On the Benefits of Multi-Domain Congestion Control in LTE Networks

Zsolt Krämer*, Sándor Molnár*, Szilárd Solymos*, Attila Mihály†
*Dept. of Telecomm. and Media Informatics, Budapest University of Technology and Economics, Hungary
{kramer, molnar, solymos}@tmit.bme.hu
†Traffic Analysis and Network Performance Laboratory, Ericsson Research, Hungary
attila.mihaly@ericsson.com

***Abstract*—Wireless networks have traditionally been a challenging environment for congestion control algorithms. Our Multi-Domain Congestion Control framework aims to outperform congestion control solutions that are based on client feedback only by leveraging cooperation between the network and the servers. In this paper, we present a performance evaluation of our approach in realistic LTE simulations. We show that Multi-Domain Congestion Control delivers significant performance gains over TCP CUBIC in short flow completion time achieving 33-84% reduction in the 1-10MB flow size range. We also demonstrate that it is able to swiftly react to increased capacity in various typical LTE environments, utilizing the new bandwidth more than 76% faster than CUBIC in a vehicular - and more than 71% faster in pedestrian scenario. Furthermore, it is demonstrated that our cooperative framework also outperforms CUBIC in long term throughput.**

***Index Terms*—Congestion Control, Cellular Network, LTE, Performance.**

I. Introduction

In Q4 2018, the number of LTE subscriptions increased by approximately 240 million, and have reached a total of around 3.6 billion. Monthly mobile data traffic grew close to 88% in 2018 which is the largest growth observed since 2013 [1]. As the key technological enablers for 5G use cases continue to evolve, it is also important to evaluate the performance of these proposals in current LTE networks.

Internet congestion control (CC) research gained momentum recently both in industry and academia. Notable proposals include the BBR (Bottleneck Bandwidth and RTT) [2] algorithm developed by Google and the PCC (Performance-oriented Congestion Control) approach described in [3] by Dong et al. Both algorithms showed promising results, however, mobile networks provide unique challenges for end-to-end congestion control and thus, the effectiveness of end-to-end solutions are limited in such environments. Moreover, [4] showed that BBR is not able to provide fairness with CUBIC.

In cellular networks, the deployment of split-connection TCP proxies appears to be a straightforward way to improve performance, mainly due to the reduced RTT between the connected endpoints and the isolation of loss events in the resulting two TCP control loops. X. Xu et al. showed the prevalence of these split-connection solutions in all four major US carrier networks, however, only in two of them were the performance gains visible [5]. It is also important to note that these performance enhancing proxies (PEPs) are expensive to maintain and contribute heavily to the ossification of the transport layer [6].

In [7] we have presented the concept of a non-ossifying, lightweight performance enhancing proxy and in [8] we describe the design, implementation and initial performance evaluation results of a Multi-Domain CC algorithm that utilizes the information sent by the Lightweight PEP.

In this paper, as a continuation of this research, we study the performance of our cooperative framework in realistic LTE simulations involving different fading channels (vehicular, pedestrian) and three key scenarios: sudden capacity increase, short flow downloads, and long flows with volatile background traffic.

The rest of the paper is organized as follows. Section II presents the related work, with Section II-C providing an overview of the components and design of our cooperative framework. Section III describes the simulation environment. Section IV and V present the performance evaluation results for the capacity increase - and the short flow scenarios, respectively. Section VI demonstrates the long term performance benefits in the presence of mixed background traffic. Section VII concludes the paper.

II. Related work

A. TCP performance in cellular networks

R. Robert et al. presented a comprehensive performance evaluation of common TCP variants in LTE networks [9]. The paper contains deep insight on the queueing delay induced by various TCP CC algorithms and it also touches on the performance of these variants in numerous scenarios, involving capacity increase, and short flows.

E. Atxutegi et al. showed a measurement-based study [10] that compares the performance of TCP CUBIC [11] and BBR in LTE networks, both with the use of an emulator and the MONROE measurement testbed [12]. The authors showed that in current LTE cellular networks, BBR minimizes latency, however, it achieved the lowest throughput in some mobility scenarios. BBR achieved good performance when the delays were low, which is a promising result for 5G deployments, however, under longer RTTs in LTE, the performance of BBR degraded.

