughput than current systems, thanks to the larger bandwidth available, but present more troublesome propagation conditions. An example is shown in Fig. 1, which compares the time evolution of the SINRs of a mmWave link at 28 GHz and an LTE link at 2.1 GHz, for a UE that moves at 2 m/s at an average distance of 75 meters from the eNB and switches from a LOS to a NLOS condition. The main differences between the mmWave and the LTE channels are that:

- the LOS to NLOS pathloss transitions are deeper for mmWave. At sub 6 GHz frequencies, these are of the order of 10-15 dB, while for the mmWave channel the fluctuation can exceed 30 dB. Therefore, the available capacity changes dramatically. Moreover, mmWave networks will be small cell networks, and mmWave links are sensitive to blockage from foliage, the human body, moving obstacles and so on, thus the transitions from LOS to NLOS for mmWave connections will be much more frequent than it is in LTE [2]. Furthermore, given the shorter range of mmWave communication, it is more probable for mmWave links than for LTE ones to experience an outage (i.e., no signal is received) because of shadowing;
- the mmWave channel has a shorter coherence time, which results in variations of the channel of the order of hundreds of microseconds [2], that are faster than those in current mobile networks. As shown in Fig. 1, the transitions of the LTE channel due to fading are much smoother than those of the mmWave link.

Currently, these issues are mainly addressed in the Radio Access Network, but they have an impact also on the performance of transport layer protocols. To date, there are only few published results on the performance of TCP over 5G mmWave links. This is due to several reasons. The first is the lack of mobile testbeds that scale beyond a single mmWave link. There are indeed commercial prototypes that can be used to set up a mmWave link between a MIMO transmitter and receiver, but the focus of these products is on the single connection. The deployment of large testbeds is limited by the lack of a definitive standard and by the cost and complexity of prototyping the electronics for these links. Moreover, an end-to-end system composed of a mmWave RAN, a core network (CN), and the internet with remote servers has a very high complexity and this prevents an insightful theoretical analysis of the system. Finally, an analytical channel model that represents the correlation of fading over time in mmWave links has not yet been developed.

As mmWave networks and 5G are not yet standardized, a good compromise between the realism of a testbed and the convenience of a theoretical analysis is system-level simulation, that can provide deep insights on a complex system while avoiding the need for expensive dedicated hardware. In this article we use an extension presented in [16] of the NYU mmWave [17] module of ns--3, a widely used open-source network simulator, together with the TCP/IP stack of the Linux kernel in order to provide a realistic insight on the interplay between TCP and mmWave links.

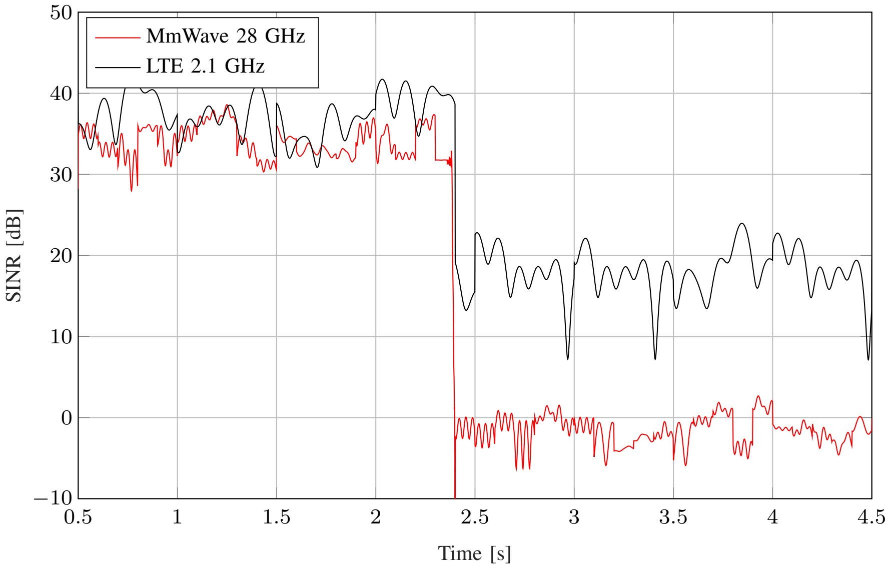

Fig. 1: *ns--3 simulated SINR for a mmWave and LTE link, with the UE moving at 2 m/s from (45, 0) to (55,0), with the eNB at coordinates (75, 50). The traces are generated using the channel model described in [2], [13] for the mmWave channel and the ns—3.*

