Computer Networks 105 (2016) 1–21

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Is multi-path transport suitable for latency sensitive traffic?

Kiran Yedugundla[a,*], Simone Ferlin[b], Thomas Dreibholz[b], Özgü Alay[b], Nicolas Kuhn[c,d], Per Hurtig[a], Anna Brunstrom[a]

[a] Dept. of Computer Science, Karlstad University, Karlstad, Sweden
[b] Simula Research Laboratory, Oslo, Norway
[c] IMT Télécom Bretagne, IRISA, Cesson-Sévigné, France
[d] Centre National d'Etudes Spatiales (CNES), Toulouse, France

ARTICLE INFO

Article history:
Received 3 December 2015
Revised 13 April 2016
Accepted 11 May 2016
Available online 21 May 2016

Keywords:
Internet
Latency
Multi-path communication
Transport protocols
MPTCP
CMT-SCTP

ABSTRACT

This paper assesses whether multi-path communication can help latency-sensitive applications to satisfy the requirements of their users. We consider Concurrent Multi-path Transfer for SCTP (CMT-SCTP) and Multi-path TCP (MPTCP) and evaluate their proficiency in transporting video, gaming, and web traffic over combinations of WLAN and 3G interfaces. To ensure the validity of our evaluation, several experimental approaches were used including simulation, emulation and live experiments. When paths are symmetric in terms of capacity, delay and loss rate, we find that the experienced latency is significantly reduced, compared to using a single path. Using multiple asymmetric paths does not affect latency – applications do not experience any increase or decrease, but might benefit from other advantages of multi-path communication. In the light of our conclusions, multi-path transport is suitable for latency-sensitive traffic and mature enough to be widely deployed.

1. Introduction

Live and interactive applications are sensitive to latency, as the user experience is negatively affected when data is delayed. For instance, freezing a live video just 1% of the video duration is sufficient to turn away 5% of the viewers [1]. Similarly, a latency of 60 ms suffices to degrade user experience in Internet gaming [2]. Multiple ways of improving the user experience of latency sensitive applications are active subjects of research. However, as far as we know, a weakly explored area is to determine whether utilizing all available network interfaces at the end host could improve such experience. In recent times, deployed devices such as tablets and smartphones are often equipped with both Wireless LAN (WLAN) and cellular 3G or 4G interfaces.

Multi-path transmission has been proposed to guarantee better resilience to link failures and a better use of resources. For instance, consider a connection using two interfaces simultaneously; if one of the interfaces (or underlying links) fails, the transmission can simply continue over the other interface. In a single-interface scenario, the transmission would be stalled and maybe require a connection re-establishment. It has also been shown that simultaneous transmission of data over multiple interfaces can increase the throughput, due to capacity aggregation [3]. Even if multi-path protocols have been shown to be more resilient to link failures and able to aggregate capacity to provide increased throughput, the impact of using multiple paths on latency has not been thoroughly investigated.

This paper fills this gap by assessing whether multi-path approaches are suitable transport protocols for applications transmitting latency-sensitive traffic, e.g., video, gaming and web traffic. Recent efforts within the Internet Engineering Task Force (IETF) include designing Multi-path TCP (MPTCP) [4] extensions to TCP [5] to enable end-to-end connections to span multiple paths simultaneously. Similarly, Concurrent Multipath Transfer for SCTP (CMT-SCTP) [6–8] is an extension to the Stream Control Transmission Protocol (SCTP) [9], enabling simultaneous multi-path communication. We therefore evaluate their suitability to carry out latency sensitive traffic.

In our experiments we consider both symmetric multi-path communication (e.g. WLAN-WLAN) as well as asymmetric (e.g. WLAN-3G). For the actual evaluations we use a combination of simulations, emulations and real experiments to ensure a correct assessment.

The remainder of this paper is structured as follows. Section 2 presents an overview of CMT-SCTP and MPTCP, and how

* Corresponding author.
E-mail addresses: kiran.yedugundla@kau.se, kirayedu@kau.se (K. Yedugundla), ferlin@simula.no (S. Ferlin), dreibh@simula.no (T. Dreibholz), ozgu@simula.no (Ö. Alay), nicolas.kuhn@cnes.fr (N. Kuhn), perhurt@kau.se (P. Hurtig), annabrun@kau.se (A. Brunstrom).

http://dx.doi.org/10.1016/j.comnet.2016.05.008
1389-1286/© 2016 Published by Elsevier B.V.

these protocols solve the core issues inherent in transport-level multi-path communication. Section 3 describes the applications used in our evaluation and their latency requirements. In Section 4, the experimental setup is detailed. Section 5 presents and explains the results obtained. In addition to that, Section 6 provides an in-depth discussion of the results. Section 7 discusses related work on multi-path transport. Finally, Section 8 concludes the paper and discusses possible future work in this area.

2. Multi-path transport

This section introduces CMT-SCTP and MPTCP, the current key multi-path transport protocols. The core issues of multi-path communication, and how these are addressed by CMT-SCTP and MPTCP, are then described.

2.1. CMT-SCTP

SCTP [9,10] is a transport protocol originally developed by the IETF Signaling Transport (SIGTRAN) Working Group [11], as part of an architecture to provide reliable and timely message delivery for Signaling System No. 7 (SS7) [12] telephony signaling information, on top of the Internet Protocol (IP) [13]. While motivated by the need to carry signaling traffic, SCTP was designed as a general purpose transport protocol on par with TCP [5] and UDP [14]. While SCTP can offer functionality similar to TCP, such as ordered and reliable transmission or congestion controlled transport, its options can be easily set so that SCTP rather features unordered transmission or multi-homing. This flexibility is one main advantage of SCTP as opposed to TCP.

The multi-homing feature of SCTP allows a single association (or connection) between two endpoints to combine multiple source and destination IP addresses. These IP addresses are exchanged and verified during the association setup, and each destination address is considered as a different path towards the corresponding endpoint. Using the Dynamic Address Reconfiguration protocol extension [15], it is also possible to dynamically add or delete IP addresses, and to request a primary-path change, during an active SCTP association.

While SCTP multi-homing [9,10] targets robustness and uses only one active path at a time, several researchers have suggested the concurrent use of all paths for sending data. Budzisz et al. [16] provides a survey of these approaches. In this paper, we consider the most complete of these proposals, Concurrent Multipath Transfer for SCTP (CMT-SCTP) [6–8]. CMT-SCTP improves the internal buffer management procedures of SCTP, transmission over multiple paths and reordering with its single sequence-number space. Assuming disjoint paths, CMT-SCTP applies the original SCTP congestion control [9] for each path independently.

2.2. MPTCP

Multi-Path TCP (MPTCP) [17] is a set of extensions to TCP [5,18] developed by the IETF MPTCP working group [19] to enable simultaneous use of multiple paths between endpoints. The motivation behind MPTCP is more efficient resource usage and improved user experience through improved resilience to network failure and higher throughput.

To use the MPTCP extensions the initiator of a connection appends a “Multipath Capable” (MP_CAPABLE) option in the SYN segment, indicating its support for MPTCP. When the connection is established, it is possible to add one TCP flow, or subflow, per available interface to this connection by using a “MPTCP Join” (MP_JOIN) option in the SYN segment. Once the MPTCP connection has been fully established, both end hosts can send data over any of the available subflows.

While MPTCP transparently divides user data among the subflows, simultaneous transmission may cause connection-level packet reordering at the receiver. To handle such reordering, two levels of sequence numbers are used. Apart from the regular TCP sequence numbers that are used to ensure in-order delivery at subflow level, MPTCP uses a 64-bit data sequence number that spans the entire MPTCP connection and can be used to order data arriving at the receiver. To ensure fairness [20] on bottleneck links shared by subflows of a MPTCP flow and other TCP flows, MPTCP extends the standard TCP congestion control. Running existing TCP congestion control algorithms independently would give MPTCP connections more than their fair share of the capacity if a bottleneck is shared by two or more of its subflows. To solve this MPTCP uses a coupled congestion control [21] that links the increase functions of each subflows’ congestion control and dynamically controls the overall aggressiveness of the MPTCP connection. The coupled congestion control also makes resource usage more efficient as it steers traffic away from more congested paths to less congested ones.

2.3. Core issues

This section presents the core issues that are related to the use of multiple paths and how they are addressed by CMT-SCTP and MPTCP.

2.3.1. Path management

As shown in Fig. 1, a path is a sequence of links between a sender and a receiver [4], over which it is possible to open a subflow. A multi-path protocol must define a path management strategy. The strategy needs to find suitable paths to open subflows over and decide whether one or more subflows should be opened over a specific path. For short or extremely time-sensitive flows, the choice of path for the initial connection establishment might be very important. For example, if (i) two paths (p_1 and p_2) are available, (ii) both paths have the same capacity and (iii) the RTT of p_1, r_1, is significantly higher than the RTT of p_2, r_2, (e.g. $r_1 > 10 \times r_2$), then whether the first subflow will be opened over p_1 or p_2 would seriously impact the latency. The number of subflows to open over a path is a problem that is not very well studied. While the Linux implementation of MPTCP supports this using its ndiffports path manager, as described later in this section, it is often regarded as unnecessary to open more than one subflow per path as they typically would traverse the same links and compete for the same network resources. However, in some specific environments, e.g. datacenters, the network might conduct load balancing between subflows, routing them over disjoint subpaths. In such situations there might be benefits of creating several subflows per path, as shown in [22].

For CMT-SCTP a path is defined by the destination IP address and port number. To manage paths, CMT-SCTP employs a simple

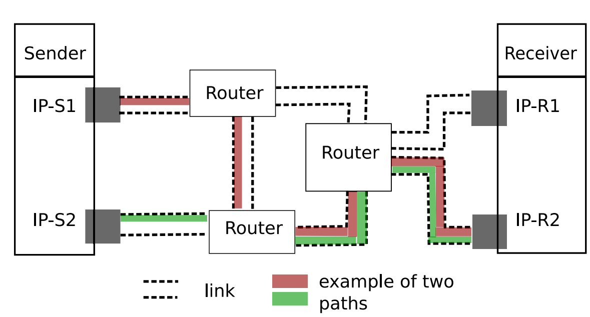

Fig. 1. Definition of a path as a sequence of links between a sender and a receiver.

strategy where the association is established during a 4-way handshake in which available IP addresses are exchanged and verified. The handshake is conducted over the default interface of the host, and after its completion each destination address is considered as a path and implicitly also as an opened subflow. The interfaces of the hosts are pairwise connected over two different subnets, resulting in two possible paths. In addition to the establishment of subflows during the association setup, there is an extension to CMT-SCTP called Dynamic Address Reconfiguration (DAR) [15] which enables an end-host to dynamically add and remove IP addresses to an existing connection.

Like CMT-SCTP, MPTCP consults the routing table to determine which interface to initiate the connection over. During the establishment phase, realised by a 3-way handshake, IP address information are exchanged between the hosts in a fashion similar to that of CMT-SCTP. However, after connection establishment, MPTCP cannot make full use of the other host's IP address information and start sending data over all paths straight away. Notably, at this point it can only use the path used for the connection establishment. Additional subflows must be opened as separate TCP connections and joined to the MPTCP connection using the "MPTCP Join" option in their SYN segments. Another difference, as compared to CMT-SCTP, is the availability of multiple path managers in MPTCP. For example, the Linux implementation of MPTCP provides four different path managers: default, full-mesh, ndiffports and binder. Using the default path manager, a host does not advertise additional IP addresses but uses the other hosts advertised IP addresses to create new subflows. The full-mesh strategy uses an opposite approach: all available IP addresses are exchanged and used to open a subflow over each and one of all the possible source-destination IP address combinations. The ndiffports manager allows a user to open *X* subflows over the default interface. Finally, the binder manager, implements Loose Source Routing as defined in [23]. Similar to CMT-SCTP, MPTCP also includes an extension to allow dynamic addition and removal of IP addresses.

2.3.2. Scheduling

If multiple subflows are available, there are different ways to schedule the transmission of data. As an example, a round-robin scheduler may iterate over the available subflows and try to transmit an entire congestion window over each subflow, while another scheduler might only consider the "fastest" available subflow. Scheduling in multi-path communication has therefore a large impact on the performance of data transmission.

One root cause for scheduling problems is the use of paths with asymmetric characteristics. Fig. 2 illustrates how asymmetric paths, in terms of RTT ($RTT_2 = 10 \times RTT_1$), can affect data transmission. Fig. 2(c), shows the so-called head-of-line blocking problem. In this scenario, packets #3 and #4 (residing in the receiver's buffer) cannot be delivered to the application as packets #1 and #2 are still in flight. Further, packets sent over the fast subflow can fill the receiver's buffer while waiting for data transferred over the slow subflow. This issue is known as receiver buffer blocking and is illustrated in Fig. 2(d).

The usual scheduler in CMT-SCTP is a round-robin scheme targeting throughput maximization: for every subflow in sequence, starting from the primary one, it sends as much data over the subflow as the congestion window allows. As mentioned earlier, asymmetric paths can be problematic and this is also true for CMT-SCTP. The problem is due to a combination of the scheduling and occupancy of the shared send and/or receive buffer space, and can cause the aforementioned problems of head-of-line blocking and receiver buffer blocking. Detailed classifications of the blocking issues are provided in [24]. To remedy this problem, other schedulers have been proposed and developed for CMT-SCTP. For example, chunk rescheduling [8,25] is a mechanism that re-injects the segment causing head-of-line blocking on a different subflow that has space available in its congestion window. Furthermore, Delay-Aware Packet Scheduling (DAPS) [26] is a scheduler that, given the RTT of the different subflows, tries to send packet sequences over them in a manner that guarantees in-order delivery at the receiver.

Similar to CMT-SCTP, several schedulers have been proposed for MPTCP. In the Linux implementation the default scheduler always tries to transmit data over the subflow with the shortest RTT, as long as there is free space in the congestion window. The default scheduling also includes a mechanism called Retransmission and Penalization (RP). This mechanism is similar to CMT-SCTP's chunk rescheduling and re-injects segments causing head-of-line blocking in a different subflow. In addition to the default scheduling mechanism, a weighted round-robin scheme is also available. The schedulers available for Linux have all been evaluated and compared in [27], identifying the shortest-RTT scheduler as the most successful in terms of throughput performance. The different schedulers are detailed in [28].

2.3.3. Congestion control for multi-path transport

When using multiple paths for transmission, different subflows cannot share a single congestion window, as each subflow is likely to have different characteristics and levels of congestion. There are, however, situations in which the subflows actually do share a bottleneck and thus have the same level of congestion. In such scenarios there must be some kind of collaboration between the congestion controllers of each subflow to ensure that the transport does not achieve more than its fair share of the network resources.