There are recent simulation studies on TCP performance in 5G mmWave cellular networks [13], [14]. Both were carried out using the mmWave module [15] developed for the ns-3 simulator. In [14] the authors prove the feasibility of using CoDel AQMs in 5G cellular networks and evaluate the performance of different TCP variants in two challenging 3GPP 5G scenarios: the high speed train and the urban deployment environments. Mateo et al. showed a very recent performance evaluation study of TCP congestion control algorithms in mmWave environments [13], focusing on the delay and throughput of the CUBIC, Scalable [16], and BBR algorithms in different scenarios, namely extensive NLOS (Non-line-of-sight), multiple NLOS and multiple short NLOS.

B. Cooperative performance optimization proposals

There are several existing approaches that leverage the cooperation between the network and the end-hosts and thus optimize congestion control by utilizing additional information from the network.

A recent approach belonging to this research direction is presented in [17]. This paper introduces Accel-Brake Control (ABC), a novel congestion control mechanism, where the main goal is to swiftly react to sudden capacity increases, while still being able to promptly react to congestion. The signaling from the network is done by reinterpreting the ECN bits. Deployment options are proposed in the paper for both networks that have legacy ECN enabled, and networks that do not, however, potential fairness issues with competing non-ABC flows are not addressed.

In [18], a cooperative framework is described to enhance TCP performance in LTE networks, and the performance of the solution is evaluated via extensive simulations. The solution is named CDBE (Client Driven Bandwidth Estimation) and the main idea behind it is placing an entity at the bottleneck (eNodeB) that reports a bandwidth estimation to the server, which then can calculate the congestion window and pacing rate accordingly. It is not clear however, how trust can be established between the server and the CBDE client.

Cooperative performance enhancement in 5G cellular networks is targeted in [19], [20]. Milliproxy, presented in [19] applies a flow window policy, modifying the advertised window values of the acknowledgements sent by the client, which are then relayed to the server. The effective congestion window is determined by the server as the minimum of the congestion window and the advertised window. The performance of this solution is evaluated by the authors using the mmWave module of the ns-3 simulator. The authors of [20] describe a different concept, targeting an Edge Cloud scenario in 5G networks. The authors exploit the observation that providing fairness is not the transport protocol's responsibility in the RAN domain, and in an Edge Cloud scenario, there is no "Internet" domain involved. A traffic probe (TP) and a traffic control function (TCF) are introduced, and the initial window is carefully inflated by these entities based on the observed state of the RAN buffer. The paper presents initial performance results, obtained by a modeling approach and the use of real

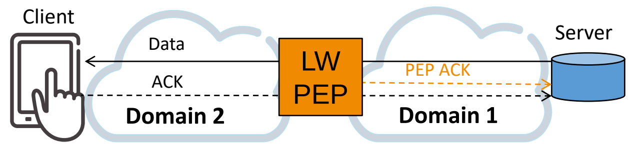

Fig. 1. The Lightweight PEP can be deployed at the border between the wired- and the wireless domains.

data provided by two different LTE network operators. Both concepts have open deployment questions. Milliproxy involves a flow buffer and needs access to the payload which could be problematic when encrypted data is considered.

C. Multi-Domain Congestion Control framework

The proposals discussed in Section II-B do not deal in detail with the concerns related to cooperative solutions in general, such as trust and privacy. We have presented a novel, lightweight cooperative approach in [7] addressing these issues. The Lightweight PEP entity sends safe-to-ignore ACKs and NACKs of the received packets to the server (see Fig. 1). After implementing our Multi-Domain CC algorithm in the Linux kernel (illustrated in Fig. 2), we published the design of the algorithm and initial performance results in [8]. Multi-Domain CC maintains two congestion windows, which are governed by two different algorithms: a conservative CUBIC algorithm and an aggressive, Scalable TCP algorithm in the wired- and wireless domains, respectively. These algorithms are clocked by two different kind of acknowledgements: the client ACKs are clocking the Scalable component, while the PEP-ACKs (the acknowledgements sent by the Lightweight PEP) are used to govern the congestion window of the CUBIC component. The effective congestion window used by the sender is always set to the minimum of the two components, thus, the Multi-Domain algorithm is able to adapt to the location of the bottleneck link. If the bottleneck is in the wired domain, the algorithm behaves conservatively, and provides fairness with other CUBIC flows. However, when the bottleneck is in the cellular domain, MD CC behaves