B. Link-Level Retransmissions and TCP

One of the main issues of mmWave links is the highly dynamic channel, that, because of blockages, experiences wide fluctuations of the perceived SINR and is therefore characterized by unreliable transmissions. Current and future mobile networks deploy different retransmission mechanisms in order to mitigate the effects of packet loss and increase the throughput of the mobile devices. When using mmWave links, these retransmission protocols become a key element in hiding the highly dynamic and consequently unstable behavior of the channel to higher layer transport protocols such as TCP.

At the MAC layer, Hybrid Automatic Repeat reQuest (HARQ) is used. When the physical layer at the receiver receives a packet, but detects the presence of some errors that prevent reliable decoding, it asks for a retransmission. The sender then transmits additional redundancy that helps retrieve the packet without any error [18]. Morever, in 3GPP-like networks (e.g., LTE), there is a layer on top of the MAC layer that may perform additional link-level retransmissions, the Radio Link Control (RLC) layer, that will also likely be a part of the 5G protocol stack [19]. Since the number of retransmissions at the MAC layer is usually limited (typically only 3 attempts are performed), the RLC layer Acknowledged Mode (AM) offers another way of recovering lost packets.

Thanks to periodic acknowledgments from the receiver, the RLC AM sender knows which packets were lost and can retransmit them. The number of attempts that RLC AM can perform is also limited. RLC Unacknowledged Mode (UM) instead does not perform any retransmission in addition to those of the HARQ at the MAC layer. Both HARQ and RLC AM operate based on information related to the link and with a greater timeliness with respect to TCP, which uses packet losses to detect congestion and operates on the larger timescale of retransmission time-outs (RTOs), of the order of a second.

A first study of TCP in mmWave cellular networks is presented in [20]. The authors study different scenarios, with either a statistical channel model, based on real measurements in New York City, or a raytracing model, with traces from a route in Bristol. They observe that, in the presence of long outage events, neither HARQ nor RLC AM retransmissions are able to mask the channel losses to the TCP sender, which triggers a RTO event, after which it may take a long time to recover the full capacity. Moreover, because of the presence of these retransmission mechanisms, TCP does not promptly track the mmWave link state when the channel between the UE and the eNB switches from a LOS to a NLOS condition. This issue is more marked in mmWave than in current cellular networks, because LOS to NLOS transitions are much more frequent and provide wider variations in the link capacity [2]. The consequence is that the packets queue in the RAN buffers, and latency increases. This is an instance of the well-known bufferbloat problem [21]: because of the large size of buffers in network devices, packets are not dropped but queued and eventually transmitted. While this improves the throughput of the mmWave link, it also increases the overall latency.

This trade-off can be better understood by considering the results in Fig. 2a, where we use the iPerf network bandwidth measurement tool on top of TCP CUBIC with and without link-level retransmissions, at different distance d between the UE and the eNB. At a small distance (i.e., $d =$ 50 m) the achievable TCP throughput is high even without retransmissions, because the UE is in LOS with very high probability. As the distance increases, two different behaviors can be observed, if retransmissions are used: (i) the achievable throughput is from 1.5 to 3 times higher than without retransmissions; however (ii) the latency increases, especially when MAC and RLC layer retransmissions are combined at very large distances.

If we consider the download of small files (from 1 to 10 MB) using wget, which can be representative of web browsing, then the difference between the download times with RLC AM (both RLC retransmissions and HARQ) and with RLC UM (HARQ only) is more noticeable than that between the values of Fig. 2b, showing that for short-lived TCP sessions it is important to perform retransmissions as fast as possible, i.e., at a layer which is as close to the radio link as possible.