For CMT-SCTP there is no default congestion control mechanism that manages the transmission based on the combined congestion state of the subflows. This is likely due to an initial design assumption that subflows do not share bottlenecks. As discussed above, such an assumption is not always true, making CMT-SCTP potentially unfair to other traffic in the network. This problem has been addressed by several researchers and coupled congestion controllers have been proposed. Examples include e.g., CMT/RPv1 [29] and CMT/RPv2 [30].

The problem of not considering shared bottlenecks was addressed already in the design phase of MPTCP. The reason to why coupled congestion control should be used, the benefits of using it, and what goals it has to achieve are all documented in [21]. To achieve these goals, various coupled congestion control schemes have been proposed for MPTCP. These include the Linked-Increases Algorithm (LIA) [21], the Opportunistic Linked-Increases Algorithm (OLIA) [31] and the BALanced Linked Adaptation (BALIA) [32]. At the time of writing, the default congestion control in the Linux MPTCP implementation is LIA.

2.3.4. Handling loss and retransmissions

When data is lost in multi-path transmission the protocol must decide whether to retransmit this data over the same subflow or over a different one.

CMT-SCTP features several schemes for retransmitting data, all detailed in [33]. A CMT-SCTP sender maintains accurate information about the working paths, as new data are transmitted over every available subflow concurrently. Therefore, many distinct strategies can be used. For example, retransmitting lost data over the same subflow, over the subflow with the largest slow-start threshold or using the subflow with the largest congestion window. While there is no default retransmission strategy for CMT-SCTP, we consider retransmissions over the subflow with lowest RTT to give latency-sensitive applications a benefit.

In MPTCP, the loss detection is performed at two levels: subflow level and MPTCP level. While loss typically is detected on subflow level, different strategies can be taken depending on how the

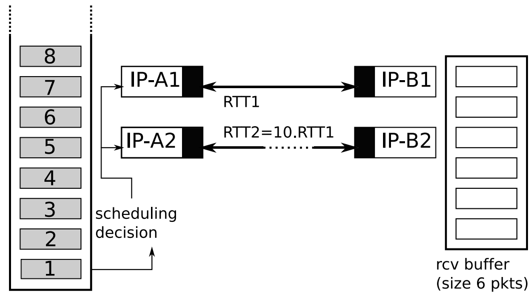

(a) At t_0; asymmetric path $RTT_2 = 10 \times RTT_1$.

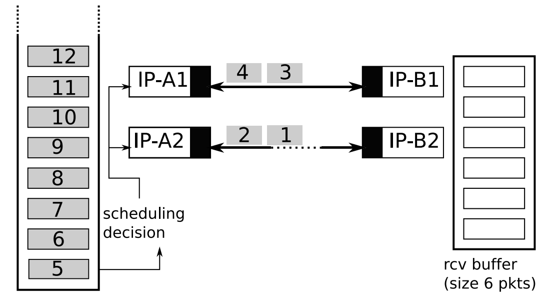

(b) At $t_1 = t_0 + \epsilon$; #1 and #2 on path 1; #3 and #4 on path 2.

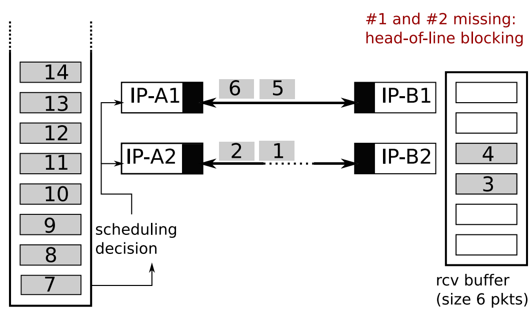

(c) At $t_2 = t_1 + RTT_1$; #3 and #4 are received but can not be read because #1 and #2 are missing $\Rightarrow$ head-of-line blocking.

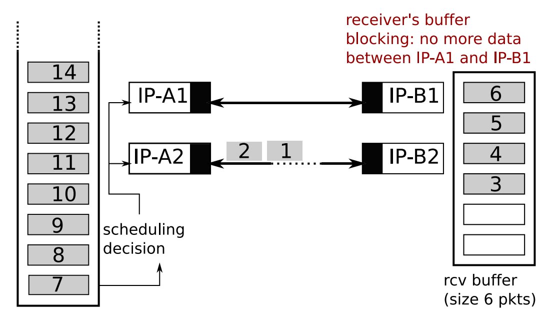

(d) At $t_3 = t_1 + 2 \times RTT_1$; #1 and #2 are not received and the receiver's buffer can not receive #7 $\Rightarrow$ receiver buffer blocking.

Fig. 2. Head-of-line blocking and receive buffer blocking.

loss was detected. If the loss is detected by the fast retransmit algorithm, data is only retransmitted over the same subflow. If, on the other hand, the loss is detected by an expiration of the RTO timer, the data can be retransmitted over both the same subflow and over an additional subflow chosen by the scheduler. The rationale for using different retransmission approaches depending on how the loss was detected is straightforward; fast retransmit only detects loss if feedback from packets sent after the lost packet(s) arrive at the receiver. Therefore, it is safe to assume that no massive congestion event or link breakage has happened, and that a retransmission will arrive safely at the receiver. If no feedback is received, however, the RTO will eventually expire, and it is then safer to retransmit the lost packet(s) over both paths in case the path over which the original transmission occurred is experiencing major congestion or other serious problems.

3. Applications and their requirements for multi-path transport

Traditionally, Internet has been dominated by web traffic running on top of short-lived TCP connections [34]. For example, Ciullo et al. [35] found that approximately 95% of the client TCP flows and 70% of the server TCP flows were less than 10 segments. Although web traffic still constitutes a large fraction of all traffic, video traffic and gaming traffic are now becoming more common. Recent measurements [36] show e.g. that more than 53% of the downstream traffic in North America is video streaming. Forecasts (cf. [37]) also show that Internet video and gaming will continue to grow with an annual compound growth rate of 29% for video traffic and 22% for gaming.

Although the aforementioned traffic classes differ significantly in many ways, they have a common property – sensitivity to latency. In this paper we will therefore use video, gaming and web traffic to assess whether multi-path protocols are suitable for latency-sensitive applications. The remainder of this section describes the main characteristics of the applications and discusses their requirements.

3.1. Video streaming

There are two main use-cases of video streaming: Video on Demand (VoD) which is not broadcast live and therefore do not have stringent latency requirements; and direct live video which is broadcast live and have requirements of low latency.

VoD applications has complete knowledge of the content to transfer, and can therefore adapt the sending rate appropriately. The quality of experience of VoD is therefore less vulnerable to one-way delay variations than the quality of experience of direct live video. As the rationale of this work is to assess whether multipath transport protocols can be used for latency sensitive applications, we will focus on direct live video as it is more sensitive to latency.

The direct live video can be divided into two sub-categories: live broadcast of TV such as BBC iPlayer[1] and private video communications such as Skype.[2] We have focused on the latter category as such applications typically are interactive in their nature and thus more sensitive to latency. Although Skype is proprietary and its communication protocol is closed, and may change over time, we use Skype-like video traffic in our evaluation. We do this for several reasons. First, Skype is a widely used application. Actually, Skype generates almost two percent of the total aggregate

[1] http://www.bbc.co.uk/iplayer/live/bbcnews/.
[2] https://www.skype.com/.

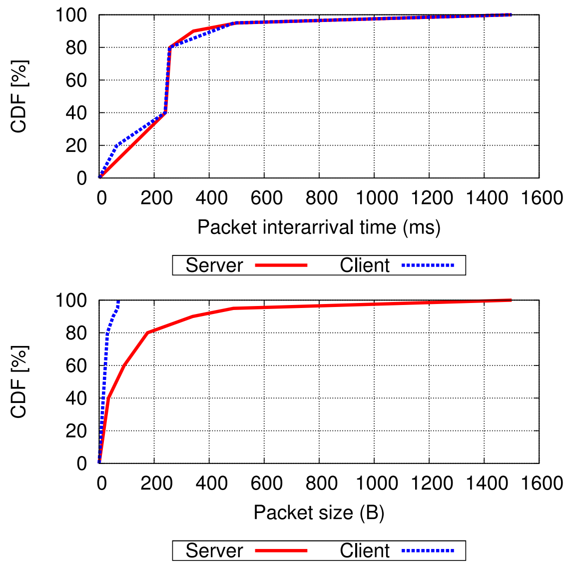

Fig. 3. Distribution of packet inter-arrival time and packet size for World of Warcraft [43].

traffic in European fixed networks [38]. Second, although Skype mainly tries to use UDP for communication NATs and firewalls often force it to use TCP, making it an interesting use case for our experiments with multi-path and reliable transports. Finally, Skype traffic is well studied and traffic characteristics have been reported by several researchers, making it relatively easy to model. According to [39,40], it dynamically adapts its sending rate to the network conditions, with a frame rate per second going from 5 frames/s to 30 frames/s and a video bit rate from 30 Kbit/s to 950 Kbit/s.

Requirements: The latency requirements for a good user experience when considering live video communication are: one-way delay should be lower than 150 ms [41] and the difference in delay between packets (jitter) should be lower than 30 ms [41].

3.2. Gaming traffic

Online gaming is often categorized into three different classes [42] each of them being characterized by specific traffic, as detailed in [43]. The classes are:

- first person avatar, e.g. First Person Shooter games (FPS)
- third person avatar, e.g. Massive Multiplayer Online games (MMO)
- omnipresent, e.g. Real Time Strategy games (RTS)

FPS games are tolerant to loss but are very delay sensitive, therefore they often use UDP as transport. MMO games, on the other hand, are less loss tolerant and require less bandwidth compared to FPS games, therefore, a mixture of TCP and UDP is used for transmission. TCP traffic of MMO games is composed of multiple thin TCP flows. Thin flows are characterized by a low transmission rate where the majority of packets are much smaller than the maximum transmission unit (MTU). An example of the traffic characteristics between the server and a client of an MMO game is illustrated in Fig. 3. For RTS games, interestingly, latency has a negligible effect on the outcome of the game, indicating that RTS game-play clearly favors strategy over the real-time aspects [44].

Considering the popularity of MMO games [43], and the fact that they use TCP, this paper assesses whether there are any benefits in using multiple paths at the transport layer to carry the traffic generated by an MMO game entitled Age of Conan.

Requirements: The requirements for a good gaming experience highly depend on the class of the game and the particular game itself. However, low latency (lower than 60 ms is indicated in [2]) and a small delay variation [45] are important for a good gaming experience.

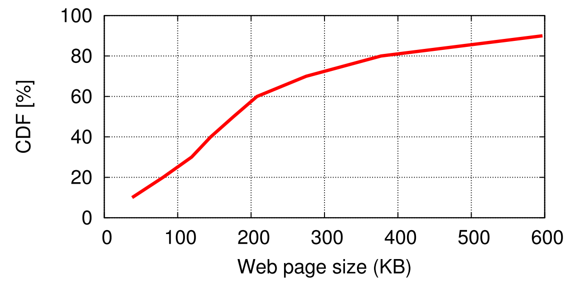

Fig. 4. Distribution of web page sizes according to [47].

3.3. Web traffic

Fig. 4 illustrates the distribution of common web site sizes. In our experiments, to be representative of the web, we have selected three web sites of different sizes: small (72 KiB), medium (1024 KiB) and large (3994 KiB). We also use software that emulates the behavior of a real browser, downloading the web sites using 6 concurrent connections over HTTP/1.1. More details on the web traffic can found in Section 4.3.3.

Requirements: The quality of user experience when accessing a web page is highly linked to the download completion time. For example, in [46], the authors report that Google measured that “an additional 500 ms to compute (a web search) [...] resulted in a 25% drop in the number of searches done by users.”. Although the download completion time may not be the most relevant metric for modern browsers, as they often start rendering pages before completion, it is the most suitable metric to use when evaluating transports as it is “browser agnostic” and therefore neutral. For good web browsing experience, the download completion time is required to be as low as possible.

4. Experiment setup

This section describes the experiment setup used for the performance evaluation of our target applications. The protocol implementations and network models used in the evaluations are also introduced.

4.1. Evaluation tool sets

In this study, we focus on the default and latest versions of the protocols. Our choice of evaluation tool sets, presented in this section, has been made on the basis of availability of source code and the fact that we wanted to consider both controlled and real life experiments.

4.1.1. Simulations, CMT-SCTP using OMNeT++

There is no stable implementation of CMT-SCTP for FreeBSD or Linux. Therefore, we could not perform emulations or real experiments using CMT-SCTP. Instead, we performed simulations using OMNeT++ [48] version 5.0b1 with the CMT-SCTP model [8,49,50], and the NetPerfMeter application model [8,51] in the latest version of the INET Framework [52], using the simulation processing tool-chain SimProcTC [53,54]. For the web traffic simulations in Section 5.3, the HttpTools [55] models provided as part of the INET Framework are used. It was only necessary to add SCTP support. The complete INET Framework sources branch used for this paper is available online.[3] Most changes have already been merged upstream.

Although there is an implementation of CMT-SCTP available for NS-2 [6], it is unmaintained and as of spring 2016 it is fairly out of date. The OMNeT++ implementation, which is used for our evaluations, includes the latest improvements and options for SCTP and is therefore representing the state-of-the-art in SCTP features.

4.1.2. Emulations, MPTCP in a controlled environment, using CORE

Because we wanted to evaluate MPTCP in a controlled environment, we ran experiments using the Linux MPTCP implementation and the Common Open Research Emulator (CORE) [56]. CORE enables the use of real protocols and applications together with emulated network links, making the evaluation of MPTCP easy to control and replicate. The Linux kernel implementation is the most complete MPTCP implementation available, so this setup also allowed us to use the most feature complete version of MPTCP. Using the same MPTCP implementation for both the controlled experiments in CORE and the real life testbed experiments also allowed us to more easily compare and validate the results.

4.1.3. Experiments, MPTCP in a real-life environment, using NorNet

Because we may drive biased conclusions if the protocols were to be evaluated only in controlled environments, we also assessed their performance using an environment where the network is used by many other applications than the one we introduce in the network.

In order to realize this, we performed real network experiments on a dedicated testbed, namely NorNet Edge (NNE) [57]. NNE is a multi-homed testbed where each node is connected to multiple UMTS operators via Huawei E392-u12 modems as well as WLAN network. More specifically, in our evaluations, we consider one operational UMTS Mobile Broadband (MBB) network in Oslo, Norway. It is labelled as "3G". The WLAN access point is a public WLAN hotspot, connecting around 100 people during work hours in a large office complex with several interfering WLAN networks. The two WLAN networks used for the WLAN-WLAN scenarios are using the same technology (IEEE 802.11ag) and sharing the same medium, therefore certain level of interference is highly likely depending on the number of users and traffic patterns. We believe, this setup reflects a realistic scenario where the users cannot control these factors. The downside of this testbed experiments is that there might be a statistically insignificant and uncontrolled behavior for few packets. Furthermore, in this paper, we focus on the transport layer, therefore we study how transport layer reacts to such realistic path characteristics. In this real-world environment, MPTCP was tested under different scenarios for different applications as in the emulation setup.