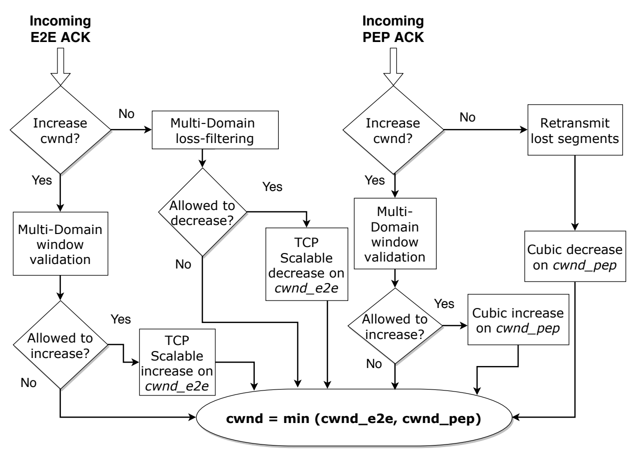

Fig. 2. The Multi-Domain Congestion Control algorithm implemented in the linux kernel

TABLE I
LTE RAN CONFIGURATION

Parameter	ns-3 value
MAC Scheduler	Proportional Fair
Pathloss model	FriisPropagationLossModel
Tx power (dBm)	10 (UE), 30 (eNodeB)
Noise level (dBm)	7 (UE), 5 (eNodeB)
AMC model	MiError
LTE band	7
RLC mode	AM (AQM enabled)
Number of RBs	100

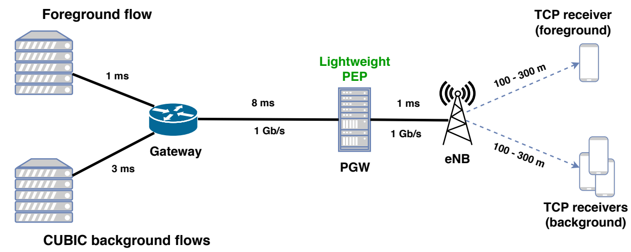

Fig. 3. Simulation topology

aggressively, taking advantage of the fact that in the cellular domain, lower layers are already taking care of the resource sharing, thus, this is not the responsibility of the congestion control algorithm in the transport layer.

III. Simulation environment

Similarly to many of the related papers mentioned in Section II, the performance evaluation was carried out using the ns-3 open source simulator [21]. We also used the Direct Code Execution (DCE) cradle [22], and the NUSE [23] userspace network stack[1]. The simulator and the NUSE stack were extended with the components of our cooperative Multi-Domain framework. In Scalable TCP, the congestion window is increased by a constant parameter after each received ACK. The recommended value for this parameter was 0.01 in [16] which the original kernel implementation changed to 0.02 to account for delayed acknowledgements. In our framework, the Scalable component is set to be more aggressive, with a constant increase parameter of 0.1.

We have used the LTE module of the ns-3 simulator to simulate realistic EPC and RAN behavior. Table I contains the parameters used to configure the RAN domain of the network. In order to enable AQM at eNodeBs, we needed the RLC-AM implementation from [15]. The target delay of the CoDel [24] AQM implemented in the RLC was increased to 20 ms in order to avoid a shallow AQM configuration that would prevent full bandwidth utilization.

We used a trace-based fading model in our simulations with two different fading traces. The EPA (Extended Pedestrian A) trace corresponds to a UE moving with 3 kmph and represents a low delay spread environment, while the EVA (Extended Vehicular A) trace models a UE moving with a speed of 60 kmph, representing a medium delay spread environment. The exact characteristics of these channels are defined in [25].

Figure 3 shows the topology of our simulations. The propagation delay between the Lightweight PEP and the eNodeB is set to 1 ms, this models a ”close” deployment to the eNodeB. Ideally, the Lightweight PEP functionality could be deployed inside an eNodeB.

IV. Reaction to sudden changes in capacity

Sudden and drastic changes in available capacity are inherent characteristics of high-frequency broadband wireless networks, and thus the phenomenon is also present to some extent in LTE, but expected to be even more pronounced in next generation (e.g., 5G) networks. Therefore we chose a specific scenario to investigate how congestion control can cope with sudden capacity changes. Our simulation models a scenario involving sudden increase in available capacity (e.g., a UE entering microcell coverage in LTE). We have implemented this by having 8 background flows stop transferring at a given time and thus leaving only one flow to utilize the resources of the cell.