The design goals of 5G networks, nonetheless, demand both high throughput and low latency. However, as shown in Fig. 2a and in [20], when the distance between eNB and UE is high or there are blockages and a NLOS channel, a high TCP throughput can be reached only with retransmissions, thus increasing latency. How to solve the latency-throughput trade-off is an issue that must be solved before mmWave networks can be considered market-ready. A first step, as proposed in [22], [23], is a cross-layer approach that makes TCP aware of the actual channel conditions. A second solution may involve the replacement of TCP with other transport protocols, or the re-design of HARQ and RLC-layer retransmissions in the 3GPP stack. Finally, an ultra-dense deployment could decrease the probability of having NLOS links towards mmWave eNBs, however, in this case the issue becomes a robust management of mobility, handovers and cell association [13].

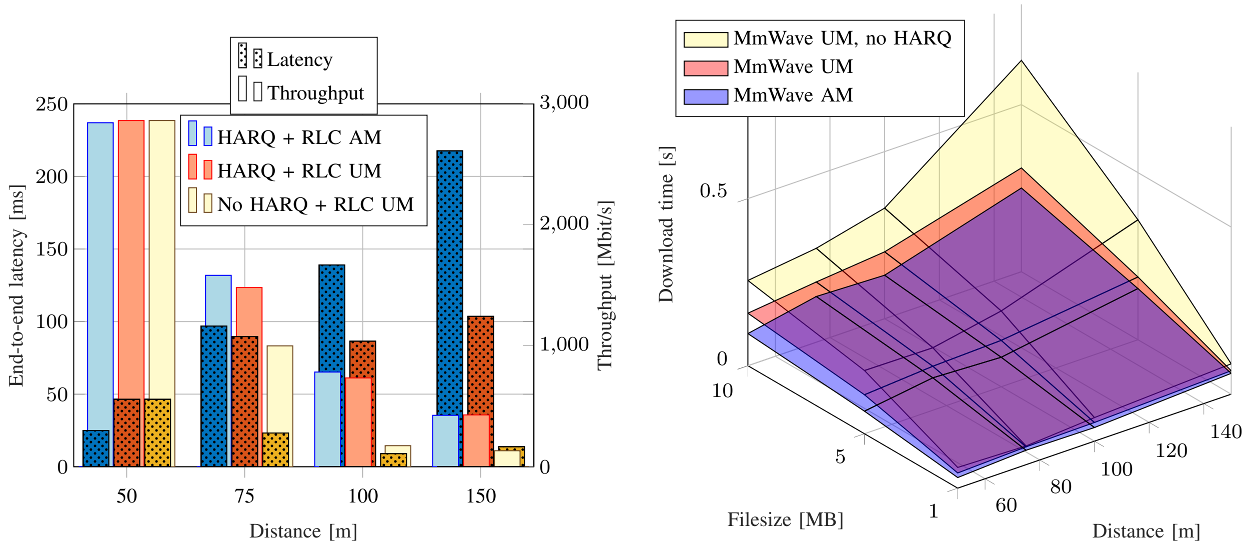

Fig. 2a: TCP latency (dotted bars) and throughput, with and without the different retransmission mechanisms of the mmWave protocol stack.

Fig. 2b: Download time as a function of the file size and of the distance, for TCP with and without lower-layer retransmissions.

Fig. 2: TCP performance on mmWave links, with and without link-level retransmissions.

C. Multipath TCP in mmWave Networks

5G devices will likely be connected to multiple Radio Access Technologies (RATs) or to multiple mmWave eNBs [13]. This is why it is interesting to analyze the performance of MP-TCP over mmWave and LTE links, since it could be used as an end-to-end solution for multi-connectivity. However, the complex interactions that exist between the reliability of the different paths and the different congestion control algorithms require a deep understanding of MP-TCP performance before deploying it on mmWave links.

