

Towards the Transport Protocols of Future Internet

Zoltán Móczár, Sándor Molnár

Abstract—End-to-end congestion control performed by the Transmission Control Protocol (TCP) is the main data transfer mechanism of today’s Internet providing reliable communication between hosts. Since the deployment of TCP the Internet has gone through a significant change due to the evolving network technologies and the diversity of applications. This process has led to a heterogeneous environment with complex traffic characteristics raising the demand for working out different TCP versions to achieve better performance in various network conditions. In addition to the traditional congestion control scheme several alternative solutions have also been proposed for reliable transport. However, the current practice of continuous modification and refinement of TCP for specific network environments does not seem to be a viable option, hence there is an increasing need for a more efficient and flexible transport protocol. In this paper we present a survey of the major data transfer mechanisms developed in the last decades, and advocate a possible direction for future research.

Index Terms—transport protocols, congestion control, fountain coding, future Internet.

I. INTRODUCTION

THE success of the Internet partially stems from the algorithms implemented in the *Transmission Control Protocol (TCP)*. This transport protocol has guaranteed the reliable transfer between end hosts and the stable operation of networks for several decades. However, network environments, applications and user behavior have changed considerably during this long period making TCP suboptimal under different conditions. As a result, a huge number of versions and enhancements of TCP have been proposed for emerging environments mainly focusing on its congestion control scheme and the related mechanisms [1].

Transport layer protocols play a significant role in the efficient and fair utilization of available network resources, and also have a great impact on the quality of user experience. Due to the importance of this research topic, thousands of researchers and developers worldwide are working on more and more efficient transport solutions. The research, development and standardization processes are managed by two large open international communities, the Internet Research Task Force (IRTF) and the Internet Engineering Task Force (IETF), which are organized into different research and working groups, respectively. The main groups focusing on the area of data transport are the *Internet Congestion Control Research Group (ICCRG)* at IRTF and the working groups of the *Transport and Services Area (TSV)* at IETF. The key goal of ICCRG is to move towards consensus on which technologies can

be considered as viable long-term solutions for the Internet congestion control architecture, and to identify the trade-off between potential benefits and costs. As opposed to ICCRG, the members of TSV work on mechanisms related to end-to-end data transfer to support various Internet applications and services that exchange potentially large volumes of traffic at high bandwidths.

This paper presents the evolution of transport protocols since the early days of Internet introducing the main pitfalls the researchers faced with during the years. Furthermore, we shed light on a promising approach, which may be able to satisfy the diverse requirements of future networks.

II. OVERVIEW OF TRANSPORT PROTOCOLS

In current Internet the Transmission Control Protocol carries the vast majority of network traffic. The history of TCP dates back to 1981 when the official protocol specification was published by the IETF in RFC 793 [2]. Over the past three decades a significant research effort has been devoted to TCP in order to meet the requirements of the continuously evolving communication networks. This process has resulted in countless TCP versions aimed to provide high performance in various environments [1]. Although, TCP determined the mainstream of the research on transport protocols, in the last years many alternative proposals have also been published to serve as the basis of reliable data communication. In this section we give an overview of the most widely known protocols including the different types of TCP and other proposals, as well.

A. Transmission Control Protocol

TCP is a connection-oriented transport protocol that provides reliable data transfer in end-to-end communication. It means that lost packets are retransmitted, and therefore, each sent packet will eventually be delivered to the destination. One of the most important features of TCP is its *congestion control* mechanism, which is used to avoid congestion collapse [20] by determining the proper sending rate and to achieve high performance. To this end, TCP maintains a congestion window (*cwnd*) that controls the number of outstanding unacknowledged packets in the network. An important aspect in the context of congestion control protocols is how they can share the available bandwidth among competing flows, also known as *fairness* property. Fairness can be interpreted between the same and different TCP versions (intra- and inter-protocol), as well as on various time scales (transient and steady-state) [21].

TCP variants can be classified based on the type of congestion indication and the target environment as shown in Table I. Most congestion control methods use packet loss information to detect congestion also called as *loss-based* TCPs. In case

Manuscript received June 16, 2014, revised August 31, 2014.

Z. Móczár and S. Molnár are with the Department of Telecommunications and Media Informatics, Budapest University of Technology and Economics, Hungary. Authors are also with the Inter-University Centre for Telecommunications and Informatics, Kassai út 26., 4028 Debrecen, Hungary.