We chose *capacity fill up time*, i.e., the time it takes to utilize 95% of the new capacity as our performance metric. We have considered three different distances between the UE and the eNodeB: 100 m, 200 m and 300 m. In this section, the capacity fill up times of MD CC and CUBIC (the current default congestion control algorithm in the Linux kernel) are compared in two cases: the vehicular - and the pedestrian sub-scenarios. Each result in Tables II and III are derived from 10 independent simulation runs.

In the vehicular case, the fading channel models a UE moving with a speed of 60 kmph. Figure 4a and Figure 4c show the throughput of the foreground flow for the 100 m and 300 m distances, respectively. The average ”steady state” throughput after the capacity increase is comparable between MD CC and CUBIC (64 Mbps for 100 m distance and 52 Mbps for the 300 m case). The time it takes to utilize 95% of this capacity, however, is significantly reduced by MD CC. Table II shows the reduction achieved by MD CC, which is between 76.6% and 78.4% in the studied range.

In the pedestrian case, the fading channel models a user moving with a speed of 3 kmph. In Figures 4b and 4d it is visible that the average throughputs achieved after the capacity increase are comparable again, however, as the distance between the UE and the eNodeB increases, MD CC starts to achieve higher throughput than CUBIC. This shows that MD CC is able to cope better with the channel quality fluctuations. It is also interesting to note that the channel conditions of the

TABLE II
AVERAGE CAPACITY FILL UP TIME OF DIFFERENT FLOWS IN VEHICULAR FADING ENVIRONMENT

Distance [m]	Capacity fill up time [s]		Reduction [%]
	CUBIC	Multi-Domain CC	
100	9.48 ± 0.54	2.22 ± 0.14	76.6
200	8.18 ± 0.34	1.88 ± 0.19	77.0
300	7.62 ± 0.42	1.64 ± 0.12	78.4