4.2. Configuration of MPTCP and CMT-SCTP

Both MPTCP and CMT-SCTP have open source implementations, making it possible to enable and/or disable specific features. Due to readability we have chosen to only present the most important features, and their settings, in this section. Additionally, Table 1 provides a short summary of this information. For a full description of the protocol configurations and for all the experimental scripts and data, please see [58].

For MPTCP, we used the state-of-the-art Linux MPTCP implementation (v0.89.3).[4] We use the default options of MPTCP, including e.g. receive buffer optimization and coupled congestion control, with an exception for the Nagle algorithm, which is turned off in all application scenarios. Turning off Nagle is common practice when running applications that require low latency [59].

The simulation uses the CMT-SCTP model for OMNeT++ that is fully described in [8,49]. As opposed to MPTCP, no default options are given for CMT-SCTP. The latest version of the SCTP simulation model [49,50] for OMNeT++ is used, implementing SCTP according to RFC 4960 [9] with all state-of-the-art features and extensions. In addition to the settings listed in Table 1, CMT-SCTP was configured with the following features:

- burst mitigation with MaxBurst=4 (default from [9, Section 15]) with "Use It or Lose It" [60] strategy (i.e. behavior like the FreeBSD SCTP implementation [50]);
- buffer splitting [24,25] to avoid buffer blocking issues [8, Section 7.5].

All data is sent in SCTP/MPTCP messages of up to 1452 bytes (resp. 1428 bytes), which corresponds approximatively to full packets (including headers) of 1500 bytes. In CMT-SCTP, the size of the payload depends on the number of chunks that are gathered in one message, therefore the full packet sizes may vary depending on the application profile.

As explained in Section 2.3.1, (i) with CMT-SCTP, one subflow can be opened on each working path as soon as the 4-way handshaking process has been operated and the primary path (i.e., the first path on which data is transmitted) has to be defined; (ii) with MPTCP, the first subflow that is opened depends on the parameterization of the Linux default interface. In our evaluations, when the paths are homogeneous (i.e., WLAN-WLAN or 3G-3G), the path on which the first subflow is opened is chosen randomly; when the paths are heterogeneous (i.e., WLAN-3G), the WLAN path is used for the first subflow.

4.3. Application traffic generation and metrics

This section presents how we generate video streaming, gaming and web traffics. The rationale for using these applications and the characteristic of the traffic that they generate are detailed in Section 3.

4.3.1. Video traffic

In this paper, we have not considered Video on Demand traffic since it is hard to accurately model or emulate this traffic in the various cases (emulation, simulation, experimentation) used in this article. Moreover, these applications are not interactive and might be seen as file transfer applications. Therefore, we considered Direct Live Video applications due to their delay sensitive nature. In order to generate Skype-like traffic we considered a constant bit rate application generating 950 Kbit/s with 30 frames/s.

4.3.2. MMO gaming traffic

For gaming traffic, we considered a set of trace files from the Massively Multiplayer Online Game Age of Conan, provided by Funcom [61]. These traffic traces extend over a very long time period. As it is extremely difficult and tedious to replay all the traces completely, we selected a set of three traces with a duration of 10 minutes each. The selection of traces was based on the possible full game play being captured in the trace. Full game play constitutes initial loading of game settings, player interaction and infrequent chunk updates depending on the game. All the traces contain a huge chunk of game setup data in the beginning of the connection followed by occasional small bursts of MTU-sized packets and small packets for the rest of the time. The selected traces were replayed using the D-ITG [62] traffic generator. In this process D-ITG is loaded with full trace and it generates packets of exact size and time sequence as seen in the trace file. This is a way of providing trace input to the experiments than generation based on a statistical setting. From here on, the traces are named Trace 1, Trace 2

[3] https://github.com/dreibh/inet/tree/td-netperfmeter-for-integration.
[4] Linux MPTCP: http://www.multipath-tcp.org.

Table 1
Options for MPTCP and CMT-SCTP.

	Path management Section 2.3.1	Scheduling Section 2.3.2	Congestion control Section 2.3.3	Handling loss Section 2.3.4
MPTCP	full-mesh	(i) LowRTT (ii) Retransmission and penalization	coupled (OLIA)	(i) Fast retransmit on the same subflow (ii) RTO on a subflow chosen by the scheduler
CMT-SCTP	one subflow per working path	(i) packet-based round-robin (ii) chunk rescheduling	uncoupled (NewReno)	retransmission on lowest RTT path

Table 2
Web traffic generation.

Size	Domain name	Number of objects	Size of objects
small	Wikipedia (www.wikipedia.org)	15	72 KiB
medium	Amazon (www.amazon.com)	54	1024 KiB
large	Huffington Post (www.huffingtonpost.com)	138	3994 KiB

and Trace 3. They have average packet inter-departure times of 181.4 ms, 74.1 ms, and 167.7 ms respectively. Furthermore, the average packet sizes are 142.7 bytes, 113 bytes, and 101.7 bytes respectively.

4.3.3. Web traffic

As previously mentioned, we consider three classes (small, medium, large) which are representative for real web sites. The classes and sites are presented in Table 2. For the experiments, we stored the files from the three sites (Wikipedia, Amazon and Huffington Post) on a local server. Each website data contained different umber of objects of different sizes. Of the three, Wikipedia content was the smallest followed by Amazon and Huffington Post. The data stored in the local server was requested and downloaded from a client with 6 concurrent connections.

4.4. Background traffic generation

The congestion level in a network has a significant impact on the behavior of protocols employing congestion control [63]. We therefore conduct experiments both with and without background traffic. Background traffic is generated with NetPerfMeter [51,64] as a mix of TCP and UDP flows constituting one long TCP flow and 4 UDP on-off flows. The TCP flow has a saturated sender sending as much data as possible with frame size of 1460 bytes. Each UDP flow generates Pareto on-off traffic with shape 1.5 and scale 0.166667, sending 25 frames per second each of size 5000 bytes. The aggregate usage of UDP background flows were maintained at 10% of the bottleneck link capacity to be realistic [65]. The UDP flows carry data at an average of 500 Kbit/s each in the WLAN-WLAN scenario and 100 Kbit/s each in the 3G-3G scenario. In each run, the background flows start before the foreground experimental traffic and end after the experimental traffic.

However, in the NorNet experiments even running one experiment with the background traffic, especially for online gaming and video streaming, can eat up all the monthly data quota. Therefore, we decided not to generate background traffic due to the limited data quotas for the 3G subscriptions and in order to provide consistent results, we have run all NorNet experiments without background traffic.

4.5. Network and system characteristics

4.5.1. Topology

Fig. 5 shows the topology used in our evaluations. The same topology was used for the simulations, emulations, and real experiments in NorNet. To understand the basic performance of protocols and highlight their characteristics, a simple topology is more useful than a complex topology. Though the topology can be seen as more general, this type of basic topology is also common practise for evaluating transport protocols, see for instance Common TCP Evaluation Suite.[5] As shown in the figure, there are two paths on the client side and a single path on the server side. In the simulations and emulations the paths that connect the server and client are, of course, modeled. For the experiments to be realistic we used a parameterisation of the paths that is based on measurements that were conducted over NorNet prior to the evaluation.

4.5.2. Path characteristics

Table 3 shows the capacity, end-to-end delay and packet loss rates of the WLAN and 3G paths that have been measured in the experimental testbed described in Section 4.1. The WLAN links are IEEE 802.11ag. The loss rate is what is experienced on the transport layer (e.g. datagrams), therefore it is the loss ratio after all link layer re-transmission schemes of underlying networks. Path 1 and Path 2 in Fig. 5 will be assigned with these characteristics depending on the technology. In Table 3, we also detail the three combinations of paths over which the multi-path protocols are evaluated, that are homogeneous WLAN (two WLAN paths), homogeneous 3G (two 3G paths) and heterogeneous (one WLAN and one 3G path).

[5] https://tools.ietf.org/html/draft-irtf-iccrg-tcpeval-01.

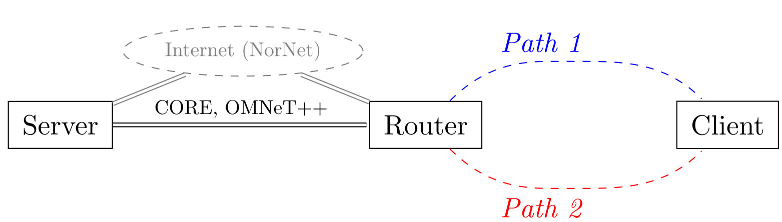

Fig. 5. Topology.

Table 3
Path characteristics and scenarios.

	WLAN	3G
Capacity [Mbit/s]	20–30	3–5
Propagation Delay [ms]	20–25	65–75
Loss [%]	1–2	0
Homogeneous (WLAN)		
Path 1	x	
Path 2	x	
Homogeneous (3G)		
Path 1		x
Path 2		x
Heterogeneous		
Path 1	x	
Path 2		x

4.5.3. Buffer sizes

System characteristics of the source and destination are known to impact the end to end performance of the flows. In order to emulate the realistic network scenarios, we use the system settings close to the standard settings for respective technologies. The TCP buffer sizes (send buffer/receive buffer) are set to be equivalent to the default Android settings, that are configured as follows:

- Homogeneous (3G): 256 KiB/256 KiB.
- Homogeneous (WLAN): 1024 KiB/2048 KiB.
- Heterogeneous (WLAN-3G):
  1024 KiB/2048 KiB.

Based on estimations from early measurement in the NorNet testbed, performed during the planning phase of the work, the queue lengths at each interface of the router (see Fig. 5) are set to 100 packets for WLAN and 3750 packets for 3G.

Note, that the 3G buffer setting of 256 KiB/256 KiB prevents an overly large bufferbloat [66,67] in a 3G/3G setup, while the setting of 1024 KiB/2048 KiB will make such a bufferbloat in the WLAN-3G case possible. We will explain this in detail with the results in Section 5 and particularly in Section 5.1.1.

5. Experiment results

This section presents the experimental evaluation and its results. The protocols are first evaluated through simulations and emulations in controlled environments, to identify the impact of various network parameters. This evaluation is then complemented with measurement results from a real environment.

For each experiment scenario, SCTP is compared to CMT-SCTP, and TCP Cubic is compared to MPTCP. For the homogeneous cases, we consider the average delay using TCP and compare it with that of MPTCP. We assume that we only have information about the technologies used. That is, for a WLAN-WLAN case, the WLAN channels might have different characteristics in terms of loss and delay, but this information is not available to the user. The user will most likely pick one of the WLANs randomly. Therefore, we consider the average WLAN TCP delay performance and compare it with the MPTCP delay performance. However, for the WLAN-3G scenario, the user will most certainly choose WLAN, since it has low delay, high capacity and is probably cheaper to use. Therefore, we compare the MPTCP delay with the TCP delay of WLAN. The evaluation of SCTP and CMT-SCTP is conducted in the same fashion.

5.1. Video streaming

First, we evaluate video traffic performance for homogeneous and heterogeneous scenarios, considering both competing and non-competing traffic as explained in Section 4.3. We use application layer message delay as the performance metric for the transport protocol latency performance.

5.1.1. CMT-SCTP simulations

Fig. 6 presents the average message delays and the variation in these delays in the form of box plots [68] for video traffic as described in Section 4.1. First, we consider the average SCTP message delay over 128 runs for video traffic without any competing traffic and illustrate the results in Fig. 2(a). Table 4 presents the percentage of traffic sent over the different paths in CMT-SCTP. Note, that although SCTP uses a certain primary path for payload data transport, there is always a small amount of control traffic (here: mainly heartbeats to check the path status, see [9]) on the other path as well.

In the homogeneous scenarios (WLAN-WLAN and 3G-3G) we observe different behaviors for WLAN-WLAN and 3G-3G cases. In the WLAN-WLAN scenario, with two similar WLAN paths, multi-path transport leads to increased latency, and also increased latency variation, mainly due to the reordering caused by retransmissions on the lossy paths (i.e., 1%–2% packet loss; see Table 1). In the 3G-3G scenario, where both paths are lossless, the performance of SCTP and CMT-SCTP are virtually the same.

For the WLAN-3G scenario, we observe that multi-path transport leads to increased delays and delay variation. The main reason for this is the reordering caused by the heterogeneous path characteristics.

For CMT-SCTP, it is important to note that SCTP has its origins as transport protocol for signaling systems (see Section 2.1), where networks are designed for specific applications. Therefore, the most important implementations – the OMNeT++ simulation model as well as CMT-SCTP implementation in FreeBSD – currently only provide a very simple scheduler; data is scheduled on the paths in a round-robin fashion. The intention of using this scheduler is to improve throughput, without caring about path delays. That is, once data to send is available, and a path's congestion window allows to send it, as much as possible is sent on this path. Then, for further data to be sent, the next path is tried. This mechanism is tightly combined with the SCTP burst mitigation that limits the amount of consecutive packets sent at once over a path. Both SCTP implementations apply burst mitigation by using the "use it or lose it" [60] strategy, with a setting of MaxBurst=4 (default from [9, Section 15]). That is, if a certain number of bytes α is acknowledged by the receiver side the sender would be allowed to send up to α new bytes. The limit of in-flight bytes is given by the congestion window. However, if a non-saturated sender does not fully utilize its allowance given by the congestion window, the congestion window is reduced to the number of in-flight bytes *plus* MaxBurst$*$MSS. As a result, using MaxBurst=4, only up to 4 packets are sent on a path before the next path is used. The round-robin scheduler together with the burst mitigation results in an increase in the message delay for CMT especially when the paths are heterogeneous or when there exists loss.

Next, the performance of CMT-SCTP was evaluated in the presence of competing background traffic (see Section 4.4) and the corresponding average message delays are illustrated in Fig. 2(b).