TABLE I
THE EVOLUTION OF TCP VARIANTS

Version	Congestion indicator		Target environment			Implementation					New features	Published
	Loss	Delay	Wired	Wireless	High-speed	BSD	Linux	Win	Mac	ns-2		
TCP Tahoe [2]	×		×			×	×				slow-start, congestion avoidance and fast retransmit	1988
TCP Reno [3]	×		×			×	×	×			fast recovery to mitigate the impact of packet losses	1990
TCP Vegas [4]		×	×				×				bottleneck buffer utilization as a congestion feedback	1995
TCP NewReno [5]	×		×			×	×		×		fast recovery, resistance to multiple losses	1999
Freeze-TCP [6]	×			×			×			×	considering radio signal quality in mobile networks	2000
TCP-Peach [7]	×			×			×				sudden start and rapid recovery for satellite networks	2001
TCP Westwood [8]	×			×						×	estimation of the available bandwidth	2001
ATCP [9]	×			×		×				×	detection of route changes in ad-hoc networks	2001
TCP Nice [10]		×	×				×				delay threshold as a secondary congestion indicator	2002
Scalable TCP [11]	×		×		×		×				MIMD congestion avoidance algorithm	2003
TCP-LP [12]	×		×				×			×	early congestion detection to react sooner than TCP	2003
HighSpeed TCP [13]	×		×		×		×			×	AIMD mechanism as the function of <i>cwnd</i>	2003
FAST TCP [14]		×	×		×		×			×	updating <i>cwnd</i> based on different equations	2003
BIC TCP [15]	×		×		×		×			×	binary search to find the proper <i>cwnd</i>	2004
Compound TCP [16]	×	×	×		×		×	×			calculation of <i>cwnd</i> using loss and delay components	2005
TCP-Illinois [17]	×	×	×		×		×				AIMD as the function of the queueing delay	2006
TCP Cubic [18]	×		×		×		×			×	control of <i>cwnd</i> by applying a cubic function	2008
LEDBAT [19]		×	×				×		×	×	congestion control for low-priority traffic	2012

of these algorithms packet loss is interpreted as the sign of a full network buffer, from which the last incoming packet was dropped, hence transmission rate should be reduced. Another group of congestion control mechanisms react to the increase observed in the round-trip time (RTT) of packets due to the building up of queues. This approach, often referred to as *delay-based* TCP, has the ability to detect congestion early rather than merely waiting until the network gets overutilized and packets are lost. In addition, *hybrid* solutions have also been proposed, which combine the beneficial properties of loss-based and delay-based algorithms.

During the years, the fast development of networks motivated researchers to optimize TCP for certain environments and purposes by modifying the traditional congestion control mechanism. Since standard TCP versions (like TCP Tahoe and Reno) failed to obtain full utilization in networks with high-bandwidth links, new algorithms have been introduced to improve the performance in such conditions. The most relevant *high-speed* TCP versions include Scalable TCP [11], High-Speed TCP [13], FAST TCP [14] and TCP Cubic [18]. On the other hand, as TCP was primarily designed for wired networks, emerging wireless communication induced a considerable research work to develop TCP versions, which can provide better performance in different kinds of *wireless* networks [22]. The performance issues experienced in such environments

stem from the unique characteristics of wireless links and the packet loss model used by TCP. The problems manifest in many applications as degradation of throughput, inefficiency in network resource utilization and excessive interruption of data transmissions. Modification of standard TCP for wireless communication has been an active research area in recent years, and many schemes have been proposed for various environments such as cellular (e.g. Freeze-TCP [6]), satellite (e.g. TCP-Peach [7]) and ad-hoc networks (e.g. ATCP [9]). In real networks a traffic mix consists of hundreds or thousands of flows generated by diverse applications and services. In order to treat low-priority traffic (e.g. background transfers like automatic software updates and data backups) differently from high-priority traffic, *low-priority* congestion control methods have been introduced. These protocols, such as TCP Nice [10], TCP-LP [12] and LEDBAT [19], respond to congestion earlier than standard TCP yielding bandwidth to competing TCP flows with higher priority.