[1]ns-3.26 and DCE version 1.9, Linux kernel version 4.7.0

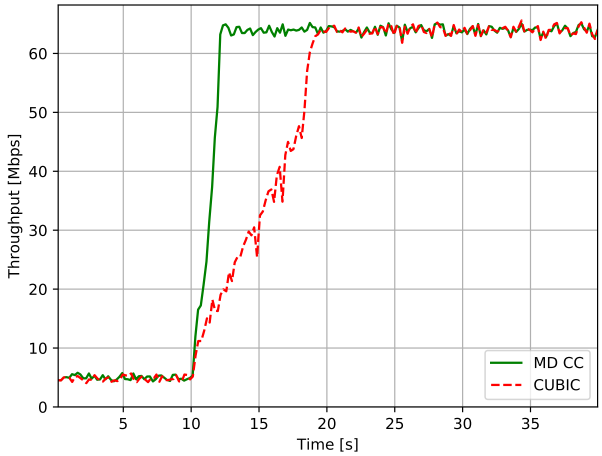

(a) Vehicular fading channel, UE at 100 m distance from the eNodeB

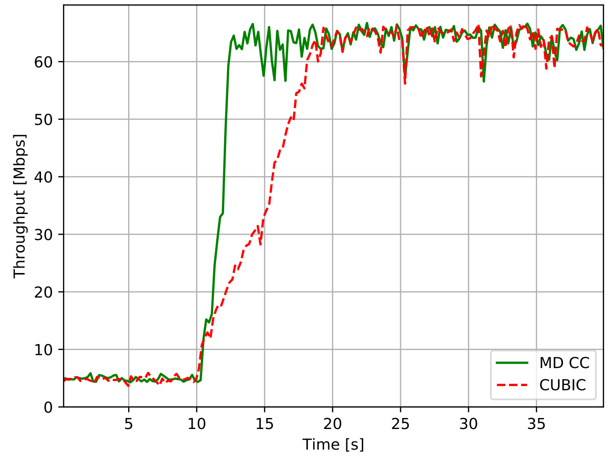

(b) Pedestrian fading channel, UE at 100 m distance from the eNodeB

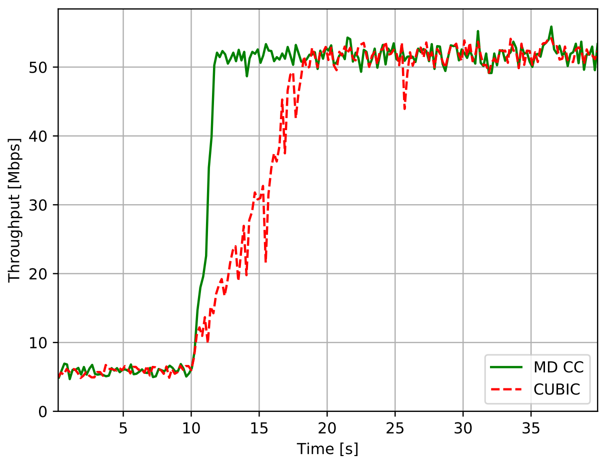

(c) Vehicular fading channel, UE at 300 m distance from the eNodeB

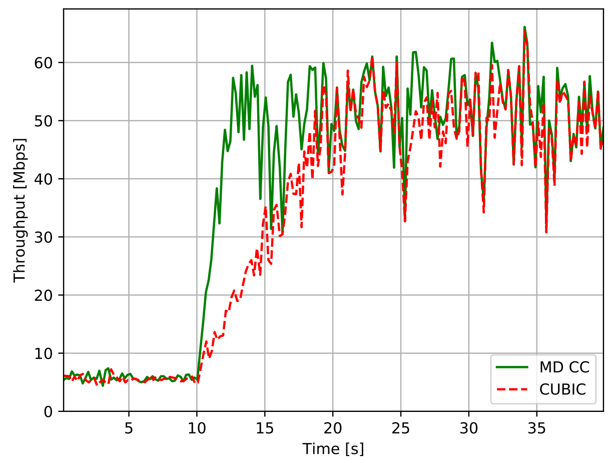

(d) Pedestrian fading channel, UE at 300 m distance from the eNodeB

Fig. 4. Throughput dynamics of different flows during a sudden capacity increase at 10 s

TABLE III
AVERAGE CAPACITY FILL UP TIME OF DIFFERENT FLOWS IN PEDESTRIAN FADING ENVIRONMENT

Distance [m]	Capacity fill up time [s]		Reduction [%]
	CUBIC	Multi-Domain CC	
100	8.80 ± 0.55	2.20 ± 0.17	75.0
200	8.13 ± 0.52	1.98 ± 0.27	75.7
300	6.89 ± 0.64	1.95 ± 0.41	71.7

vehicular model result in more stability in TCP throughput. For larger distances in the pedestrian scenario, the throughput of CUBIC decreases compared to the vehicular case, which is consistent with [26]. As shown in Table III, reduction in capacity fill up time achieved by MD CC is between 71.7% and 75.7%.

V. Short flow completion times

As our Multi-Domain CC algorithm enables faster increase of the congestion window in the slow start phase, it is expected to provide enhanced performance when the downlink traffic consists of short flows. In this section, we investigate and quantify the performance gains.

Figure 3 depicts the simulation scenario for our short flow performance comparison. The delays, link capacities and LTE RAN configurations remain the same, and we chose to use the EVA fading channel for this scenario, based on the observations that it provides higher and more stable throughput. The distance between the UE and the eNodeB is set to 300 m. The foreground traffic consists of short flows (either utilizing MD CC or CUBIC) and the background traffic is a long CUBIC flow. We measured the flow completion time of 50 short flows both with MD CC and CUBIC.

Figure 5 shows the cumulative distribution functions of flow completion times for 500 kB and 1 MB flows using MD CC and CUBIC. It can be seen that for 500 kB flows, the gain is negligible, however, it becomes significant as the size of the short flows increases. For 1 MB flows, MD CC provides a 32.8% average reduction in flow completion times.