In the WLAN-WLAN scenario, similar to the case without background traffic, multi-path transport leads to increased latency and latency variation due to the reordering caused by retransmissions. We also see that the background traffic has no visible impact in the WLAN/WLAN scenario due to the random packet loss on the path. As the background flows are experiencing loss, the TCP background flow backs off before it causes any noticeable congestion, and as the UDP background traffic is only 10% of the link capacity it also has very limited impact. Hence in the WLAN-WLAN scenario there was no noticeable effect of congestion losses or queuing delay on

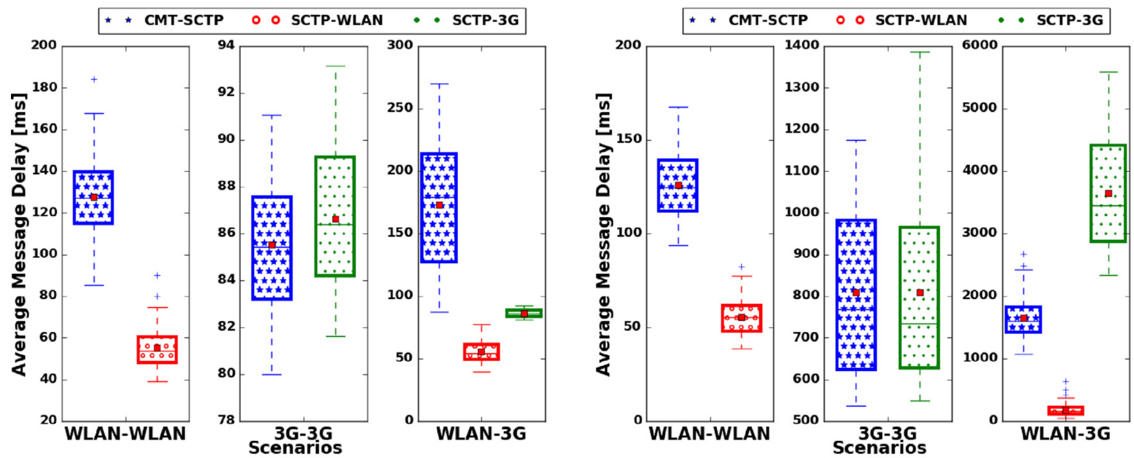

(a) Without Background Traffic (b) With Competing Background Traffic

Fig. 6. Average message delay for CBR video traffic over CMT-SCTP.

Table 4
Path 1 Traffic share (in %) for CBR video traffic over CMT-SCTP.

Traffic	Background	WLAN-WLAN		3G-3G		WLAN-3G	
		WLAN	WLAN	3G	3G	WLAN	3G
Video on CMT-SCTP	✗	17.5	82.5	6.3	93.7	16.4	83.6
	√	17.5	82.5	4.5	95.5	79.2	20.8

the foreground flows as compared to the corresponding scenario without background traffic.

For the 3G-3G scenario, also similar to non-competing traffic, we observe no significant performance difference for CMT-SCTP as compared to SCTP. Note that in this scenario, the send and receive buffer sizes are 256 KiB (see Section 4.5.3), and the background traffic leads to an increased delay due to bufferbloat. We observe an average delay of around 800 ms when background traffic exists as compared to 80-90 ms when there is no competing traffic.

The WLAN-3G results differ quite much from the non-background results. Note that for WLAN-3G case, we have different buffer settings: a send buffer of 1024 KiB and a receive buffer of 2048 KiB (see Section 4.5.3). These buffer settings allow the queue on the 3G path to grow, causing a significant bufferbloat. Here, we observe the delay on the 3G path jumping to values of almost 4 s, making any interactivity virtually impossible. Using CMT-SCTP leads to a significant reduction of the delay – due to the additional usage of the low-latency WLAN path – to values of around 1.8 s. However, this delay is still much higher compared to the delay of SCTP on the WLAN path.

5.1.2. MPTCP emulation

Fig. 7 (a) presents the average message delay for video traffic in all the scenarios considered, i.e., WLAN-WLAN, 3G-3G, WLAN-3G. Each plot represents the data for 30 repetitions when no background traffic was present.

In the WLAN-WLAN scenario, the path delay difference between the two paths is small (i.e., 20 ms – 25 ms; see Table 1). Thus, the main factor that determines the average message delay is the link losses. Loss on one path causes the scheduler to push data on the other path and eventually exploiting the availability of multiple paths. If both paths have similar delay and loss as per the setup, data is sent on both paths causing traffic to oscillate between paths. Such oscillation of traffic between paths is known as flapping. In some of the repetitions, we observed that flapping and losses caused data to arrive out-of-order at the receiver, resulting in increased delays. However, on average, we observed that MPTCP improves the delay performance and reduces the delay variation as compared to TCP. This is in contrast to the relationship between CMT-SCTP and SCTP discussed above. The difference is due to the fact that MPTCP uses a lowest-RTT scheduler (see Section 4.5.3) that moves the sending of data between the paths in a better way as loss occurs. The limitation imposed by MaxBurst in CMT-SCTP also results in a less suitable distribution of the data as compared to MPTCP where the congestion window stays larger.

Table 5
Video traffic data share per path using MPTCP.

Traffic	Background	WLAN-WLAN		3G-3G		WLAN-3G	
		WLAN	WLAN	3G	3G	WLAN	3G
Video	✗	54.19	45.81	99.84	0.16	51.96	48.04
	√	55.67	44.33	85.24	14.76	51.98	48.02

In the 3G-3G scenario, there are no losses and all the data is sent over only one path. Hence, there is no performance differences between MPTCP and TCP. Table 5 provides some insights on the share of data over each 3G path. The delay difference between the paths is small, but still enough to make the scheduler use only one of the paths. Certain configurations start with a non-optimal interface as default, and in those cases the scheduler eventually switches to the other path. This is evident in Table 5, where 0.16% of the data was sent on the other 3G path.

In the heterogeneous scenario (WLAN-3G), the WLAN path clearly has a lower average delay than the 3G path, even for small amounts of loss on the WLAN link. The behavior of the default scheduler ensures that MPTCP uses the path with lowest RTT. However, the performance of MPTCP was observed to be worse than that of TCP in this scenario. MPTCP uses both paths due to losses in the WLAN, triggering transmission over the 3G path which otherwise would not be used due to the large path delay differences. In the case of video traffic considered for the experiments, the data share shown in Table 5 should be identical to the

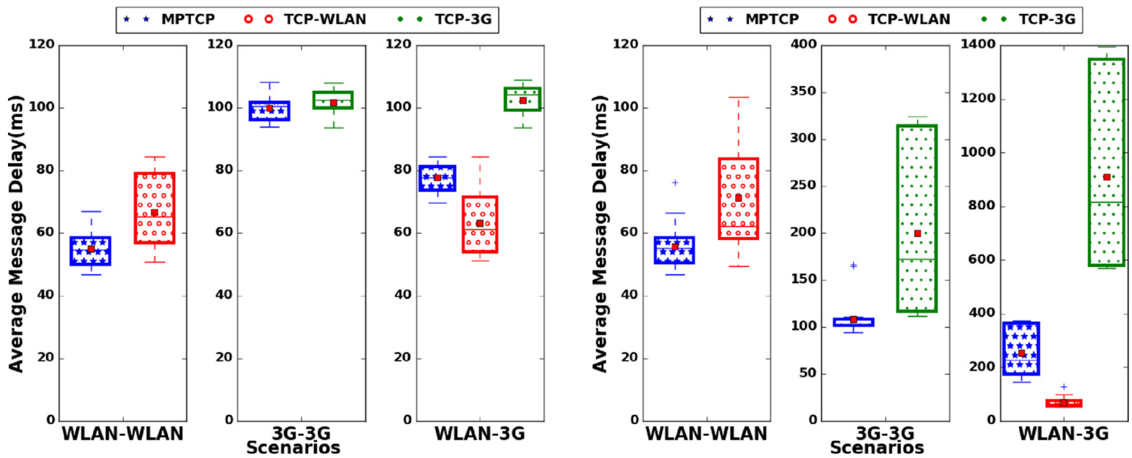

(a) Without Background Traffic (b) With Competing Background Traffic

Fig. 7. Average message delay for CBR video traffic over MPTCP in CORE emulation.

packet share, due to the fixed size of the packets. The large amount of data transmitted on the 3G link provokes head-of-line blocking and the resulting application-level latency prevents MPTCP from reducing the latency.

To analyse the performance in the case of competing traffic, we considered experiments with background flows as specified in Section 4.4. The results are presented in Fig. 7(b). Similar to the CMT-SCTP case, the background traffic has negligible impact over the WLAN paths as the loss encountered by the background flows prevents congestion from forming. Hence there is no impact of background traffic on the WLAN-WLAN scenario.

In the 3G-3G scenario, there is an improvement in the performance of MPTCP which was not visible without background traffic (or for CMT-SCTP). MPTCP has a less varying average message delay than TCP in this scenario mainly due to the use of multiple paths and the lowest-RTT scheduling: since the distribution of the traffic considers the delay of each path, it is affected by the current congestion level of each path. The MPTCP scheduler used both paths and the data distribution among the paths is 85/15% as shown in Table 5. The path with the shorter base RTT may not always be the best path as the background traffic builds up queues in the network, leading to a somewhat more even share of the data between the paths as compared to the scenario without background traffic.

In the WLAN-3G scenario, MPTCP increases the delay compared to single path TCP over WLAN as the data was split between asymmetric paths. The underlying issue is the same as when there is no background traffic, but the background traffic causes the delays to get larger. As in the CMT-SCTP scenario, in Fig. 7(b), the delays for MPTCP and for TCP over 3G is higher than those of the 3G-3G scenario, due to the larger receive buffers in the WLAN-3G scenario.

5.1.3. *MPTCP real measurements*

We have run over 30 experiments in the NorNet Edge (NNE) testbed for the video traffic and illustrated the delay measurements in Fig. 8. Note that different from WLAN links, in mobile broadband networks, each user has a dedicated channel. Therefore, we assume that the user is only streaming video without running any other bandwidth demanding applications.

For the homogeneous cases (e.g., WLAN-WLAN and 3G-3G), we observe that the paths can have quite different delay values although we are using the same technology. This results in delay differences between the paths, compared to the emulations, and impact the performance of MPTCP. For example, in the 3G-3G scenario, we observe that delay with MPTCP lies between the delay

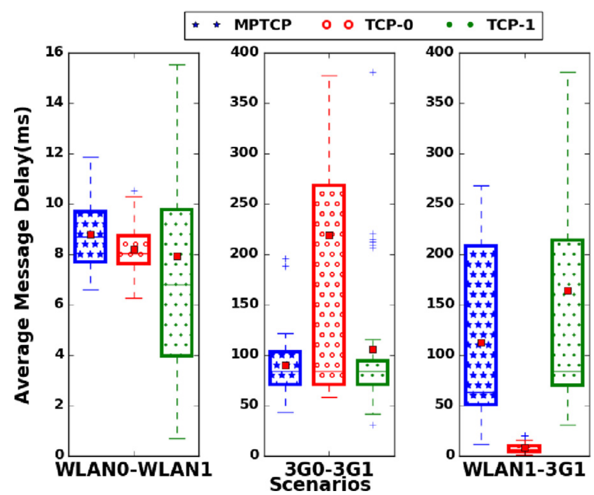

Fig. 8. Average message delay for CBR video traffic over MPTCP in NorNet experiment.

Table 6
Video traffic data share per path using MPTCP with NorNet.

	WLAN-WLAN		3G-3G		WLAN-3G	
	WLAN	WLAN	3G	3G	WLAN	3G
Video	65.05	34.9	93.8	6.19	20.41	79.5

with TCP of the two 3G paths and MPTCP provide delay values much closer to the 3G path with lower delay with TCP. Similarly, for the WLAN-WLAN scenario, the delay with MPTCP is on average closer to the delay of the WLAN path with the lower TCP delay. We further observe large variations in the delay values among different experiment runs. This is in fact an expected result in real networks where the channel conditions can be very dynamic. These observations can further be verified by looking into the video traffic data share tabulated in Table 6.

For the heterogeneous scenario (WLAN-3G), we observe that the delays achieved by MPTCP is higher than the delay of TCP on the WLAN path, since MPTCP occasionally uses 3G path that has higher delay values as compared to WLAN. Although we have higher variations in the experimental results, this behavior is in general consistent with the emulation results.

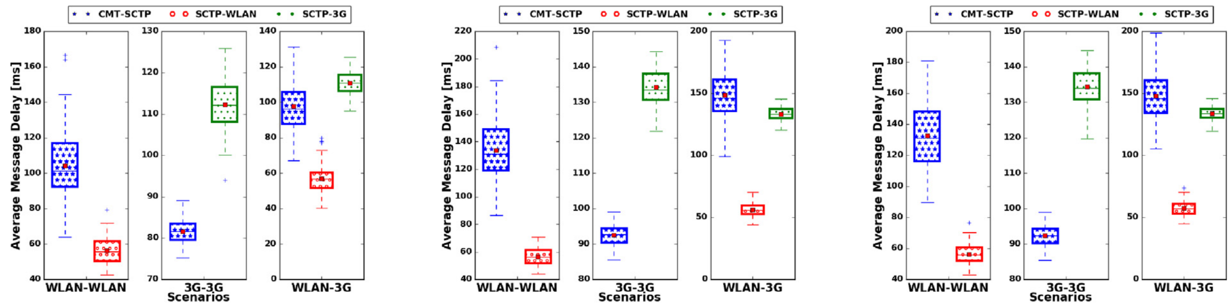

(a) Trace 1 (b) Trace 2 (c) Trace 3

Fig. 9. Average message delay for gaming traffic over CMT-SCTP (without background traffic).

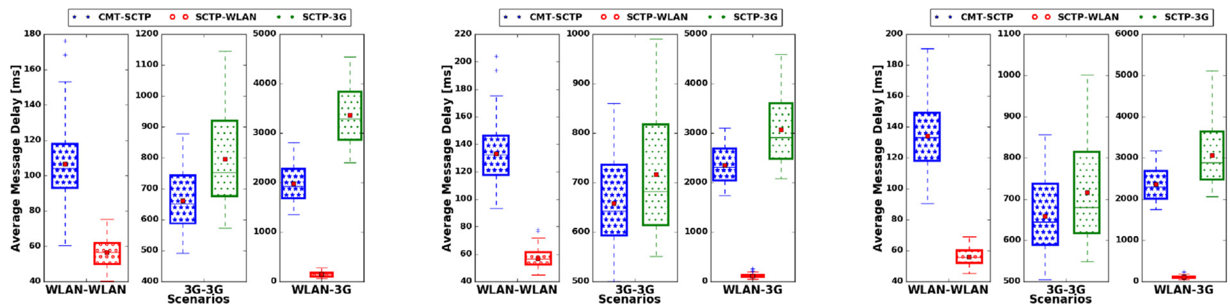

(a) Trace 1 (b) Trace 2 (c) Trace 3

Fig. 10. Average message delay for gaming traffic over CMT-SCTP (with competing background traffic).

Table 7
Path 1 Traffic share (in %) for gaming traffic over CMT-SCTP.

Trace	Background	WLAN-WLAN		3G-3G		WLAN-3G	
		WLAN	WLAN	3G	3G	WLAN	3G
Gaming T1	✗	24.1	75.9	30.9	69.1	33.7	66.3
	√	24.4	75.6	32.3	67.7	57.9	42.1
Gaming T2	✗	36.8	63.2	37.5	62.5	47.4	52.6
	√	36.9	63.1	23.9	76.1	63.9	36.1
Gaming T3	✗	36.4	63.6	37.3	62.7	48.4	51.6
	√	36.5	63.5	22.9	77.1	65.0	35.0

5.2. Gaming traffic

Gaming traffic is the second application traffic type that we consider. We start by presenting the simulation results for CMT-SCTP and continue with emulation and live experimentation results for MPTCP. We use the gaming traffic presented in Section 4.3.2. Same as for video traffic, we use application layer message delay as the performance metric.