1) *Loss-based TCP Versions*: One of the earliest approaches to handle congestion was introduced in TCP Tahoe [20], which was also served as the first practical implementation of these control schemes in the BSD operating system. The proposal is based on the original TCP specification [2] and introduces new algorithms called slow-start, congestion avoidance (AIMD: additive increase multiplicative

decrease [23]) and fast retransmit, as well as an improved method for round-trip time estimation. These mechanisms allow the sender to detect available network resources and adjust the transmission rate accordingly. However, reducing the congestion window to one packet when a packet loss occurs, as done by Tahoe, is a very aggressive solution.

TCP Reno [3] tackles this problem by applying a novel method referred to as fast recovery algorithm. In case of Reno a lost packet is detected and retransmitted if triple duplicate acknowledgements are received or a timeout event occurs at the sender. This mechanism makes TCP Reno effective to recover from a single packet loss, but it still suffers from performance degradation when multiple packets are dropped from a window of data. To overcome this limitation a selective acknowledgement (SACK) option has been proposed in [24].

TCP NewReno [5] is a variant of TCP Reno intended to improve its performance when a burst of packets is lost. To this end, NewReno modifies Reno's fast recovery algorithm making it possible to recover without a retransmission timeout by resending one packet per each round-trip time until all of the lost packets from the window have been retransmitted.

TCP Cubic [18], being an enhanced version of its predecessor, BIC TCP [15], is one of the most widely used TCP versions today since it serves as the default congestion control algorithm of Linux operating systems. BIC TCP was originally designed to solve the well-known RTT unfairness problem by combining two schemes called additive increase and binary search. TCP Cubic simplifies the window control of BIC and it applies a cubic function in terms of the elapsed time from the last loss event, which provides good stability and scalability. Furthermore, it keeps the window growth rate independent of RTT making the protocol TCP-friendly under both short and long RTT paths.

Beside the congestion control algorithms described above, many other solutions have been worked out to improve the performance of standard TCP. One of the main issues is that it takes a long time to make a full recovery from packet loss for high-bandwidth, long-distance connections, because the congestion window builds up very slowly. In order to cope with this limitation HighSpeed TCP (HSTCP) [13] was proposed, which can achieve better performance on high-capacity links by modifying the congestion control algorithm for use with large congestion windows. Scalable TCP (STCP) [11] applies a multiplicative increase and multiplicative decrease (MIMD) algorithm to obtain performance improvement in high-speed networks and it can also guarantee the scalability of the protocol. TCP Westwood [8] is a sender-side modification of the congestion control mechanism that improves the performance of TCP Reno both in wired and wireless networks. The main problem is that TCP Reno equally reacts to random and congestion losses, thus cannot distinguish between them. In fact, TCP Westwood shows moderate sensitivity to random errors, therefore the improvement is most significant in wireless networks with lossy links. MultiPath TCP (MPTCP) [25] is a recent approach for enabling the simultaneous use of multiple IP addresses or interfaces by a modification of TCP that presents a regular TCP interface to applications, while in fact spreading data across several subflows.

2) *Delay-based TCP Versions:* TCP Vegas [4], as a pioneer of delay-based TCPs, measures the difference (δ) between the expected and actual throughput based on round-trip delays. If δ is less than a lower threshold denoted by α , Vegas assumes that the path is not congested and increases the sending rate. If δ is larger than an upper threshold denoted by β , it is regarded as the strong indication of congestion, hence Vegas reduces the transmission rate. The expected throughput is calculated by dividing the current congestion window by the minimum RTT.

FAST TCP [14] is a congestion avoidance algorithm especially targeted for long-distance, high-latency links. FAST determines the current congestion window size based on both round-trip delays and packet losses over a path. The algorithm estimates the queueing delay of the path using RTTs and if the delay falls below a threshold, it increases the window aggressively. If it gets closer to the threshold, the algorithm slowly reduces the increasing rate.

3) *Hybrid Solutions:* Compound TCP (CTCP) [16], implemented in several Microsoft Windows operating systems, is a synergy of delay-based and loss-based approaches extending the standard TCP Reno congestion avoidance algorithm by a scalable, delay-based component. CTCP exploits the information about both packet loss and delay to control the transmission rate. The delay-based component can rapidly increase the sending rate when the network path is underutilized, but ease if the bottleneck queue becomes full. This mechanism provides good scalability in terms of bandwidth, and a reasonably fair behavior.

