# TRANSPORT LAYER INNOVATIONS FOR FUTURE NETWORKS



Michele Polese Peyman Teymoori Jing Zhu

The communication technologies used for Internet connectivity have radically changed since the 1980s, when the current design of the TCP/IP stack was introduced. Most of the traffic will soon be generated by mobile devices, connected to different generations of wireless networks, which now can provide to end users a multi-gigabit-per-second data rate. Mobile devices are more and more heterogeneous, and support connectivity over different interfaces and networks. Similarly, the capacity of fixed and backhaul networks has dramatically increased, thanks to advances in optical communications.

This has facilitated the development of new applications, such as Virtual and Augmented Reality, Cloud Gaming, and Industry 4.0, among others, that impose stringent end-to-end performance requirements on networks, in terms of, for example, throughput, latency, and reliability. This, along with the increased heterogeneity of the network, has called for an evolution of decades-old transport layer paradigms, bringing the transport layer into the spotlight once again. Indeed, the end-to-end performance and the quality of the service perceived by the users does not depend only on the capabilities of the communication technologies, but also on the interaction among the network, the applications, and the transport layer. Therefore, there has been a recent renewed interest in the design and performance evaluation of new transport protocols, congestion control mechanisms, and solutions to maximize the end-to-end performance of next-generation wireless and wired networks.

This *IEEE Communications Magazine* Special Issue (SI) investigates recent advances at the transport layer that will satisfy the performance requirements of future applications over recent and future networks in the next decade. Notably, the research in this area has focused on three main thrusts. First, the design and proposal of new transport protocols, which can represent an alternative to the common Transmission Control Protocol (TCP). An example is QUIC, a protocol originally introduced by Google and then standardized by the IETF, which introduces multi-stream capabilities, integrates cryptographic functionalities, and connection establishment with zero Round Trip Times (RTTs). Two papers of this SI focus on QUIC, with details on its standardization process and outlook, and on multi-stream scheduler design to maximize the Value of Information (VoI). A third paper of the SI, instead, focuses on the redesign of the Applications Programming Interfaces (APIs) that can be used to communicate with and control the transport layer.

The second trend is represented by the evolution of congestion control techniques, considering also in-network performance enhancing proxies, cross-layer, and/or learning-based approaches. This SI features a paper discussing how in-network approaches based on network coding can enable transport services in a mesh network, and a paper that uses a learning-based method for scheduling at the transport layer.

Finally, the third trend sees the introduction of native multipath capabilities at the transport layer. A notable example is represented by Multipath TCP (MPTCP), the multipath extension of TCP, which is discussed in three papers of this SI. They address the evaluation and design of schedulers for MPTCP for 5G mobile networks, client-based path selection, and application of MPTCP in Internet of Things (IoT) scenarios.

## SUBMISSION STATISTICS AND ORGANIZATION OF THE SI

This SI describes the evolution of the design and technologies for the transport layer, offering the readers a complete overview on the research directions related to multipath transport design, advanced congestion control, TCP alternatives, and also on standardization efforts. The call for papers for the *IEEE Communications Magazine* SI "Transport Layer Innovations for Future Networks" was published in July 2020, and it attracted 19 high-quality submissions from the top researchers in the transport layer area from around the world. After an extensive review process, the editorial team selected seven papers for publication. They can be grouped into three themes, namely: (i) evolution of transport layer APIs and QUIC; (ii) network-based approaches to improve end-to-end performance; and (iii) multipath transport strategies. A brief description of each paper is provided below.

## KEY CONTRIBUTIONS OF ACCEPTED PAPERS

### Evolution of TCP APIs and QUIC

The article "Transport Services: A Modern API for an Adaptive Internet Transport Layer" introduces a new protocol-independent and adaptive transport layer API recently developed by the IETF TAPS (Transport Services) working group, and compares three known implementations in terms of key features and supported protocol stacks. It also demonstrates the flexibility and ease of use of this new API with a Python example from an open-source implementation.

The authors of "Beyond QUIC v1 – A First Look at Recent Transport Layer IETF Standardization Efforts" analyze current transport standardization efforts in the Internet Engineering

Task Force (IETF). In particular, they focus on improvements and increased support for QUIC, including QUIC enhancements, Multiplexed Application Substrate over QUIC Encryption (MASQUE) and WebTransport. MASQUE aims at enabling arbitrary application tunneling over HTTP/3 and QUIC. WebTransport, instead, aims at improving the mapping between QUIC and HTTP/3 and web interfaces.