5.2.1. CMT-SCTP

The average SCTP message delays over 128 runs are presented in Fig. 9 (without background traffic) and Fig. 10 (with background traffic according to Section 4.4) for the three gaming traces (Trace 1, Trace 2 and Trace 3) described in Section 4.3.2. The traffic share between the paths (i.e. the first path) is provided in Table 7.

A particular property of the gaming traffic is its mix of traffic patterns due to the different phases of game play (see Section 4.3.2). The CMT-SCTP scheduler provides no gain for the small packets sent during the game. Therefore, when using symmetric WLAN paths in the WLAN-WLAN scenario, CMT-SCTP only leads to additional delay caused by reordering for the occasional small bursts of packets sent.

For the 3G-3G scenario, as explained previously for the video traffic, the small send and receive buffer settings of 256 KiB keeps bufferbloat and packet reordering small. Therefore, the effort for reordering messages remains small as well. However, due to the higher network latency, the burst mitigation handling leads to some performance gain with CMT-SCTP; the non-saturated sender does not fully utilize its allowance given by the congestion window. Therefore, the congestion window is reduced by the burst mitigation. This limitation keeps the congestion window small, allowing CMT-SCTP to send messages more quickly on the two independent paths when small bursts of packets need to be sent.

Again, as explained in Section 5.1 there is high bufferbloat on the 3G path in the WLAN-3G case, due to the send/receive buffer sizing of 1024 KiB/2048 KiB. This leads to a significant reordering for CMT-SCTP during the initial large burst of the game. Therefore, no significant performance improvement is achieved when no background traffic is present. However, in the scenario with background traffic, CMT-SCTP leads to some improvement due to the distribution of traffic over two paths. Nevertheless, the latencies caused by the bufferbloat of at least 2 s make any gaming interactivity impossible.

5.2.2. MPTCP emulation

Fig. 11 shows the delay values calculated over 30 runs for each gaming trace and the distribution of the data on each path is shown in Table 8.

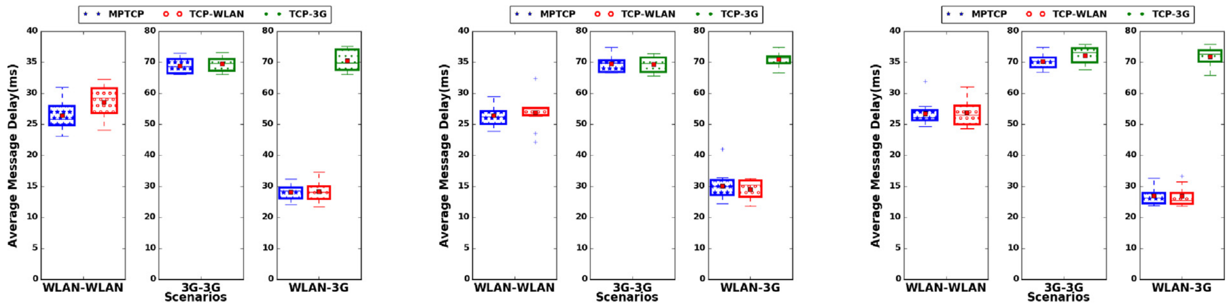

(a) Trace 1 (b) Trace 2 (c) Trace 3

Fig. 11. Average message delay for gaming traffic over MPTCP in CORE emulation (without background traffic).

Table 8
Gaming traffic data share per path using MPTCP.

Traffic	Background	WLAN-WLAN		3G-3G		WLAN-3G	
		WLAN	WLAN	3G	3G	WLAN	3G
Gaming T1	✗	64.88	35.12	80.85	19.15	94.03	5.97
	√	78.58	21.42	85.93	14.07	94.27	5.73
Gaming T2	✗	76.79	23.21	83.40	16.60	97.98	2.02
	√	67.65	32.35	93.42	6.58	99.23	0.77
Gaming T3	✗	78.30	21.70	92.01	7.99	99.29	0.71
	√	69.33	30.67	87.76	12.24	99.31	0.69

In the WLAN-WLAN scenario, it is clear that one path is more used than the other, which was not the case for the video traffic. This is due to the limited amount of data to transmit for the many small packets. Losses on the WLAN have minimal impact on performance as there is less data to send on the other path in the event of loss during large portions of the game. In the 3G-3G scenario, the average delay using MPTCP is similar to that of TCP. We also observe a split of data over the paths, which was not the case for the video traffic.

Though the scenarios WLAN-WLAN and 3G-3G are symmetric in nature based on the characteristics, one of the interfaces is the best in any given configuration. The path delay settings are randomly drawn from a range of values (see Table 1). The maximum possible difference between two path delays are 5ms in WLAN-WLAN and 10ms in 3G-3G, respectively. This difference is sufficient for the MPTCP scheduler to estimate, adapt and change outgoing path for a packet. If the default interface is the best of the two available interfaces, then the flow uses mostly this interface. However, when the default interface is not the best of the two available interfaces, the MPTCP scheduler will send the first few packets on the default interface before settling with the other interface. Due to the initial large burst and the long-tail nature of the gaming traffic, the first few packets that were sent on a possibly sub-optimal interface represents a large chunk of the data share, although most packets are transferred on the other interface. This also results in very similar performance for MPTCP and TCP.

It is also worth noting the difference to CMT-SCTP here, which showed a clear gain in the 3G-3G scenario. The difference for CMT-SCTP is that all paths are available for transmission immediately and it was able to spread the initial burst of data over both paths to reduce the delay. As MPTCP sets up the second subpath in parallel with starting the data transfer, and also uses a larger initial congestion window, there was no gain during the initial burst of data from the game setup.

In the asymmetric WLAN-3G scenario, the average MPTCP delay values are similar to the TCP delay of the WLAN path. This is due to most of the data being sent over the default (better) path, which in this case is the WLAN.

Table 9
Gaming traffic data share per path using MPTCP in NorNet.

Traffic	WLAN-WLAN		3G-3G		WLAN-3G	
	WLAN	WLAN	3G	3G	WLAN	3G
Gaming T1	77.67	22.32	0	100.0	40.74	59.25
Gaming T2	59.10	40.89	0	100.0	71.39	28.6
Gaming T3	54.57	45.42	0	100.0	30.31	69.69

Fig. 12 presents the average delay of MPTCP versus TCP for each gaming trace with competing background traffic. The impact of background traffic for gaming is very similar to the video traffic case. In the 3G-3G scenario, there is a lower and less variable average message delay using MPTCP, due to the scheduling. For the WLAN-3G scenario, MPTCP increases the delay compared to single path TCP over WLAN as some of the data is sent over the 3G path. Still, the degradation is smaller and the improvement in relation to TCP over 3G is larger than in the video traffic scenario. This as a smaller fraction of data is sent on the 3G path in this scenario as discussed above. Again, background traffic has no impact in the WLAN-WLAN scenario.

5.2.3. MPTCP real measurements

We illustrate gaming traffic delay performance for real-world measurements in Fig. 13 and the traffic distribution over the paths is presented in Table 9. For the homogeneous scenarios, we observe that the average TCP delay is very similar to the MPTCP delay. There are small variations among different traces, where in one trace MPTCP's delay is a little bit lower than the average TCP delay and in one trace it is a little bit higher. We observed that the packet shares are slightly more split towards the better path in the emulations compared to the experiments.

For WLAN-3G, we observe that MPTCP delay is slightly higher than the TCP WLAN delay. We observe that the packet share on the 3G path is higher in the experiments compared to the emulations, increasing the relative MPTCP delay slightly compared to the emulation results. Overall, the general trends seen between the protocols are still similar in the experiments and in the emulations.

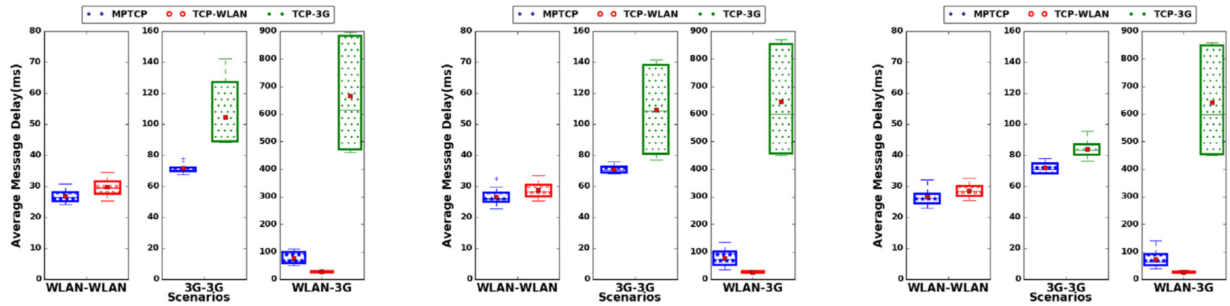

(a) Trace 1 (b) Trace 2 (c) Trace 3

Fig. 12. Average message delay for gaming traffic over MPTCP in CORE emulation (with competing background traffic).

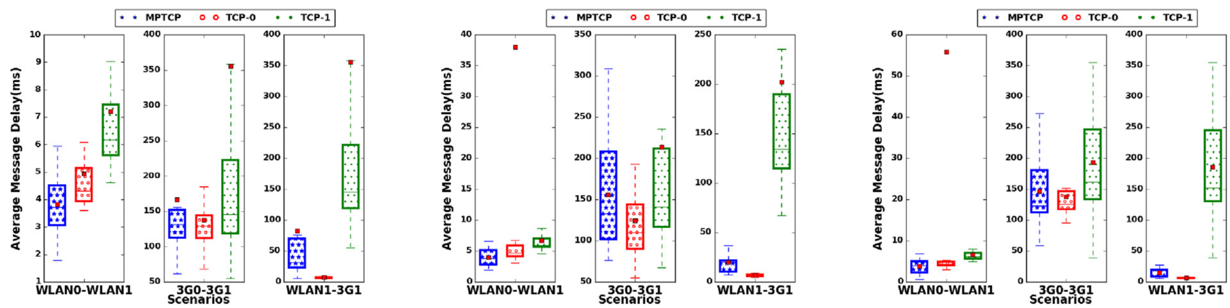

(a) Trace 1 (b) Trace 2 (c) Trace 3

Fig. 13. Average message delay for gaming traffic over MPTCP in NorNet experiment.

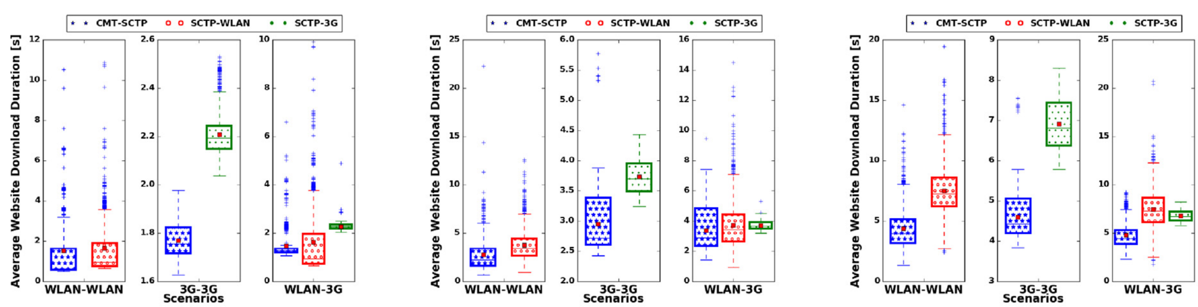

(a) Wikipedia (www.wikipedia.org) (b) Amazon (www.amazon.com) (c) Huffpost (www.huffingtonpost.com)

Fig. 14. Website download times over CMT-SCTP (without background traffic).

5.3. Web traffic

Finally, we evaluate the latency of web traffic for the homogeneous and heterogeneous scenarios, both with and without background traffic. For the web traffic, we chose the web site download time as the metric for transport protocol latency performance.

5.3.1. CMT-SCTP simulations

The web site download time results for the three web site scenarios are presented in Fig. 14 (without background traffic; 256 runs) and Fig. 15 (with background traffic according to Section 4.4; 1024 runs). Table 10 provides the corresponding traffic share for the two paths. Clearly, the benefit of CMT-SCTP usage increases as the web site size grows. The Wikipedia site (see Table 2), having only 72 KiB of payload data, is the smallest of the three sites. Therefore, the benefit of using CMT-SCTP for this web site is only small.

In the two scenarios with 3G path(s), a slight benefit can be seen: the 3G path has a small capacity and also a higher latency. Therefore, combining this 3G path with another 3G path, or even with a WLAN path, results in a faster download of the Wikipedia web site. As expected, for the Amazon (1 MiB) and the Huffington Post (3.9 MiB) web sites, CMT-SCTP reaches a significant download time reduction in most cases. However, for the WLAN-3G scenario with background traffic the path asymmetry is too large and CMT-SCTP performs worse than SCTP over WLAN for all web sites. Here the negative effect from head-of-line blocking dominates the gain from load balancing.

5.3.2. MPTCP emulation

The average web site download times over 30 runs for the three web site scenarios, without background traffic, are presented in Figure 16. The corresponding results with background traffic are presented in Fig. 17. Comparing the delay performance of MPTCP to that of TCP, MPTCP only provides limited improvements in download time for Wikipedia, but larger gains for both Amazon and Huffington Post (especially in 3G-3G scenarios). As also seen for CMT-SCTP, the results indicate that the size of the web site is critical to the total download time. With concurrent connections (6 in our setup), small web sites such as Wikipedia, can mostly be transferred within the initial window of TCP, not allowing MPTCP to exploit multiple paths.

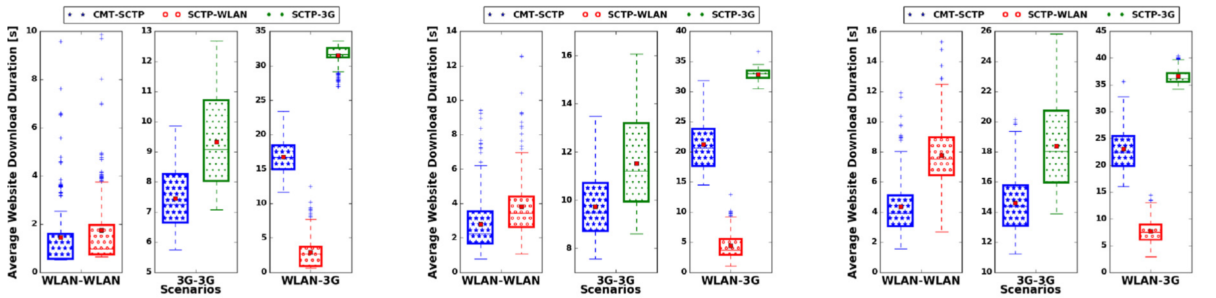

(a) Wikipedia (www.wikipedia.org) (b) Amazon (www.amazon.com) (c) Huffpost (www.huffingtonpost.com)

Fig. 15. Website download times over CMT-SCTP (with competing background traffic).