TCP-Illinois [17] uses packet loss information to determine whether the congestion window size should be increased or decreased, and measures the queueing delay to determine the amount of increment or decrement. This hybrid solution makes it possible to obtain high throughput and fair resource allocation, while being compatible with standard TCP.

B. Other Proposals

Beyond the Transmission Control Protocol several approaches have also been suggested for reliable data transport in communication networks. Some of these protocols are partially based on the concept of TCP, or use similar mechanisms.

Internet traffic has a complex characteristics investigated in many papers in the last decade. Recent studies showed that most flows are small carrying *only several kilobytes* of data and short lasting *less than a few seconds* [26]. Rate Control Protocol (RCP) [27] is a congestion control algorithm designed to significantly speed up the download of short-lived flows generated by typical applications. For example, a mid-size flow contains 1000 packets and TCP makes them last nearly 10 times longer than it would be necessary. RCP enables flows to finish close to the minimum possible, leading to a notable improvement for web users and distributed file systems.

eXplicit Control Protocol (XCP) [28] uses direct congestion notification instead of the indirect congestion indicators such as packet loss or delay. XCP delivers the highest possible application performance over a broad range of network infrastructures including high-speed and high-delay links where

TCP performs poorly. It also introduces a novel way for separating the efficiency and fairness policies of congestion control, enabling routers to quickly make use of available bandwidth while conservatively managing the allocation of the available bandwidth to competing flows. XCP carries the per-flow congestion state in the packet header allowing the sender to request a desired throughput for its transmission, and XCP-capable routers inform the senders about the degree of the congestion at the bottleneck.

Stream Control Transmission Protocol (SCTP) [29] is a reliable transport protocol that provides stable, ordered delivery of data between two endpoints by using congestion control like TCP and also preserves data message boundaries like UDP. However, unlike TCP and UDP, SCTP offers additional services such as multi-homing, multi-streaming, security and authentication.

Quick UDP Internet Connections (QUIC) [30] is a recent approach for data transfer announced by Google in 2013. QUIC is currently under development and has been integrated in Google Chrome for evaluation purposes. The new protocol supports a set multiplexed connections over UDP, and was designed to provide security protection equivalent to TLS/SSL, along with reduced connection and transport latency. It also implements and applies a bandwidth estimation algorithm in each direction in order to avoid network congestion. QUIC's main goal is to optimize the performance of connection-oriented web applications and services by reducing the connectivity overhead to zero RTT.

III. DIGITAL FOUNTAIN BASED COMMUNICATION

Over the years, the issues of TCP motivated researchers to find alternative ways for data transfer beside the traditional congestion control based approach. In 2007, GENI (Global Environment for Network Innovations) [31] published a research plan, in which they recommend the omission of the congestion control mechanism and suggest to use efficient erasure coding to cope with network congestion. Since then, the questions related to this idea have been investigated only in a few papers. Raghavan and Snoeren argue in [32] that it may not be necessary to keep the network uncongested to achieve good performance and fairness. They introduce a concept called decongestion control and presume that a protocol relying upon greedy, high-speed transmission has the potential to perform better than TCP. Bonald et al. studied the consequences of operating a network without congestion control [33], and concluded that it does not inevitably lead to congestion collapse as believed earlier.

In this section we review the related work carried out in the field of erasure code driven data transport, then introduce and describe a possible data transfer paradigm for Future Internet, which applies a fundamentally different principle compared to that of TCP. According to our concept presented in [34], congestion control can be completely omitted from the transport layer if efficient fountain coding is used as a replacement. We propose a novel network architecture where each host communicates by a digital fountain based transport protocol, while fair schedulers deployed in routers are responsible for

fair bandwidth sharing among competing traffic flows. We show that this new paradigm has many benefits and also introduce some possible future application areas.