The article "QUIC-EST: A Transmission Scheme to Maximize VoI of Multi-Stream Correlated Data Flows" proposes a new framework to combine the congestion control and multi-stream features of the recently proposed QUIC transport protocols. It quantifies the utility of the application data through the framework of Value of Information (VoI) and shows the benefits combining QUIC with a new scheduling mechanism that takes VoI and temporal/cross-sensor correlation into consideration.

**Network-based Approaches to Improve End-to-end Performance**

The article "Bringing Network Coding into SDN: A Case-study for Highly Meshed Heterogeneous Communications" presents an architecture for highly-meshed heterogeneous Multi-Source Multi-Destination (MS-MD) networks applicable to 5G wireless networks and beyond. The architecture uses Software-Defined Networking (SDN) principles to deploy adaptive causal network coding in heterogeneous and highly-meshed communication networks. This is an attempt toward solving the challenging problem of utilizing network coding in meshed networks to exploit all available infrastructure/communication mediums.

**Multipath Transport Strategies**

The article "Multipath Scheduling for 5G Networks: Evaluation and Outlook" compares state-of-the-art multipath schedulers with real 5G trace data, considering both static and mobile scenarios. It then introduces M-Peekaboo, a learning-based multipath scheduler, and extends its usage to 5G, and shows the benefit of adopting ML in multipath scheduling decisions, especially in 5G networks with rapidly changing dynamics.

The article "IoT Traffic Offloading with MultiPath TCP" presents an IoT multi-access reference architecture, empowered with a Multipath TCP (MPTCP) gateway connected to multiple cellular networks to offload Machine-to-Machine (M2M) data traffic. This allows devices such as sensors, actuators, and smart meters/monitors to efficiently collect data on remote servers. Performance results of this system show a throughput increase and significant latency reduction as the number of network sensors increases.

The article "Exploiting Client Inference in Multipath TCP over Multiple Cellular Networks" targets the problem of path heterogeneity across LTE networks deployed by different operators. It shows that MPTCP receivers might experience out-of-order queueing, leading to degraded throughput. It then presents a client-based MPTCP (cMPTCP) framework, built upon the Linux MPTCP implementation, that also utilizes the mobile client's network inference information. Experimental results confirm that cMPTCP outperforms existing algorithms in dynamic channel conditions with path heterogeneity.

## Conclusions and Acknowledgment

This SI on transport layer innovations has collected significant contributions that summarize the state of the art on the topic, and also suggest interesting future directions for research on transport layer innovations. We would like to thank the authors of the 19 papers submitted to the SI, and the reviewers who helped us with timely and insightful reviews. We would also like to thank Tarek S. El-Bawab and Antonio Sanchez-Esguevillas, Editor-in-Chief and Associate-Editor-in-Chief, as well as the *IEEE Communications Magazine* editorial staff for their guidance in organizing and preparing the SI.

### Biographies

Michele Polese [M'20] (m.polese@northeastern.edu) is a principal research scientist at the Institute for the Wireless Internet of Things, Northeastern University, Boston. He obtained his Ph.D. from the University of Padova, Italy, in 2020, where he also was a postdoctoral researcher and adjunct professor. During his Ph.D., he visited NYU, AT&T, and Northeastern University. He collaborated with several academic and industrial research partners, including Intel, InterDigital, AT&T Labs, University of Aalborg, King's College and NIST. His research interests are in the analysis and development of protocols and architectures for future generations of cellular networks (5G and beyond), in particular for millimeter-wave communication, and in the performance evaluation of complex networks.

Peyman Teymoori (peymant@ifi.uio.no) received his B.S. and M.S degrees in computer engineering from Ferdowsi University of Mashhad and Amirkabir University of Technology, Iran, in 2001 and 2004, respectively. He has led several software development projects in industry after the Masters' study. He received his Ph.D. degree in computer communication networks from the University of Tehran in 2013. He was a visiting researcher at Gwangju Institute of Science and Technology, South Korea during 2011–2013. He started his position at the University of Oslo (UiO), Norway as a postdoctoral researcher, and he is now a research fellow in the Network and Distributed Systems group, Department of Informatics of UiO. His research interests include network optimization, congestion control, recursive networks, and performance evaluation.

Jing Zhu [M'04, SM'12] (jing.z.zhu@intel.com) 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 software engineer at Intel Corporation. His main research interests are system design, performance optimization, and applications for heterogeneous multi-access/radio networks, including 4G/5G cellular systems, high-density wireless LANs, satellite, and mobile ad hoc networks. He has authored 100+ patent applications and 60+ peer reviewed publications.