Table 10
Path 1 Traffic share (in %) for website download over CMT-SCTP

Website	Background	WLAN-WLAN		3G-3G		WLAN-3G	
		WLAN	WLAN	3G	3G	WLAN	3G
Wikipedia	✗	23.5	76.5	21.7	78.3	24.5	75.5
	√	23.7	76.3	23.3	76.7	33.0	67.0
Amazon	✗	24.6	75.4	23.6	76.4	31.2	68.8
	√	24.3	75.7	24.1	75.9	44.2	54.8
Huffington Post	✗	23.9	76.1	23.2	76.8	31.9	68.1
	√	23.9	76.1	24.6	76.4	50.5	49.5

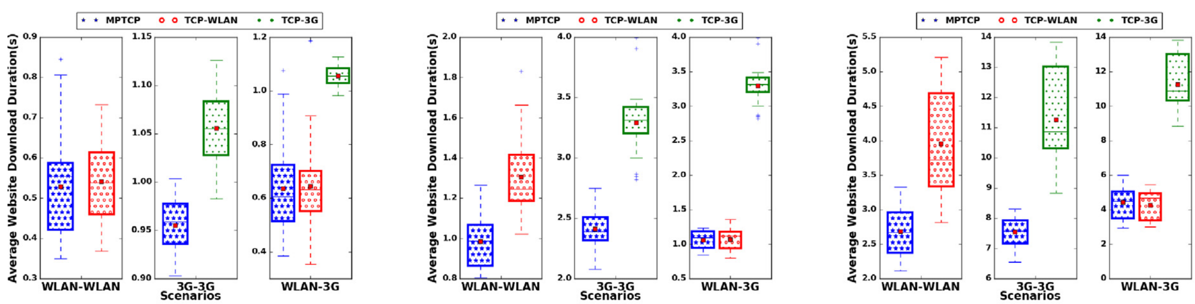

(a) Wikipedia (www.wikipedia.org) (b) Amazon (www.amazon.com) (c) Huffpost (www.huffingtonpost.com)

Fig. 16. Website download times over MPTCP in CORE emulation (without background traffic).

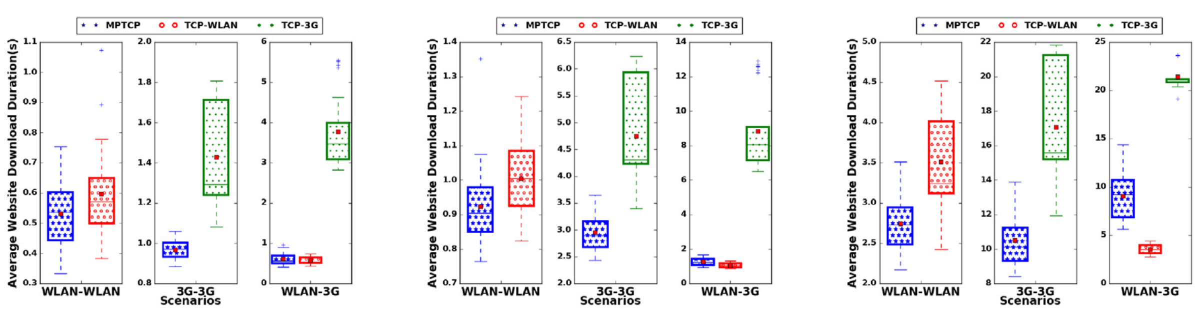

(a) Wikipedia (www.wikipedia.org) (b) Amazon (www.amazon.com) (c) Huffpost (www.huffingtonpost.com)

Fig. 17. Website download times over MPTCP in CORE emulation (with Competing background traffic).

With background traffic, the performance trends are similar to that of the non-background case. For the 3G paths, background traffic significantly increases the download time as well as the variation in download time. Background traffic has very little impact over the WLAN paths, but the random loss over the WLAN links still leads to a large variation in download times. The position of the loss in the short web flows can have a huge impact on the total web site download duration.

Table 11 indicates that in symmetric scenarios, data transfer uses both paths. In WLAN-WLAN, random losses causes data to be sent over the second path even in small web sites like Wikipedia. The delay variance on the paths is the primary reason for use of the second path in 3G-3G scenarios. In asymmetric scenarios, most of the data uses the primary faster path WLAN.

5.3.3. MPTCP real measurements

In Fig. 18, we illustrated the results of download times for web traffic. For WLAN-WLAN, we observe that the MPTCP delay is lower than the average TCP delay for Amazon whereas it is higher than the average TCP delay for the other two web sites: Wikipedia and

Table 11
Web traffic data share per path using MPTCP.

	Background	WLAN-WLAN		3G-3G		WLAN-3G	
		WLAN	WLAN	3G	3G	WLAN	3G
Wikipedia	✗	90.22	9.78	86.98	13.02	99.94	0.06
	√	86.98	13.02	83.86	16.14	99.94	0.06
Amazon	✗	76.02	23.98	56.73	43.27	99.81	0.19
	√	70.45	29.55	72.64	27.36	99.9	0.1
Huffington Post	✗	85.3	14.7	68.1	31.9	99.9	0.1
	√	83.0	17	79.13	20.87	98.86	1.14

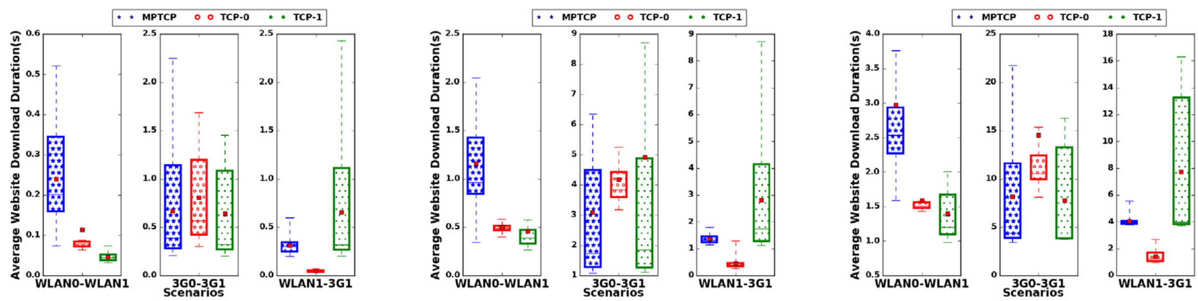

(a) Wikipedia (www.wikipedia.org) (b) Amazon (www.amazon.com) (c) Huffpost (www.huffingtonpost.com)

Fig. 18. Website download times over MPTCP in NorNet experiment.

Table 12
Web traffic data share per path using MPTCP in NorNet.

	WLAN-WLAN		3G-3G		WLAN-3G	
	WLAN	WLAN	3G	3G	WLAN	3G
Wikipedia	94.3	5.6	100.0	0.0	100.0	0.0
Amazon	26.1	73.8	96.9	3.1	100.0	0.0
Huffington Post	14.5	85.4	100.0	0.0	100.0	0.0

Huffington Post. Similar observations can be made for the 3G-3G scenario. We observe that when there is enough data to be transmitted, MPTCP can provide benefits. For the WLAN-3G case, similar to other traffic, we observe that for all the web sites, MPTCP delay is a little bit higher than the WLAN delay.

The traffic distribution of data over the different paths is shown in Table 12. We observed that for the heterogeneous cases, almost all traffic is transferred over the WLAN path whereas for the homogeneous cases, the distribution depends on the size of the web site.

6. Discussion of results

In this paper, we have run extensive measurements to evaluate the capability of multi-path transport protocols to carry latency sensitive application traffic. More specifically, we have analysed the application delays for video traffic, online gaming and web services both with and without competing traffic. Furthermore, we considered end-hosts experiencing multiple homogeneous paths as well as heterogeneous ones. The results are summarized in Fig. 19, where the performance of multi-path, as compared to single path, is categorized into four types. We next elaborate on this table and recap our findings from Section 5.

The currently used round-robin scheduler in CMT-SCTP is optimised for throughput, not for low latency. Therefore, when there is a significant delay difference between the two paths, we observe performance degradation with CMT-SCTP compared to SCTP. However, for homogeneous paths, especially when the paths are not very lossy, we observe that CMT-SCTP can significantly reduce latency, especially in the web scenario. For video traffic CMT-SCTP provides similar or higher delay values as compared to SCTP. For example, in WLAN-WLAN scenarios, reordering due to distribution over multiple lossy paths increases the delay both with and without background traffic. On the other hand, in 3G-3G scenarios, we observe similar delay values for CMT-SCTP and SCTP. For the WLAN-3G scenario, packet reordering causes increased delay for CMT-SCTP even when there is no competing traffic. For gaming traffic, while CMT-SCTP leads to a latency increase in case of two similar, low-latency WLAN paths, it becomes beneficial in case of high-latency paths with background traffic; in comparison of using only the higher-delay path, CMT-SCTP is able to take advantage of the lower-delay path to reduce latency. However, as observed for the video traffic, its scheduler is optimised for throughput maximization without taking care of path delay. Therefore, while the latency is lower than using only the high-delay path, it is still much higher than using only the low-delay path in the heterogeneous WLAN-3G case. For web traffic, we observe that using CMT-SCTP improves the web site download speed, especially for homogeneous paths. Since the web traffic is saturated (i.e., send as much data as possible), the round-robin scheduler that is used by the CMT-SCTP implementation performs reasonably well by ensuring that both paths are utilized, although it does not always chose the path with the lowest RTT. The larger the web site, the better the performance improvement achieved by CMT-SCTP. However, for the WLAN-3G case, especially when there is background traffic, CMT-SCTP cannot handle the delay difference between the 3G and WLAN paths, resulting in poor performance for CMT-SCTP compared to SCTP.

For MPTCP, the default scheduler is based on delay and it has been shown to achieve low and stable latency [69]. For different type of traffics, we observe similar or lower delay values for MPTCP compared to TCP, for homogeneous paths. However, the main factor that determines the delay performance of MPTCP is indeed the path heterogeneity, and for heterogeneous paths we observe performance degradation whose degree depends on the traffic and whether there exists background traffic. More specifically, for video traffic, when the links are lossy, e.g. the WLAN-WLAN case, we observe delay gains due to link aggregation. For the 3G-3G case, where there are no losses, MPTCP selects the best available path

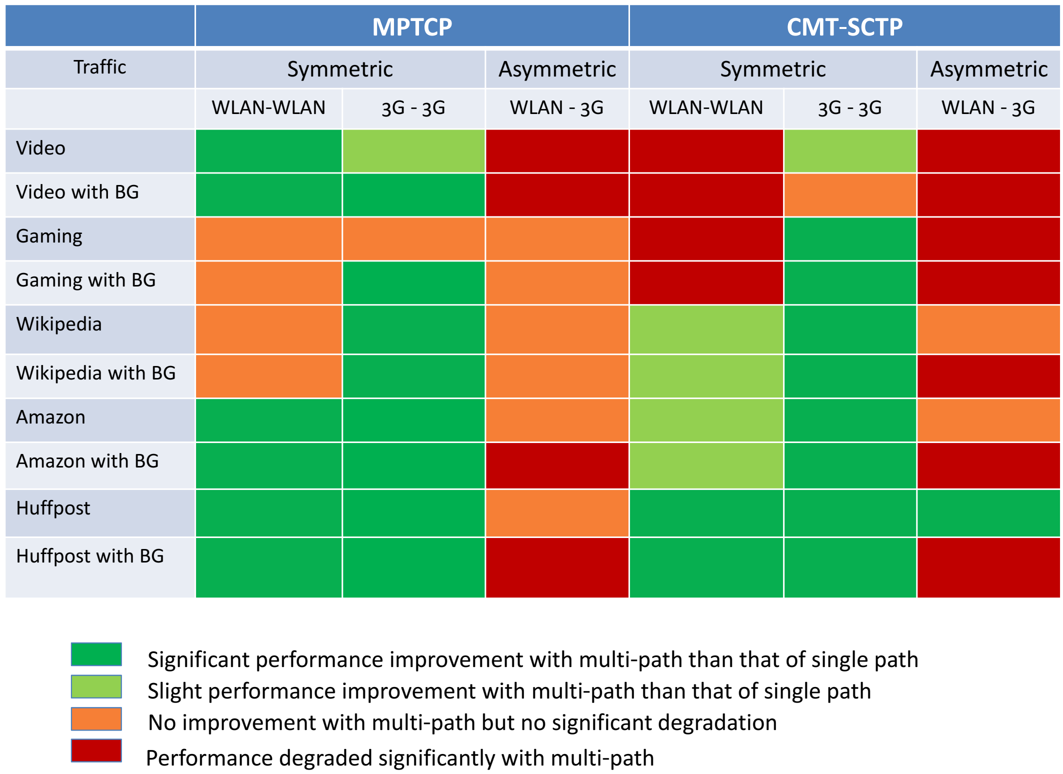

Fig. 19. Multi-path versus single path transport protocols depending on the latency sensitive traffic: Summary table.

resulting in minor gains compared to TCP. In the presence of background traffic, the delay values are in general higher but the benefits of using MPTCP are consistent with that of the non-competing traffic. We further observe that these emulation results are mostly consistent with the NorNet experiments. Here, the main difference is that in real networks, paths have more diverse characteristics, although using the same technology. This results in slight heterogeneity and the delay using MPTCP becomes higher than using TCP over the best path only, but still lower than the average TCP delay. For gaming traffic, we observe very similar delay values to TCP for almost all cases due to the very limited amount of data. The background traffic did not induce enough loss in the foreground flows, therefore the delay values are similar to that of no background traffic. One exception is the 3G-3G scenario where MPTCP keeps sending over one path as long as there is no loss, therefore, providing some delay gains. For the WLAN-3G case, MPTCP uses the WLAN at almost all times, therefore, the MPTCP delay is similar to the TCP delay of WLAN. Similarly, for the results of the real experiments, we observe similar behavior to the emulations, especially for the homogeneous scenarios, with slight variations among different trace files. For the heterogeneous scenario, there is slightly higher delay in real experiments compared to emulations, due to some traffic is being transferred over the slower 3G path. For the web traffic, we observe that MPTCP provides lower delay values, especially for web sites with many objects. The lower delay is a consequence of MPTCP's scheduler which always tries to use the path with the lowest RTT. However, when the paths are very heterogeneous in terms of delay and loss as in the WLAN-3G case, losses in WLAN forces MPTCP to use the suboptimal 3G path; therefore, the MPTCP delay becomes higher than the TCP delay of WLAN. These results hold for emulations with and without background traffic. Similar to the previous applications, the results of the real experiments are mostly consistent with the emulation results. Due to the differences in the paths for the homogeneous cases (e.g. 3G-3G and WLAN-WLAN), MPTCP delay is higher than the best path while still much lower than the average TCP delay.