A. Error-Correcting Codes in Data Transport

In recent times, many research works have focused on the application of erasure codes in data transport. A theoretical fountain based protocol (FBP) was investigated in [35]. The authors showed that a Nash equilibrium can be reached in a network with FBP-based hosts resulting in a performance similar to the case when each host uses TCP. Kumar et al. proposed a transport protocol for wireless networks using fountain codes [36] and analyzed its performance by a Markovian stochastic model. They demonstrated through packet-level simulations that their protocol may perform better or worse than TCP depending on the redundancy parameter, the number of nodes in a WLAN cell and the wireless channel conditions. The authors of [37] designed a new TCP version on the basis of rateless erasure codes to enhance its operation in lossy environments. According to their results, such modification of TCP has proven to be effective in case of high packet loss rate. Y. Cui and his colleagues proposed FMTCP (Fountain code-based Multipath TCP) in [38], which exploits the advantage of the fountain coding scheme to avoid the performance degradation caused by frequent retransmissions applied in MPTCP. The authors introduced an algorithm to flexibly allocate encoded symbols to different subflows based on the expected packet arrival time over different paths.

B. A Novel Data Transfer Paradigm

1) *Architecture and Protocol*: The key component of our recent proposal is a new transport mechanism called *Digital Fountain based Communication Protocol (DFCP)* [34], which uses digital fountain codes to recover lost packets instead of traditional retransmissions. Fountain codes [39] are rateless erasure codes with the property that a potentially limitless sequence of encoded symbols can be generated from a given set of source symbols, such that the original source symbols can ideally be recovered from any subset of the encoded symbols of size equal to or only slightly larger than the number of source symbols. Raptor codes [40] are the most efficient ones in the family of fountain codes as they offer linear time encoding and decoding complexity, hence the latest version of DFDP implements this scheme.

Our vision of the future network architecture relying on DFDP is shown in Figure 1. There are multiple senders exchanging information with the corresponding receivers, and each host is allowed to send at its maximum transmission rate. Senders generate a potentially infinite stream of encoded symbols from the original message of size k by adding a redundancy of $\epsilon > 0$. When any subset of size $\lceil (1 + \epsilon)k \rceil$ encoded symbols arrive to the receiver, high probability decoding becomes possible, and fountain coding ensures that each received packet at the destination increases the probability of successful decoding. This approach makes it possible to leave the network congested resulting in fully utilized links. To ensure equal bandwidth sharing among competing flows we




Fig. 1. The network architecture built upon DFCP

suggest the use of *fair schedulers* in the network nodes. Several implementations approximating the ideal fair scheduling, such as Deficit Round Robin (DRR) [41], are available and can be configured easily in routers. The feasibility of this solution is supported by the scalability of per-flow fair queueing [42]. We note that maximal rate sending does not mean the full utilization of the transmission capacity available at the sender side in all cases since it would lead to the so-called dead packet problem. This phenomenon happens when a source transmits at a higher speed than its fair share of the bottleneck link needlessly wasting the bandwidth on the whole path from concurrent flows. However, there are many possible ways to avoid this undesirable behavior, for example, by carrying a feedback signal in the acknowledgements about the degree of congestion at the bottleneck (see, e.g. [43]), or by estimating the available bandwidth on the path [44].

2) *Potential Benefits*: The operation of DFCP has been validated on three independent testing platforms [45] including a laboratory testbed, the Emulab network emulation environment [46] and the ns-2 network simulator [47]. In addition, a comprehensive performance evaluation study comparing DFCP to the most relevant TCP versions was also presented in [45]. The main purpose of our investigation was to understand the nature of digital fountain based communication and to reveal its features. Measurements were performed both on simple topologies and in multi-bottleneck networks. The results pointed out that the new paradigm has many potential benefits. One of the main fundamental properties of DFCP is its high resistance to packet losses while it also shows a moderate sensitivity to delay. The latter feature makes it possible to eliminate the well-known RTT unfairness problem of TCP as DFCP can provide fair bandwidth sharing among competing flows independently of their RTTs. Another excellent improvement compared to congestion control based

data transmission is that the concept of DFCP avoids the issues introduced by TCP’s slow-start algorithm, and hence enables both short-lived and long-lived flows to complete faster [48]. Furthermore, our protocol is able to obtain maximum performance even with small buffers, which could make it attractive for all-optical networks. Finally, digital fountain based transport guarantees good scalability and stability as well, both in terms of performance and fairness for increasing number of flows and link capacity.

3) *Possible Applications*: The proposed network architecture can provide an efficient framework for numerous applications. For example, our scheme supports *multipath communication*, which makes it possible to achieve better network resiliency and load balancing. Since DFCP is insensitive to packet loss and delay, it is a good candidate for *wireless networks*, as well. Moreover, the new data transfer paradigm is closely aligned to the high utilization requirement of *data centers* and the concept of *all-optical networking* where only small buffers can be realized.