Figure 6 shows an example for the evolution of the conges-

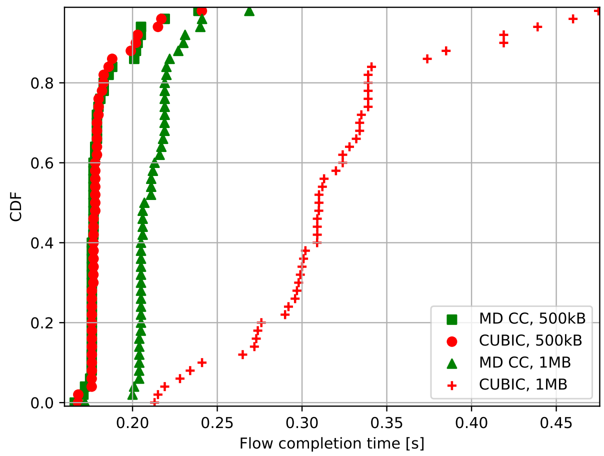

Fig. 5. CDF of flow completion times for flow sizes of 500 kB and 1 MB

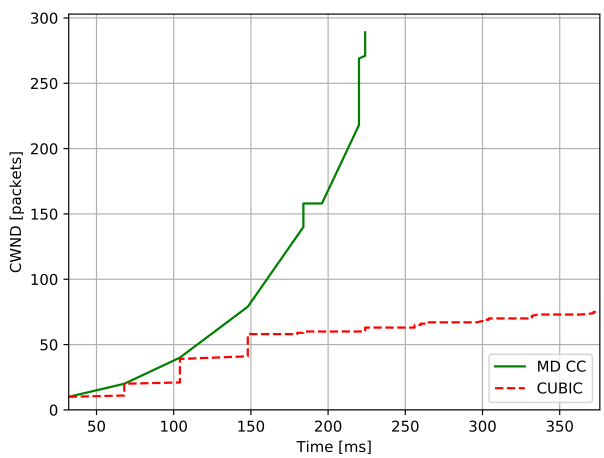

Fig. 6. The congestion window of MD CC and CUBIC during a 1 MB downlink flow

tion window when 1 MB is downloaded from the server. The RTT between the server and the PEP is significantly smaller than the RTT between the sender and the client, thus MD CC is clocked much faster by the PEP-ACKs. The Hybrid Slow Start algorithm in CUBIC switches to congestion avoidance after 150 ms due to the observed increase in packet delays. As this jitter is not present in the "Internet" domain, MD CC avoids the early transition to congestion avoidance and it is able to send the whole 1 MB flow in slow start.

Table IV summarizes the results for the short flow scenario. It can be seen that in the studied range of flow sizes between 500 kB and 10 MB, the reduction in flow completion time ranges between 0% and 84.2%. The presented performance improvement has the potential to deliver highly significant QoE improvements for users in LTE networks for most of the considered flow sizes. Note that the flow sizes with considerable gains fall in the range of downlink TCP flow sizes during web browsing when using HTTP/2 [27].

TABLE IV
AVERAGE COMPLETION TIMES OF SHORT FLOWS

Flow size	Completion time [s]		Reduction [%]
	CUBIC	Multi-Domain CC	
500 kB	0.182 ± 0.013	0.182 ± 0.013	0.0
1 MB	0.316 ± 0.056	0.213 ± 0.014	32.8
2 MB	0.551 ± 0.192	0.214 ± 0.017	61.2
10 MB	1.593 ± 0.415	0.242 ± 0.2	84.2

VI. LONG TERM PERFORMANCE IN THE PRESENCE OF COMPETING MIXED TRAFFIC

The demonstrated performance gains in the capacity increase and short flow use-cases still leave the question of long term performance open. In this section, we illustrate the viability of MD CC in this use-case as well. In order to maintain consistency, the cellular network configuration remains the same as in Section V.

In this scenario, instead of CUBIC -, we consider 3 UDP background flows, each of them generating a realistic aggregate traffic mix. The traffic is generated according to a Poisson Pareto Burst Process (PPBP), which is able to capture the Long-Range Dependent characteristics of flows. The PPBP model consists of overlapping constant bit-rate bursts, where the bursts arrive according to a Poisson process and their length follows a Pareto distribution. This traffic generator was implemented by the authors of [28]. The parameters for the PPBP senders were chosen as follows. We have selected 15 Mbps for the rate, 100 ms for the mean length of the bursts and 3 for the mean number of active bursts. The Hurst parameter was set to 0.7 in accordance with [28].