One conclusion of our study is that multi-path transport protocols can hardly reduce the latency for all the tested applications, when there is some asymmetry between the paths. Moreover, in this case, multi-path transport may increase the latency, mainly because of head-of-line blocking. However, it is worth pointing out that in most symmetric scenarios, multi-path transport protocols enable a significant latency reduction.

7. Related work

This section discusses work related to ours. While there are numerous articles on MPTCP and its performance in relation to TCP, not much has been written on the relation between CMT-SCTP and SCTP. Instead, most articles on CMT-SCTP propose various optimizations to the protocol itself. There are, of course exceptions; Aydin et al. [70] elaborates the importance of TCP friendliness for single homed SCTP and evaluates the TCP friendliness of CMT-SCTP. Arianpoo et al. [71] propose an adaptive network coding mechanism for CMT-SCTP to desensitize the receiver against packet reordering and consequently eliminate the receiver buffer blocking problem. They claim to have improve the CMT-SCTP performance by 62% over the original implementation in cases of severe path asymmetry.

For MPTCP, a closely related research work is [72], which measures MPTCP performance with the aim to understand the benefit of using two interfaces with MPTCP over using either one of the interface with TCP. This study also focuses on the impact that flow size has on the average latency, and provides insights into the effect of path characteristic diversity on application level performance. Their conclusions are consistent with ours: using multi-path becomes more and more beneficial when the size of the data to transmit increases. We extend their work by considering more application scenarios. In [73], S. Deng et al., studies the performance of MPTCP over wireless technologies using Android application traffic. Their study focuses on energy efficiency and provides new challenges such as dynamic decision making at the mobile applications to select appropriate network technology depending on the flow size and traffic pattern. Handover performance was seen as a potential MPTCP performance impairment especially when the path characteristics are different. Andrei et al. [74] provided a simultaneous association solution using MPTCP for WLAN that avoids fast handover. It also provides possible modifications at the client side implementation, to mitigate the throughput loss in cases where the WLAN characteristics differ due to channel specification. Such approach of reducing the occurrence of handovers is a necessary improvement: it is essential to improve performance where multi-path transport increases the latency, and that has been identified in this paper.

Grinnemo et al. [75] provides a first comprehensive evaluation of MPTCP performance with latency as the quality of experience metric for cloud-based applications. They study three different applications: Netflix, Google Maps and Google Docs, representing high, mid and low intensity cloud-based traffic. The authors conclude that MPTCP provides significant performance gains for high and mid intensity traffic. Furthermore, it is noted that the variation in RTTs among network paths causes higher application latency, and the current Linux standard scheduler is seen as the primary cause of increased latency in such cases.

Raiciu et al. [76] proposed a mobility architecture to allow MPTCP to switch between different technologies and handle mobility at the transport layer instead of at the network layer. The mobility of MPTCP was evaluated with simulations and indoor mobility experiments. The criteria for the evaluation was measured throughput on TCP and MPTCP using WLAN-3G, and power efficiency of both protocols. The study concludes that MPTCP provides performance improvements over TCP when multiple interfaces are used in parallel. Power efficiency of MPTCP depends on the underlying interface power consumption and should be tuned for better performance. Later, the power efficiency of MPTCP drew much attention in [77], which analyses the energy consumption and handover performance of MPTCP in the different operational modes: Full MPTCP Mode, Backup Mode and Single path Mode. This work again provides experimental evaluations using the Linux implementation of MPTCP and commercial access networks providing 3G and broadband access on static nodes. The study concludes that MPTCP handovers might have small impact on application delay and goodput in different operational modes.

With a few exceptions discussed above, most of the prior research on CMT-SCTP or MPTCP measurements focused on the performance of the protocol in terms of throughput, energy consumption, handover performance and RTTs. To the best of our knowledge, this paper is the first to provide a comprehensive analysis of multi-path transport performance with latency as the main metric.

As seen throughout this paper, the performance of multi-path transport is highly dependent on doing efficient scheduling. Paasch et al. [27] provide a detailed study of schedulers and their impact on performance. The authors implement a generic modular framework for evaluating MPTCP schedulers in Linux. Using this framework, different schedulers are then evaluated using various performance metrics and different types of traffic, including bulk and application limited traffic.

8. Conclusions and future work

For an increasing number of applications, latency plays an important role as it directly impacts their performance. Still, most work considering multi-path communication is solely focused on resilience and throughput maximization. The work presented in this paper tries to bridge this gap by evaluating whether multi-path communication can help latency-sensitive applications satisfy their users' requirements. Three latency-sensitive applications have been considered: video, gaming and web traffic. Performance have been evaluated using 3G–3G, 3G–WLAN, and WLAN–WLAN paths, in both simulated, emulated and real-life environments considering both CMT-SCTP and MPTCP.

The results indicate that multi-path communication can reduce latency significantly, but only when paths are symmetric in terms of delay and loss rate. The potential gain comes mainly from two factors: the possibility to distribute short bursts of data over multiple interfaces and the ability to select the best of the available paths for data transmission. In asymmetric scenarios where the latency reduction is not as significant (or non-existent), applications may still benefit from other properties of multi-path communication, without increasing latency. This is, however, highly dependent on the scheduling mechanism used. As seen in some of the CMT-SCTP experiments, a scheduler designed mainly for throughput maximization, may lead to increased latency in some scenarios. Considering the importance of scheduling, this is where we direct our attention for future work, and we are currently designing a scheduler targeting latency-sensitive traffic.

Acknowledgments

The authors are funded by the European Community under its Seventh Framework Programme through the Reducing Internet Transport Latency (RITE) project (ICT-317700). The views expressed are solely those of the authors.

References

[1] S.S. Krishnan, R.K. Sitaraman, Video stream quality impacts viewer behavior: inferring causality using quasi-experimental designs, in: Proceedings of the 12th ACM Internet Measurement Conference (IMC), ACM, New York/U.S.A., 2012, pp. 211–224, doi:10.1145/2398776.2398799.
[2] P. Quax, P. Monsieurs, W. Lamotte, D.D. Vleeschauwer, N. Degrande, Objective and subjective evaluation of the influence of small amounts of delay and jitter on a recent first person shooter game, in: Proceedings of 3rd ACM SIGCOMM Workshop on Network and System Support for Games (NetGames), ACM, Portland, Oregon/U.S.A., 2004, pp. 152–156, doi:10.1145/1016540.1016557. ISBN 1-58113-942-X.
[3] J. Liu, Y. Tan, F. Fu, T. Dreibholz, X. Zhou, Y. Bai, X. Yang, W. Du, Study on MPTCP and CMT-SCTP congestion control mechanism, Comput. Eng. 41 (4) (2015) 117–124, doi:10.3969/j.issn.1000-3428.2015.04.022.
[4] A. Ford, C. Raiciu, M. Handley, O. Bonaventure, TCP extensions for multipath operation with multiple addresses, RFC 6824, IETF, 2013, doi:10.17487/RFC6824. ISSN 2070-1721.
[5] J.B. Postel, Transmission control protocol, Standards Track RFC 793, IETF, 1981. ISSN 2070-1721.
[6] J.R. Iyengar, P.D. Amer, R.R. Stewart, Concurrent multipath transfer using SCTP multihoming over independent end-to-end paths, IEEE/ACM Trans. Netw. 14 (5) (2006) 951–964, doi:10.1109/TNET.2006.882843. ISSN 1063-6692.
[7] P.D. Amer, M. Becke, T. Dreibholz, N. Ekiz, J.R. Iyengar, P. Natarajan, R.R. Stewart, M. Tüxen, Load sharing for the stream control transmission protocol (SCTP), Internet Draft draft-tuexen-tsvwg-sctp-multipath-11, IETF, Individual Submission, 2015.
[8] T. Dreibholz, Evaluation and optimisation of multi-path transport using the stream control transmission protocol, University of Duisburg-Essen, Faculty of Economics, Institute for Computer Science and Business Information Systems, 2012 Habilitation treatise.
[9] R.R. Stewart, Stream control transmission protocol, RFC 4960, IETF, 2007, doi:10.17487/RFC4960. ISSN 2070-1721.
[10] T. Dreibholz, I. Rüngeler, R. Seggelmann, M. Tüxen, E.P. Rathgeb, R.R. Stewart, Stream control transmission protocol: past, current, and future standardization activities, IEEE Commun. Mag. 49 (4) (2011) 82–88, doi:10.1109/MCOM.2011.5741151. ISSN 0163-6804.
[11] IETF SIGTRAN WG, IETF signaling transport working group, 2012.
[12] ITU-T, Introduction to CCITT signalling system no. 7, Recommendation Q.700, International Telecommunication Union, 1994.
[13] J.B. Postel, Internet protocol, Standards Track RFC 791, IETF, 1981. ISSN 2070-1721.
[14] J.B. Postel, User Datagram Protocol, Standards Track RFC 768, IETF, 1980. ISSN 2070-1721.