IV. SUMMARY

In this paper we have reviewed the evolution of transport protocols since the introduction of TCP till today, and discussed the main principles of different congestion control schemes optimized for various target network environments. We have also presented several alternative data transport mechanisms developed in the past decades including recent approaches. We claim that the main lesson learned from this long research history of transport protocols is a need for a paradigm shift. We advocate a promising data transfer method for Future Internet based on digital fountain codes, which can provide more efficient and flexible operation than TCP and may open the way to a broad range of application areas.

ACKNOWLEDGEMENT

This work was partially supported by the European Union and the European Social Fund through project FIRST (grant no. TÁMOP-4.2.2.C-11/1/KONV-2012-0001) organized by ETIK Debrecen. The research also received support from the High Speed Networks Laboratory (HSNLab) at Budapest University of Technology and Economics.

REFERENCES

- [1] A. Afanasyev, N. Tilley, P. Reiher, L. Kleinrock, "Host-to-Host Congestion Control for TCP", *IEEE Communications Surveys and Tutorials*, vol. 12, no. 3, pp. 304–342, 2010.
- [2] J. Postel, "Transmission Control Protocol", *RFC 793, IETF*, 1981.
- [3] K. Fall, S. Floyd, "Simulation-based Comparisons of Tahoe, Reno, and SACK TCP", *ACM SIGCOMM Computer Communication Review*, vol. 26, no. 3, pp. 5–21, 1996.
- [4] L. S. Brakmo, L. L. Peterson, "TCP Vegas: End-to-End Congestion Avoidance on a Global Internet", *IEEE Journal on Selected Areas in Communications*, vol. 13, no. 8, pp. 1465–1480, 1995.
- [5] S. Floyd, T. Henderson, A. Gurtov, "The NewReno Modification to TCP's Fast Recovery Algorithm", *RFC 3782, IETF*, 2004.
- [6] T. Goff, J. Moronski, D. Phatak, "Freeze-TCP: A True End-to-End Enhancement Mechanism for Mobile Environments", *Proceedings of the 19th IEEE International Conference on Computer Communications*, pp. 1537–1545, Tel-Aviv, Israel, 2000.
- [7] I. F. Akyildiz, G. Morabito, S. Palazzo, "TCP-Peach: A New Congestion Control Scheme for Satellite IP Networks", *IEEE/ACM Transactions on Networking*, vol. 9, no. 3, pp. 307–321, 2001.
- [8] S. Mascolo, C. Casetti, M. Gerla, M. Y. Sanadidi, R. Wang, "TCP Westwood: Bandwidth Estimation for Enhanced Transport over Wireless Links", *Proceedings of the 7th ACM MOBICOM Conference*, pp. 287–297, 2001.
- [9] J. Liu, S. Singh, "ATCP: TCP for Mobile Ad Hoc Networks", *IEEE Journal on Selected Areas in Communications*, vol. 19, no. 7, pp. 1300–1315, 2001.
- [10] A. Venkataramani, R. Kokku, M. Dahlin, "TCP Nice: A Mechanism for Background Transfers", *ACM SIGOPS Operating Systems Review*, vol. 36, pp. 329–344, 2002.
- [11] T. Kelly, "Scalable TCP: Improving Performance in High-Speed Wide Area Networks", *ACM Computer Communication Review*, vol. 33, no. 2, pp. 83–91, 2003.
- [12] A. Kuzmanovic, E. W. Knightly, "TCP-LP: A Distributed Algorithm for Low Priority Data Transfer", *Proceedings of the 22th IEEE International Conference on Computer Communications*, pp. 1691–1701, San Francisco, CA, USA, 2003.
- [13] S. Floyd, "HighSpeed TCP for Large Congestion Windows", *RFC 3649, IETF*, 2003.
- [14] C. Jin, D. X. Wei, S. H. Low, "FAST TCP: Motivation, Architecture, Algorithms, Performance", *IEEE/ACM Transactions on Networking*, vol. 14, no. 6, pp. 1246–1259, 2006.