Figure 7 shows the throughput dynamics of CUBIC in the presence of the above described background flow in a 100 second long simulation and also the the same results of another simulation where the CUBIC is replaced with MD CC. It is apparent that for a significant portion of the simulations, the two algorithms perform similarly, however, when the

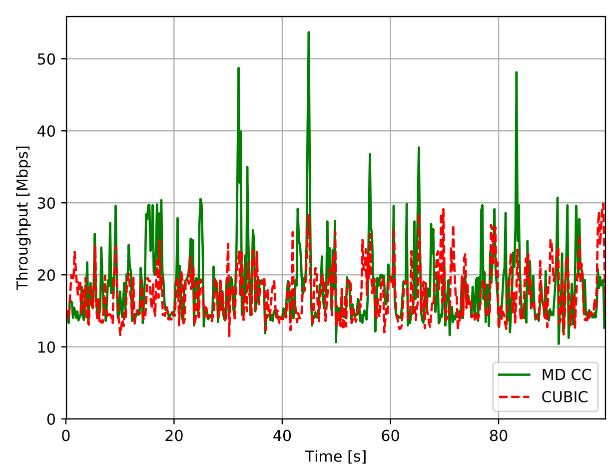

Fig. 7. Throughput of long MD CC and CUBIC flows in the presence of mixed background traffic

available capacity increases for a short time, the Multi-Domain framework is able to capitalize, and thus, overall, improve the average long term throughput. For 10 independent runs, the improvement in throughput was between 4.2% and 9.8%, with an average of 7.3%.

VII. Conclusion

We have proposed Multi-Domain Congestion Control with a cooperative framework where a lightweight, non-ossifying PEP is placed at the border between the wired and the wireless domain of a cellular network. The congestion control algorithm at the sender utilizes the feedback from this PEP to optimize congestion control for both domains.

In this paper, we have evaluated the performance of this Multi-Domain Congestion Control framework in realistic LTE scenarios, and compared the performance to the currently deployed end-to-end TCP CUBIC algorithm.

We have shown that Multi-Domain CC significantly outperforms CUBIC when short flows or sudden capacity increases are considered. Performance gains in the presence of mixed background traffic are also demonstrated. The implications of the results are twofold. First, this means that our cooperative framework could be a non-ossifying alternative to split-connection performance enhancing proxies in current LTE networks. The results also highlight the usefulness of such methods in improving the performance for the more volatile high-bandwidth broadband wireless access networks.

References

[1] Ericsson, "Ericsson mobility report - q4 update," 2019.

[2] N. Cardwell, Y. Cheng, C. S. Gunn, S. H. Yeganeh, and V. Jacobson, "Bbr: Congestion-based congestion control," *Queue*, vol. 14, no. 5, p. 50, 2016.

[3] M. Dong, Q. Li, D. Zarchy, P. B. Godfrey, and M. Schapira, "Pcc: Re-architecting congestion control for consistent high performance." in *NSDI*, vol. 1, no. 2.3, 2015, p. 2.

[4] M. Hock, R. Bless, and M. Zitterbart, "Experimental evaluation of bbr congestion control," in *Network Protocols (ICNP), 2017 IEEE 25th International Conference on*. IEEE, 2017, pp. 1–10.

[5] X. Xu, Y. Jiang, T. Flach, E. Katz-Bassett, D. Choffnes, and R. Govindan, "Investigating transparent web proxies in cellular networks," in *International Conference on Passive and Active Network Measurement*. Springer, 2015, pp. 262–276.

[6] G. Papastergiou, G. Fairhurst, D. Ros, A. Brunstrom, K.-J. Grinnemo, P. Hurtig, N. Khademi, M. Tüxen, M. Welzl, D. Damjanovic *et al.*, "De-ossifying the internet transport layer: A survey and future perspectives," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 1, pp. 619–639, 2017.

[7] A. Mihály, S. Nádas, S. Molnár, Z. Krämer, R. Skog, and M. Ihlar, "Supporting multi-domain congestion control by a lightweight pep," in *Internet of Things, Embedded Systems and Communications (IINTEC), 2017 International Conference on*. IEEE, 2017, pp. 105–110.

[8] Z. Krämer, S. Molnár, A. Mihály, and S. Nádas, "Towards multi-domain congestion control in next-generation networks," in *2019 IEEE Wireless Communications and Networking Conference (WCNC)*. IEEE, 2019. [Online]. Available: http://hsnlab.tmit.bme.hu/~molnar/files/wcnc2019.pdf

[9] R. Robert, E. Atxutegi, A. Arvidsson, F. Liberal, A. Brunstrom, and K.-J. Grinnemo, "Behaviour of common tcp variants over lte," in *2016 IEEE Global Communications Conference (GLOBECOM)*. IEEE, 2016, pp. 1–7.