[15] R.R. Stewart, Q. Xie, M. Tüxen, S. Maruyama, M. Kozuka, Stream control transmission protocol (SCTP) dynamic address reconfiguration, Standards Track RFC 5061, IETF, 2007, doi:10.17487/RFC5061. ISSN 2070-1721.
[16] L. Budzisz, J. Garcia, A. Brunstròm, R. Ferrús, A taxonomy and survey of SCTP research, ACM Comput. Surveys (CSUR) 44 (4) (2012) 18:1–18:36, doi:10.1145/2333112.2333113.
[17] S. Barré, C. Paasch, O. Bonaventure, MultiPath TCP: from theory to practice, in: Proceedings of the 10th International IFIP Networking Conference, Valencia/Spain, 2011, pp. 444–457. http://dx.doi.org/10.1109/JPROC.2010.2093850. ISBN 978-3-642-20756-3.
[18] M. Allman, V. Paxson, E. Blanton, TCP congestion control, Standards Track RFC 5681, IETF, 2009, doi:10.17487/RFC5681. ISSN 2070-1721.
[19] IETF MPTCP WG, IETF multipath TCP working croup, 2012.
[20] M. Becke, T. Dreibholz, H. Adhari, E.P. Rathgeb, On the fairness of transport protocols in a multi-path environment, in: Proceedings of the IEEE International Conference on Communications (ICC), Ottawa, Ontario/Canada, 2012, pp. 2666–2672, doi:10.1109/ICC.2012.6363695. ISBN 978-1-4577-2052-9.
[21] C. Raiciu, M. Handley, D. Wischik, Coupled congestion control for multipath transport protocols, RFC 6356, IETF, 2011a, doi:10.17487/RFC6356. ISSN 2070-1721.
[22] C. Raiciu, S. Barré, C. Pluntke, A. Greenhalgh, D. Wischik, M. Handley, Improving datacenter performance and robustness with multipath TCP, in: Proceedings of the ACM SIGCOMM Conference, Toronto, Ontario/Canada, 2011b, pp. 266–277, doi:10.1145/2018436.2018467. ISBN 978-1-4503-0797-0.
[23] L. Boccassi, M. Fayed, M. Marina, Binder: a system to aggregate multiple internet gateways in community networks, in: Proceedings of the ACM MobiCom Workshop on Lowest Cost Denominator Networking for Universal Access (LCDNet), Miami, FL, USA, 2013, pp. 3–8.
[24] H. Adhari, T. Dreibholz, M. Becke, E.P. Rathgeb, M. Tüxen, Evaluation of concurrent multipath transfer over dissimilar paths, in: Proceedings of the 1st International Workshop on Protocols and Applications with Multi-Homing Support (PAMS), Singapore, 2011, pp. 708–714, doi:10.1109/WAINA.2011.92. ISBN 978-0-7695-4338-3.
[25] T. Dreibholz, M. Becke, E.P. Rathgeb, M. Tüxen, On the use of concurrent multipath transfer over asymmetric paths, in: Proceedings of the IEEE Global Communications Conference (GLOBECOM), Miami, Florida/U.S.A., 2010, pp. 1–6, doi:10.1109/GLOCOM.2010.5683579. ISBN 978-1-4244-5637-6.
[26] N. Kuhn, E. Lochin, A. Mifdaoui, G. Sarwar, O. Mehani, R. Boreli, DAPS: intelligent delay-aware packet scheduling for multipath transport, in: Proceedings of the IEEE International Conference on Communications (ICC), Sydney, New South Wales/Australia, 2014, pp. 1222–1227.
[27] C. Paasch, S. Ferlin, Özgü Alay, O. Bonaventure, Experimental evaluation of multipath TCP schedulers, in: Proceedings of the ACM SIGCOMM Capacity Sharing Workshop (CSWS), Chicago, Illinois/U.S.A., 2014, pp. 27–32, doi:10.1145/2630088.2631977. ISBN 978-1-4503-2991-0.
[28] C. Raiciu, C. Paasch, S. Barré, A. Ford, M. Honda, F. Duchêne, O. Bonaventure, M. Handley, How hard can it be? Designing and implementing a deployable multipath TCP, in: Proceedings of the 9th USENIX Conference on Networked Systems Design and Implementation (NSDI), San Jose, California/U.S.A., 2012, pp. 1–14.
[29] T. Dreibholz, M. Becke, J. Pulinthanath, E.P. Rathgeb, Applying TCP-friendly congestion control to concurrent multipath transfer, in: Proceedings of the 24th IEEE International Conference on Advanced Information Networking and Applications (AINA), Perth, Western Australia/Australia, 2010, pp. 312–319, doi:10.1109/AINA.2010.117. ISBN 978-0-7695-4018-4.
[30] T. Dreibholz, M. Becke, H. Adhari, E.P. Rathgeb, On the impact of congestion control for concurrent multipath transfer on the transport layer, in: Proceedings of the 11th IEEE International Conference on Telecommunications (ConTEL), Graz, Steiermark/Austria, 2011, pp. 397–404. ISBN 978-953-184-152-8.
[31] R. Khalili, N. Gast, M. Popović, J.-Y.L. Boudec, MPTCP is not pareto-optimal: performance issues and a possible solution, IEEE/ACM Trans. Netw. 21 (5) (2013) 1651–1665, doi:10.1109/TNET.2013.2274462. ISSN 1063-6692.
[32] A. Walid, Q. Peng, J. Hwang, S.H. Low, Balanced linked adaptation congestion control algorithm for MPTCP, Internet Draft draft-walid-mptcp-congestion-control-02, IETF, Individual Submission, 2015.
[33] J.R. Iyengar, P.D. Amer, R.R. Stewart, Retransmission policies for concurrent multipath transfer using SCTP multihoming, in: Proceedings of the 12th IEEE International Conference on Networks (ICON), 2, 2004, pp. 713–719, doi:10.1109/ICON.2004.1409269.
[34] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, F. Jahanian, M. Karir, Atlas internet observatory 2009 annual report, 47th NANOG, 2009.
[35] D. Ciullo, M. Mellia, M. Meo, Two schemes to reduce latency in short lived TCP flows, IEEE Commun. Lett. 13 (10) (2009) 806–808.
[36] Sandvine intelligent broadband networks, Global internet phenomena report, 2013.
[37] Cisco visual networking index, Forecast and methodology, 2012–2017, 2013.
[38] Sandvine Intelligent Broadband Networks, Global internet phenomena report, 2015.
[39] X. Zhang, Y. Xu, H. Hu, Y. Liu, Z. Guo, Y. Wang, Profiling skype video calls: rate control and video quality, in: Proceedings of the IEEE INFOCOM, Orlando, Florida/U.S.A., 2012, pp. 621–629, doi:10.1109/INFCOM.2012.6195805.
[40] L.D. Cicco, S. Mascolo, V. Palmisano, Skype video congestion control: an experimental investigation, Computer Netw. 55 (3) (2011) 558–571. http://dx.doi.org/10.1016/j.comnet.2010.09.010.
[41] T. Szigeti, C. Hattingh, End-to-End QoS Network Design: Quality of Service in LANs, WANs, and VPNs Quality of Service Design Overview, Cisco Press, 2004. ISBN 978-1587051760.
[42] M. Claypool, K. Claypool, Latency and player actions in online games, Commun. ACM 49 (11) (2006) 40–45, doi:10.1145/1167838.1167860.
[43] X. Che, B. Ip, Review: packet-level traffic analysis of online games from the genre characteristics perspective, J. Netw. Comput. Appl. 35 (1) (2012) 240–252, doi:10.1016/j.jnca.2011.08.005.
[44] N. Sheldon, E. Girard, S. Borg, M. Claypool, E. Agu, The effect of latency on user performance in Warcraft III, in: Proceedings of the 2nd Workshop on Network and System Support for Games (NetGames), Redwood City, California/U.S.A., 2003, pp. 3–14.
[45] K.-T. Chen, P. Huang, G.-S. Wang, C.-Y. Huang, C.-L. Lei, On the sensitivity of online game playing time to network QoS, in: Proceedings of the 25th IEEE International Conference on Computer Communications (INFOCOM), Barcelona, Catalonia/Spain, 2006, pp. 1–12, doi:10.1109/INFOCOM.2006.286. ISSN 0743-166X.
[46] O3b networks and sofrecom, Why latency matters to mobile backhaul, 2013.
[47] Google developers, Web metrics: size and number of resources, 2010.
[48] A. Varga, OMNeT++ discrete event simulation system user manual – version 4.6, 2014.
[49] T. Dreibholz, M. Becke, J. Pulinthanath, E.P. Rathgeb, Implementation and evaluation of concurrent multipath transfer for SCTP in the INET framework, in: Proceedings of the 3rd ACM/ICST International Workshop on OMNeT++, Torremolinos, Málaga/Spain, 2010, pp. 15–22, doi:10.4108/ICST.SIMUTOOLS2010.8673. ISBN 978-963-9799-87-5.
[50] I. Rüngeler, SCTP – evaluating, improving and extending the protocol for broader deployment, University of Duisburg-Essen, Faculty of Economics, Institute for Computer Science and Business Information Systems, 2009 Ph.D. thesis.
[51] T. Dreibholz, M. Becke, H. Adhari, E.P. Rathgeb, Evaluation of a new multipath congestion control scheme using the NetPerfMeter tool-chain, in: Proceedings of the 19th IEEE International Conference on Software, Telecommunications and Computer Networks (SoftCOM), Hvar/Croatia, 2011, pp. 1–6. ISBN 978-953-290-027-9.
[52] A. Varga, INET framework for OMNeT++, 2015.
[53] T. Dreibholz, X. Zhou, E.P. Rathgeb, SimProcTC – the design and realization of a powerful tool-chain for OMNeT++ simulations, in: Proceedings of the 2nd ACM/ICST International Workshop on OMNeT++, Rome/Italy, 2009, pp. 1–8, doi:10.4108/ICST.SIMUTOOLS2009.5517. ISBN 978-963-9799-45-5.
[54] T. Dreibholz, E.P. Rathgeb, A powerful tool-chain for setup, distributed processing, analysis and debugging of OMNeT++ simulations, in: Proceedings of the 1st ACM/ICST International Workshop on OMNeT++, Marseille, Bouches-du-Rhône/France, 2008, pp. 74–81, doi:10.4108/ICST.SIMUTOOLS2008.2990. ISBN 978-963-9799-20-2.
[55] K.V. Jónsson, HttpTools: a toolkit for simulation of web hosts in OMNeT++, in: Proceedings of the 2nd ACM/ICST International Workshop on OMNeT++, Rome/Italy, 2009, pp. 1–8, doi:10.4108/ICST.SIMUTOOLS2009.5589. ISBN 978-963-9799-45-5.
[56] J. Ahrenholz, Comparison of CORE network emulation platforms, in: Military Communications Conference (MILCOM), San Jose, California/U.S.A., 2010, pp. 166–171, doi:10.1109/MILCOM.2010.5680218.
[57] A. Kvalbein, D. Baltrūnas, K.R. Evensen, J. Xiang, A.M. Elmokashfi, S. Ferlin, The NorNet edge platform for mobile broadband measurements, Comput. Netw. Special Issue on Future Internet Testbeds 61 (2014) 88–101, doi:10.1016/j.bjp.2013.12.036. ISSN 1389-1286.
[58] T. Dreibholz, S. Ferlin, O. Alay, N. Kuhn, K. Yedugundla, P. Hurtig, A. Brunstrom, Is multi-path transport suitable for latency sensitive traffic, 2015.
[59] J.C. Mogul, G. Minshall, Rethinking the tcp nagle algorithm, SIGCOMM Comput. Commun. Rev. 31 (1) (2001) 6–20, doi:10.1145/382176.382177.
[60] M. Allman, E. Blanton, Notes on burst mitigation for transport protocols, ACM SIGCOMM Comput. Commun. Rev. (CCR) 35 (2) (2005) 53–60, doi:10.1145/1064413.1064419. ISSN 0146-4833.
[61] Funcom, Age of Conan, 2015.
[62] A. Botta, A. Dainotti, A. Pescapé, A tool for the generation of realistic network workload for emerging networking scenarios, Comput. Netw. 56 (15) (2012) 3531–3547, doi:10.1016/j.comnet.2012.02.019.
[63] S. Ha, L. Le, I. Rhee, L. Xu, Impact of background traffic on performance of high-speed TCP variant protocols, Comput. Netw. 51 (7) (2007) 1748–1762, doi:10.1016/j.comnet.2006.11.005.
[64] T. Dreibholz, NetPerfMeter homepage, 2015.
[65] S. Hassayoun, J.R. Iyengar, D. Ros, Dynamic window coupling for multipath congestion control, in: Proceedings of the 19th IEEE International Conference on Network Protocols (ICNP), 2011, pp. 341–352, doi:10.1109/ICNP.2011.6089073.
[66] V. Cerf, V. Jacobson, N. Weaver, J. Gettys, BufferBloat: what's wrong with the internet? ACM Queue 9 (12) (2011) 10–20, doi:10.1145/2076796.2076798. ISSN 1542-7730.
[67] J. Gettys, Bufferbloat – Dark Buffers in the Internet, 2011.
[68] D.F. Williamson, R.A. Parker, J.S. Kendrick, The box plot: a simple visual method to interpret data, Ann. Internal Med. 110 (11) (1989) 916–921, doi:10.7326/0003-4819-110-11-916.
[69] C. Paasch, S. Ferlin, Özgü Alay, O. Bonaventure, Experimental evaluation of multipath TCP schedulers, in: Proceedings of the ACM SIGCOMM Workshop on Capacity Sharing Workshop (CSWS), ACM, New York/U.S.A., 2014, pp. 27–32, doi:10.1145/2630088.2631977. ISBN 978-1-4503-2991-0.

[70] I. Aydin, J. Iyengar, P. Conrad, C.-C. Shen, P. Amer, Evaluating tcp-friendliness in light of concurrent multipath transfer, Comput. Netw. 56 (7) (2012) 1876–1892. http://dx.doi.org/10.1016/j.comnet.2012.01.010.

[71] N. Arianpoo, I. Aydin, V.C.M. Leung, Network coding: a remedy for receiver buffer blocking in the concurrent multipath transfer of data over multi-hop wireless networks, in: Communications (ICC), 2014 IEEE International Conference on, 2014, pp. 3258–3263, doi:10.1109/ICC.2014.6883823.

[72] Y.-C. Chen, Y. Lim, R.J. Gibbens, E.M. Nahum, R. Khalili, D. Towsley, A measurement-based study of multiPath TCP performance over wireless networks, in: Proceedings of the 13th ACM Internet Measurement Conference (IMC), Barcelona/Spain, 2013, pp. 455–468, doi:10.1145/2504730.2504751. ISBN 978-1-4503-1953-9.

[73] S. Deng, R. Netravali, A. Sivaraman, H. Balakrishnan, WiFi, LTE, or both? Measuring multi-homed wireless internet performance, in: Proceedings of the 14th ACM Internet Measurement Conference (IMC), ACM, Vancouver, British Columbia/Canada, 2014, pp. 181–194, doi:10.1145/2663716.2663727. ISBN 978-1-4503-3213-2.

[74] A. Croitoru, D. Niculescu, C. Raiciu, Towards wifi mobility without fast handover, in: Proceedings of the 12th USENIX Symposium on Networked Systems Design and Implementation (NSDI), USENIX Association, Oakland, CA, 2015, pp. 219–234.

[75] K.J. Grinnemo, A. Brunstrom, A first study on using mptcp to reduce latency for cloud based mobile applications, in: 2015 IEEE Symposium on Computers and Communication (ISCC), 2015, pp. 64–69, doi:10.1109/ISCC.2015.7405495.

[76] C. Raiciu, D. Niculescu, M.B. Braun, M.J. Handley, Opportunistic mobility with multipath TCP, in: Proceedings of the Sixth International Workshop on MobiArch, in: MobiArch '11, ACM, Bethesda, Maryland/U.S.A., 2011, pp. 7–12, doi:10.1145/1999916.1999919. ISBN 978-1-4503-0740-6.

[77] C. Paasch, G. Detal, F. Duchêne, C. Raiciu, O. Bonaventure, Exploring mobile/WiFi handover with multipath TCP, in: Proceedings of the ACM SIGCOMM Workshop on Cellular Networks (CellNet), 2012, pp. 31–36, doi:10.1145/2342468.2342476. ISBN 978-1-4503-1475-6.

Kiran Yedugundla is a PhD student at the Computer Science Department, Karlstad University, Sweden. He received his masters by resarch in Computer science and engineering in 2007 from Indian Institute of Technology Madras (IITM), India. His current research interests include transport protocols, multipath communications. He has previously worked in the areas of optical burst switched networks and network management.

Simone Ferlin is an engineer and doctoral candidate at Simula Research Laboratory and university of Oslo. Her interests are primarily in computer networks; her work includes networked systems, transport protocols and congestion control, multipath transport, measurements and data analysis. Her past research includes signal processing for wireless communication. Her dissertation focuses on improving robustness in multipath transport. Although the multipath promises are great, increasing both connection throughput and resilience, it comes with challenges when the underlying network paths are heterogeneous – challenges which Simone's dissertation addresses through working with Multi-Path TCP (MPTCP), which has been under standardisation at the IETF and gaining increasing commercial deployment. Simone received her Dipl.-Ing. in Information Technology from Friedrich-Alexander Universitaet Erlangen-Nuernberg, Germany in 2010. She is now working in the NorNet project, which is a real-world large-scale multi-homing testbed for experimental research funded by the Norwegian research council. She is also very interested in biodiversity conservation and ways to integrate technology into environmental monitoring.

Thomas Dreibholz has received his Diplom (Dipl.-Inform.) degree in Computer Science from the University of Bonn in Bonn, Germany in 2001. Furthermore, he has received his Ph.D. degree (Dr. rer. nat.) in 2007 as well as his Habilitation (Priv.-Doz.) degree in 2012 from the University of Duisburg-Essen in Essen, Germany. Now, he works as a researcher on computer networks in the Network Systems Group of the Simula Research Laboratory in Fornebu, Norway. He has published and presented more than 45 research contributions at international conferences and in journals, on the topics of Reliable Server Pooling (RSerPool), the Stream Control Transmission Protocol (SCTP) and Quality of Service (QoS). Furthermore, he has contributed multiple Working Group and Individual Submission Drafts to the IETF standardisation processes of RSerPool and SCTP. He is also co-author of multiple RFC documents published by the IETF. In addition, he has written the RSerPool reference implementation.

Ozgu Alay received the B.S. and M.S. degrees in Electrical and Electronics Engineering from Middle East Technical University, Ankara, Turkey, in 2003 and 2006 respectively, and the Ph.D. degree in Electrical and Computer Engineering from Tandon School of Engineering of NYU, in 2010. She is currently a senior research scientist at Networks Department of Simula Research Laboratory, Norway. Her research interests lie in the areas of mobile broadband networks, multi-path protocols, robust multimedia transmission and cooperative communications.

Nicolas Kuhn received his Master degree in aeronautical engineering in 2010. From 2010 to 2013, he was a Ph.D. student at the Institut Supérieur d'Aéronautique et de l'Espace (ISAE) and at the National ICT Australia (NICTA) to obtain a Ph.D. from the University of Toulouse - EDMITT in December 2013. From January 2014 to September 2015, he was the principal investigator for Institut Mines-Télécom, Télécom Bretagne for the RITE european project. Since September 2015, he works at Centre National d'Etudes Spatiales (CNES) as a research engineer.

Per Hurtig received a M.Sc. and Ph.D. in Computer Science from Karlstad University, Sweden, in 2006 and 2012, respectively. He is currently Assistant Professor at the Department of Computer Science at Karlstad University. His research interests include transport protocols, low-latency internet communication, multi-path transport and network emulation. He has participated in several externally funded international research projects, within these areas, and also lead a number of national projects.

Anna Brunstrom received a B.Sc. in Computer Science and Mathematics from Pepperdine University, CA, in 1991, and a M.Sc. and Ph.D. in Computer Science from College of William & Mary, VA, in 1993 and 1996, respectively. She joined the Department of Computer Science at Karlstad University, Sweden, in 1996, where she is currently a Full Professor and Research Manager for the Distributed Systems and Communications Research Group. Her research interests include transport protocol design, techniques for low latency Internet communication, cross-layer interactions, multi-path communication and performance evaluation of mobile broadband systems. She has lead several externally funded research projects within these areas and served as the principal investigator and coordinator from Karlstad University (KaU) in additional national and international projects. She is currently the KaU principal investigator within three EU projects, the RITE FP7 project focused on reducing Internet transport latency, the NEAT H2020 project aiming to design a new, evolutive API and transport-layer architecture for the Internet, and the MONROE H2020 project proposing to design and operate an European transnational open platform for independent, multi-homed, large-scale monitoring and assessment of mobile broadband performance. She has authored/coauthored ten book chapters and over 100 international journal and conference papers.