- [15] L. Xu, K. Harfoush, I. Rhee, "Binary Increase Congestion Control (BIC) for Fast Long-Distance Networks", *Proceedings of the 23rd IEEE International Conference on Computer Communications*, vol. 4, pp. 2514–2524, Hong Kong, China, 2004.
- [16] K. Tan, J. Song, Q. Zhang, M. Sridharan, "A Compound TCP Approach for High-Speed and Long Distance Networks", *Technical Report, Microsoft Research*, pp. 1–12, 2005.
- [17] S. Liu, T. Basar, R. Srikant, "TCP-Illinois: A Loss and Delay-Based Congestion Control Algorithm for High-Speed Networks", *Proceedings of the 1st International Conference on Performance Evaluation Methodologies and Tools*, pp. 1–13, Pisa, Italy, 2006.
- [18] I. Rhee, L. Xu, "CUBIC: A New TCP-Friendly High-Speed TCP Variant", *Proceedings of the 3rd International Workshop on Protocols for Fast Long-Distance Networks*, pp. 1–6, Lyon, France, 2005.
- [19] S. Shalunov, G. Hazel, J. Iyengar, M. Kuehlewind, "Low Extra Delay Background Transport (LEDBAT)", *RFC 6817, IETF*, 2012.
- [20] V. Jacobson, "Congestion Avoidance and Control", *Proceedings of the 1988 ACM SIGCOMM Symposium*, pp. 314–329, Stanford, CA, USA, 1988.
- [21] S. Molnár, B. Sonkoly, T. A. Trinh, "A Comprehensive TCP Fairness Analysis in High Speed Networks", *Computer Communications, Elsevier*, vol. 32, no. 13–14, pp. 1460–1484, 2009.
- [22] Y. Tian, K. Xu, N. Ansari, "TCP in Wireless Environments: Problems and Solutions", *IEEE Communications Magazine*, vol. 43, no. 3, pp. 27–32, 2005.
- [23] D.-M. Chiu, R. Jain, "Analysis of the Increase and Decrease Algorithms for Congestion Avoidance in Computer Networks", *Journal of Computer Networks and ISDN Systems*, vol. 17, no. 1, pp. 1–14, 1989.
- [24] M. Mathis, J. Mahdavi, S. Floyd, A. Romanow, "TCP Selective Acknowledgment Options", *RFC 2018, IETF*, 1996.
- [25] A. Ford, C. Raiciu, M. Handley, S. Barre, J. Iyengar, "Architectural Guidelines for Multipath TCP Development", *RFC 6182, IETF*, 2011.
- [26] S. Molnár, Z. Móczár, "Three-dimensional Characterization of Internet Flows", *Proceedings of the 2011 IEEE International Conference on Communications*, pp. 1–6, Kyoto, Japan, 2011.
- [27] N. Dukkipati, N. McKeown, "Why Flow-Completion Time is the Right Metric for Congestion Control", *ACM SIGCOMM Computer Communication Review*, vol. 36, no. 1, pp. 59–62, 2006.
- [28] D. Katabi, M. Handley, C. Rohrs, "Congestion Control for High Bandwidth-Delay Product Networks", *ACM SIGCOMM Computer Communication Review*, vol. 32, no. 4, pp. 89–102, 2002.
- [29] R. Stewart, Q. Xie, K. Morneault, C. Sharp, H. Schwarzbauer, T. Taylor, I. Rytyina, M. Kalla, L. Zhang, V. Paxson, "Stream Control Transmission Protocol", *RFC 2960, IETF*, 2000.
- [30] J. Brodtkin, "Google Making the Web Faster with a Protocol that Reduces Round-Trips", <http://arstechnica.com/information-technology/2013/06/google-making-the-web-faster-with-protocol-that-reduces-round-trips/>, 2013.
- [31] D. Clark, S. Shenker, A. Falk, "GENI Research Plan (Version 4.5)", April 23, 2007.
- [32] B. Raghavan, A. C. Snoeren, "Decongestion Control", *Proceedings of the 5th ACM Workshop on Hot Topics in Networks*, pp. 61–66, Irvine, CA, USA, 2006.
- [33] T. Bonald, M. Feuillet, A. Proutiere, "Is the 'Law of the Jungle' Sustainable for the Internet?", *Proceedings of the 28th IEEE Conference on Computer Communications*, pp. 28–36, Rio de Janeiro, Brazil, 2009.
- [34] S. Molnár, Z. Móczár, A. Temesváry, B. Sonkoly, Sz. Solymos, T. Csicsics, "Data Transfer Paradigms for Future Networks: Fountain Coding or Congestion Control?", *Proceedings of the IFIP Networking 2013 Conference*, pp. 1–9, New York, NY, USA, 2013.