Firstly, we consider whether using LTE or mmWave as a secondary subflow yields a higher throughput. In Fig. 3 we show the throughput performance of different MP-TCP congestion control algorithms over different connections, with respect to the baseline of a SP-TCP connection with TCP CUBIC on a mmWave link. When the UE has a high probability of being in LOS (i.e., $d <= 50$ m), the solution with Multipath TCP on mmWave-only links outperforms SP-TCP, with a gain between 800 Mbit/s and 1 Gbit/s (about 30-40%). The LTE link, instead, has a much smaller rate than a mmWave link when $d <= 50$ m and, therefore, the throughput of MP-TCP on LTE and mmWave subflows is close to or worse than that of the reference SP-TCP. However, the performance of an LTE link is less dependent on the distance d than that of a mmWave link, thus, for higher distances, MP-TCP over LTE and mmWave links performs better than MP-TCP with only mmWave connections. The 73 GHz mmWave subflow, indeed, provides a larger theoretical capacity than the LTE one, but has a lossy behavior that penalizes the overall throughput, except for small distances. As shown in Fig. 3, for $d = 150$ m, MP-TCP with LTE and 28 GHz mmWave offers a gain of more than 450 Mbit/s (about 100%) with respect to SP-TCP (i.e., more than the LTE link throughput), without increasing the latency. This shows that link diversity, i.e., the presence of the secondary and reliable LTE path, improves the throughput of the overall connection.

Secondly, we compare the performance of two different congestion control algorithms for MP-TCP: BALIA, i.e., a coupled CC algorithm, and CUBIC, i.e., an uncoupled one. As shown in Fig. 3, MP-TCP with the coupled BALIA CC algorithm fails to meet the first target of MP-TCP design, since in many cases its throughput is lower than that of SP-TCP. When coupling a mmWave and an LTE link, indeed, the congestion control algorithm sees the losses on the 28 GHz mmWave connection as congestion, and it steers the whole traffic to the LTE subflow, degrading the performance of the end-to-end connection. Instead, the uncoupled congestion control algorithm is not affected by this issue, since each path behaves independently. However, in this case the MP-TCP flow may be unfriendly to SP-TCP flows on shared bottlenecks.

Therefore, the definition of a congestion control algorithm that meets the three design goals of MP-TCP for mmWave links is still an open area of research. The current CC algorithms either waste the potential gain given by multi-connectivity, because of sub-optimal balancing between flows, or harm the fairness among network users. Therefore, before deploying these algorithms in future mmWave networks, a proper congestion control algorithm should be designed and thoroughly tested.