[10] E. Atxutegi, F. Liberal, H. K. Haile, K.-J. Grinnemo, A. Brunstrom, and A. Arvidsson, "On the use of tcp bbr in cellular networks," *IEEE Communications Magazine*, vol. 56, no. 3, pp. 172–179, 2018.

[11] 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, 2008.

[12] Ö. Alay, A. Lutu, M. Peón-Quirós, V. Mancuso, T. Hirsch, K. Evensen, A. Hansen, S. Alfredsson, J. Karlsson, A. Brunstrom *et al.*, "Experience: An open platform for experimentation with commercial mobile broadband networks," in *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking*. ACM, 2017, pp. 70–78.

[13] P. Jimenez Mateo, C. Fiandrino, and J. Widmer, "Analysis of tcp performance in 5g mm-wave mobile networks," in *Communications (ICC), 2019 IEEE International Conference on*, 2019.

[14] M. Zhang, M. Polese, M. Mezzavilla, J. Zhu, S. Rangan, S. Panwar, and M. Zorzi, "Will tcp work in mmwave 5g cellular networks?" *IEEE Communications Magazine*, vol. 57, no. 1, pp. 65–71, 2019.

[15] M. Mezzavilla, M. Zhang, M. Polese, R. Ford, S. Dutta, S. Rangan, and M. Zorzi, "End-to-end simulation of 5g mmwave networks," *IEEE Communications Surveys & Tutorials*, vol. 20, no. 3, pp. 2237–2263, 2018.

[16] T. Kelly, "Scalable tcp: Improving performance in highspeed wide area networks," *ACM SIGCOMM computer communication Review*, vol. 33, no. 2, pp. 83–91, 2003.

[17] P. Goyal, M. Alizadeh, and H. Balakrishnan, "Rethinking congestion control for cellular networks," in *Proceedings of the 16th ACM Workshop on Hot Topics in Networks*. ACM, 2017, pp. 29–35.

[18] Z. Zhong, I. Hamchaoui, A. Ferrieux, R. Khatoun, and A. Serhrouchni, "Cdbe: A cooperative way to improve end-to-end congestion control in mobile network," in *2018 14th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*. IEEE, 2018, pp. 216–223.

[19] M. Polese, M. Mezzavilla, M. Zhang, J. Zhu, S. Rangan, S. Panwar, and M. Zorzi, "milliproxy: a tcp proxy architecture for 5g mmwave cellular systems," *arXiv preprint arXiv:1712.02700*, 2017.

[20] Å. Arvidsson and L. Westberg, "Fast transport for edge computing in 5g networks," in *2018 26th International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*. IEEE, 2018, pp. 1–5.

[21] "The ns3 simulator," Available: http://www.nsnam.org/, 2016.

[22] H. Tazaki, F. Uarbani, E. Mancini, M. Lacage, D. Camara, T. Turletti, and W. Dabbous, "Direct code execution: Revisiting library os architecture for reproducible network experiments," in *Proceedings of the ninth ACM conference on Emerging networking experiments and technologies*. ACM, 2013, pp. 217–228.

[23] H. Tazaki, R. Nakamura, and Y. Sekiya, "Library operating system with mainline linux network stack," *Proceedings of netdev*, 2015.

[24] K. Nichols, V. Jacobson, A. McGregor, and J. Iyengar, "Controlled delay active queue management," Tech. Rep., 2018.

[25] E. Lte, "Evolved universal terrestrial radio access (e-utra); base station (bs) radio transmission and reception (3gpp ts 36.104 version 8.6. 0 release 8), july 2009," *ETSI TS*, vol. 136, no. 104, p. V8, 2009.

[26] Y. Barayan and I. Kostanic, "Performance evaluation of proportional fairness scheduling in lte," in *Proceedings of the World congress on engineering and computer science*, vol. 2, 2013, pp. 712–717.

[27] "Http archive - state of the web," 2019. [Online]. Available: https://httparchive.org/reports/state-of-the-web

[28] D. Ammar, T. Begin, and I. Guerin-Lassous, "A new tool for generating realistic internet traffic in ns-3," in *Proceedings of the 4th International ICST Conference on Simulation Tools and Techniques*. ICST (Institute for Computer Sciences, Social-Informatics and , 2011, pp. 81–83.