- [35] L. López, A. Fernández, V. Cholvi, "A Game Theoretic Comparison of TCP and Digital Fountain Based Protocols", *Computer Networks, Elsevier*, vol. 51, no. 12, pp. 3413–3426, 2007.
- [36] D. Kumar, T. Chahed, E. Altman, "Analysis of a Fountain Codes Based Transport in an 802.11 WLAN Cell", *Proceedings of the 21st International Teletraffic Congress*, pp. 1–8, Paris, France, 2009.
- [37] A. Botos, Z. A. Polgar, V. Bota, "Analysis of a Transport Protocol Based on Rateless Erasure Correcting Codes", *Proceedings of the 2010 IEEE International Conference on Intelligent Computer Communication and Processing*, vol. 1, pp. 465–471, Cluj-Napoca, Romania, 2010.
- [38] Y. Cui, X. Wang, H. Wang, G. Pan, Y. Wang, "FMTCP: A Fountain Code-Based Multipath Transmission Control Protocol", *Proceedings of the 32nd IEEE International Conference on Distributed Computing Systems*, pp. 366–375, Macau, China, 2012.
- [39] D. J. C. MacKay, "Fountain Codes", *IEE Proceedings – Communications*, vol. 152, no. 6, pp. 1062–1068, 2005.
- [40] A. Shokrollahi, "Raptor Codes", *IEEE Transactions on Information Theory*, vol. 52, no. 6, pp. 2551–2567, 2006.
- [41] M. Shreedhar, G. Varghese, "Efficient Fair Queuing Using Deficit Round-Robin", *IEEE/ACM Transactions on Networking*, vol. 4, no. 3, pp. 375–385, 1996.
- [42] A. Kortebi, L. Muscariello, S. Oueslati, J. Roberts, "On the Scalability of Fair Queuing", *Proceedings of the 3rd ACM Workshop on Hot Topics in Networks*, pp. 1–6, San Diego, CA, USA, 2004.
- [43] B. Briscoe, R. Woundy, A. Cooper, "Congestion Exposure (ConEx) Concepts and Use Cases", *RFC 6789, IETF*, 2012.
- [44] C. D. Guerrero, M. A. Labrador, "On the Applicability of Available Bandwidth Estimation Techniques and Tools", *Computer Communications, Elsevier*, vol. 33, no. 1, pp. 11–22, 2010.
- [45] Z. Móczár, S. Molnár, B. Sonkoly, "Multi-Platform Performance Evaluation of Digital Fountain Based Transport", *IEEE Science and Information Conference 2014*, London, UK, 2014.
- [46] Emulab Network Emulation Testbed, <http://www.emulab.net/>
- [47] ns-2 Network Simulator, <http://www.isi.edu/nsnam/ns/>
- [48] S. Molnár, Z. Móczár, B. Sonkoly, "How to Transfer Flows Efficiently via the Internet?", *Proceedings of the 3rd IEEE International Conference on Computing, Networking and Communications*, pp. 462–466, Honolulu, HI, USA, 2014.



Zoltán Móczár received his M.Sc. degree in computer science from the Budapest University of Technology and Economics (BME), Hungary in 2011. Currently he is a Ph.D. student at the High Speed Networks Laboratory, Department of Telecommunications and Media Informatics, BME. His research interests include traffic engineering, performance evaluation of high-speed transport protocols and Future Internet design. He has participated in several academic and applied research projects.



Sándor Molnár is an Associate Professor at the Department of Telecommunications and Media Informatics, Budapest University of Technology and Economics (BME). In electrical engineering he received his M.Sc. in 1991 and Ph.D. in 1996 from BME. Dr. Molnár is serving in the editorial board of international journals and is working in the TPC of IEEE conferences. He has been participating in various EU projects and was the BME project leader of the Gold Award winner 2009 CELTIC project titled "Traffic Measurements and Models in Multi-Service Networks (TRAMMS)". His main interests include teletraffic analysis and performance evaluation of modern communication networks.