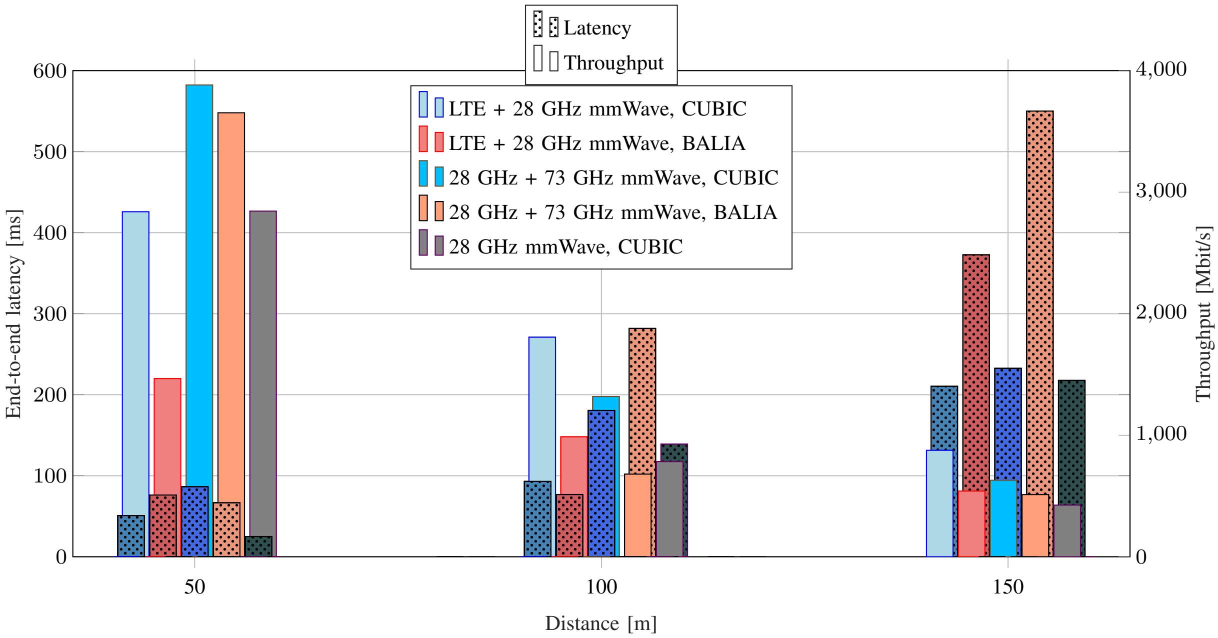

Fig. 3: MP-TCP throughput and latency for different distances d and different MP-TCP options. The rightmost bar in each group shows the performance of a SP-TCP connection with TCP CUBIC over a 28 GHz mmWave link as a reference.

IV. CONCLUSIONS

In this article, we provided an overview of the issues that may arise from the complex interactions between TCP, either single or multipath, and mmWave links. The interplay between communications at such high frequencies and legacy transport protocols could be a limitation that prevents the full exploitation of the potential of mmWave communications. For mmWave it is fundamental to mask the channel losses to the higher TCP layer with link-level retransmissions, in order to reach high throughput, but these mechanisms introduce additional latency. An important contribution could be given by multi-connectivity through MP-TCP at the transport layer, where path diversity offers promising results, but a suitable congestion control algorithm has not yet been identified.

An optimized interaction between transport protocols and mmWave networks will be a fundamental enabler of 5G networks. In particular, the main issue is related to how to reduce latency while maintaining a high and stable throughput while using TCP and MP-TCP over mmWave networks, so that the end users can effectively exploit the huge amount of resources available in fifth generation cellular networks.

In order to reach this goal, in the next few years several open research questions will have to be addressed, such as: (i) how to design reliable handover mechanisms that do not impact TCP performance in high mobility scenarios; (ii) how to optimally exploit multi-connectivity, either with MP-TCP, or with different solutions at a lower layer in the RAN, also studying the trade-off between energy consumption and very high throughput; and (iii) how to combine the stability and reactiveness of TCP CUBIC with a MP-TCP congestion control algorithm coupled over 5G mmWave and legacy links.

REFERENCES

[1] M. Iwamura, “NGMN View on 5G Architecture,” in 2015 IEEE 81st Vehicular Technology Conference (VTC Spring), May 2015, pp. 1–5.

[2] S. Rangan, T. S. Rappaport, and E. Erkip, “Millimeter-wave cellular wireless networks: Potentials and challenges,” Proc. IEEE, vol. 102, no. 3, pp. 366–385, Mar 2014.

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

[4] M. Allman, V. Paxson, and E. Blanton, “TCP congestion control,” RFC 5681, 2009.

[5] C. Callegari, S. Giordano, M. Pagano, and T.Pepe, “A survey of congestion control mechanisms in Linux TCP,” in Distributed Computer and Communication Networks. Springer, Mar 2014, pp. 28–42.

[6] S. Ha, I. Rhee, and L. Xu, “CUBIC: a new TCP-friendly high-speed TCP variant,” ACM SIGOPS Operating Systems Review, vol. 42, no. 5, pp. 64–74, July 2008.

[7] Y.-C. Chen, Y.-s. Lim, R. J. Gibbens, E. M. Nahum, R. Khalili, and D. Towsley, “A measurement-based study of Multipath TCP performance over wireless networks,” in Proceedings of the 2013 Conference on Internet measurement. ACM, 2013, pp. 455–468.

[8] A. Ford, C. Raiciu, M. Handley, and O. Bonaventure, “TCP Extensions for Multipath Operation with Multiple Addresses,” RFC 6824, 2013.

[9] A. Ford, C. Raiciu, M. Handley, S. Barre, and J. Iyengar, “Architectural Guidelines for Multipath TCP Development,” RFC 6182, 2011.

[10] C. Raiciu, M. Handley, and D. Wischik, “Coupled congestion control for multipath transport protocols,” RFC 6356, 2011.

[11] R. Khalili, N. Gast, M. Popovic, and J. Y. L. Boudec, “MPTCP is Not Pareto-Optimal: Performance Issues and a Possible Solution,” IEEE/ACM Transactions on Networking, vol. 21, no. 5, pp. 1651–1665, Oct 2013.

[12] Q. Peng, A. Walid, J. Hwang, and S. H. Low, “Multipath TCP: Analysis, Design, and Implementation,” IEEE/ACM Transactions on Networking, vol. 24, no. 1, pp. 596–609, Feb 2016.

[13] M. Polese, M. Giordani, M. Mezzavilla, S. Rangan, and M. Zorzi, “Improved Handover Through Dual Connectivity in 5G mmWave Mobile Networks,” to appear in IEEE JSAC Special Issue on Millimeter Wave Communications for Future Mobile Networks, 2017 [Online]. Available: https://arxiv.org/abs/1611.04748

[14] B. Nguyen, A. Banerjee, V. Gopalakrishnan, S. Kasera, S. Lee, A. Shaikh, and J. Van der Merwe, “Towards understanding TCP performance on LTE/EPC mobile networks,” in Proceedings of the 4th workshop on All things cellular: operations, applications, & challenges. ACM, 2014, pp. 41–46.

[15] R. Li, M. Shariat, and M. Nekovee, “Transport protocols behaviour study in evolving mobile networks,” in 2016 International Symposium on Wireless Communication Systems (ISWCS), Sep 2016, pp. 456–460.

[16] M. Polese, M. Mezzavilla, and M. Zorzi, “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, 2016, pp. 118–123.

[17] R. Ford, M. Zhang, S. Dutta, M. Mezzavilla, S. Rangan, and M. Zorzi, “A Framework for End-to-End Evaluation of 5G mmWave Cellular Networks in ns-3,” in Proceedings of the Workshop on Ns-3, June 2016, pp. 85–92.

[18] S. Choi and K. G. Shin, “A class of adaptive hybrid ARQ schemes for wireless links,” IEEE Transactions on Vehicular Technology, vol. 50, no. 3, pp. 777–790, May 2001.

[19] Korea Telecom (KT) 5G-SIG, “5G.322, KT 5th Generation Radio Access, Radio Link Control protocol (5G pre-specifications),” 2016.

[20] M. Zhang, M. Mezzavilla, R. Ford, S. Rangan, S. S. Panwar, E. Mellios, D. Kong, A. R. Nix, and M. Zorzi, “Transport Layer Performance in 5G mmWave Cellular,” in 2016 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS), Apr 2016, pp. 730–735.

[21] J. Gettys and K. Nichols, “Bufferbloat: Dark buffers in the internet,” ACM Queue, vol. 9, no. 11, pp. 40:40–40:54, Nov 2011.

[22] M. Zhang, M. Mezzavilla, J. Zhu, S. Rangan, and S. Panwar, “The Bufferbloat Problem over Intermittent Multi-Gbps mmWave Links,” Nov. 2016. [Online]. Available: https://arxiv.org/abs/1611.02117

[23] T. Azzino, M. Drago, M. Polese, A. Zanella, M. Zorzi, “X-TCP: A Cross Layer Approach for TCP Uplink Flows in

mmWave Networks," in Proceedings of the 16th Annual Mediterranean Ad Hoc networking workshop, June 2017. [Online]. Available: https://arxiv.org/abs/1706.00904

Michele Polese received his B.Sc. (2014) and M.Sc. (2016) in Telecommunication Engineering from the University of Padova, Italy. Since October 2016 he is a Ph.D. student at the Department of Information Engineering of the University of Padova. In April 2017 he spent a month as Visiting Academic at NYU WIRELESS. His research interests are in the analysis and development of protocols and architectures for the 5G mmWave networks.

Rittwik Jana is a Director Inventive Science at AT&T Labs Research. His research interests span intelligent service composition of VNFs using TOSCA, model driven control loop and automation in ONAP, networked video streaming and wireless communications and systems. Rittwik obtained a Ph.D. in Telecommunications Engineering from the Australian National University, Canberra, Australia in 2000.

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.