

Technical University Berlin
Telecommunication Networks Group

Performance Evaluation of a
Common Congestion Controller for
TCP Connections

Michael Savorić, Holger Karl

{savoric,karl}@ee.tu-berlin.de

Berlin, May 2002

TKN Technical Report TKN-02-005

TKN Technical Reports Series
Editor: Prof. Dr.-Ing. Adam Wolisz

Abstract

Today's standard Internet transport protocol implementations perform flow and congestion control separately for each data stream, in isolation from all other data streams. It is advantageous in terms of improving the overall performance of the data streams, i.e., the throughput and fairness, to reuse network information and—as an extension—to establish a common congestion control between some of the data streams of an end system.

In this technical report, we describe the design goals and explain the algorithms of a common congestion control approach for TCP connections called "ensemble flow congestion management" (EFCM). In addition, we investigate the performance gain of the EFCM approach compared to the standard TCP congestion and flow control under different network conditions. Simulations with the EFCM approach show a considerable increase in throughput and fairness without increasing the aggressiveness of a set of TCP connections; the effect on background traffic is also negligible. The proposed EFCM controller algorithms are easy to implement and have a low additional complexity.

Keywords: TCP, Congestion Control, Flow Control, Network Information Reuse, Common Congestion Control

Contents

1	Introduction	3
2	Ensemble Flow Congestion Management (EFCM)	5
2.1	The EFCM design constraints	5
2.2	The EFCM controller	5
2.3	The EFCM jointly controlled TCP variables	6
2.4	The EFCM control algorithms	6
2.4.1	The network information reuse of the EFCM controller for a new TCP connection	6
2.4.2	The common congestion control of the EFCM controller for concurrent TCP connections	7
3	Simulation Model	9
3.1	Simulated network topology	9
3.2	Traffic load models	10
4	Evaluation Metric	12
5	Simulation Results	14
5.1	Overview	14
5.2	Scenario 1: 10 ms router-router delay, reliable last hop	16
5.3	Scenario 2: 10 ms router-router delay, unreliable last hop	18
5.4	Scenario 3: 30 ms router-router delay, reliable last hop	20
5.5	Scenario 4: 30 ms router-router delay, unreliable last hop	22
5.6	Scenario 5: 50 ms router-router delay, reliable last hop	24
5.7	Scenario 6: 50 ms router-router delay, unreliable last hop	26
5.8	Summary and discussion	28
6	Conclusion and Outlook	39
A	Complexity of the current EFCM controller	41
B	Statistical Evaluation	42
B.1	Statistical evaluation method	42
B.2	Statistical evaluation results	42

B.3 Summary of the statistical evaluation	46
C Current developments of the EFCM controller	47

Chapter 1

Introduction

In today's Internet, most of the data streams are using the TCP or UDP transport protocols. One of the most important responsibilities of a transport protocol are flow and congestion control. A TCP stream performs congestion control by slowly increasing its sending rate, probing the network's capacity and trying to adapt to it in order not to overload the network; in a sense, a TCP stream collects information about the current network conditions. For a UDP stream, these tasks are left to the application, since UDP has no built-in congestion and flow control.

For both types of transport protocols, all data streams of an end system act separately and independently of each other. This means that, for example, TCP data streams do not share their information about the current network conditions and UDP data streams do not adjust their sending rate to the current network conditions as detected by some TCP data streams. As a result, the overall performance of these data streams can be suboptimal, since the performance of each data stream is optimized only by using its locally available network information.

Exploiting such network information that can be present within an end system's protocol stack and sharing this information between multiple streams should improve overall performance. This information sharing can happen once at the start of a new data stream to initialize its flow and congestion control variables with more adequate values. This approach is called one-time network information reuse. As an extension, information can be shared continuously among several data streams during their whole lifetime in order to *jointly* control them; this second approach is called common congestion control. This technical report focuses on common congestion control.

Sharing network information between data streams can only happen among streams that use the same network path. Hence, a common congestion approach is only reasonable between data streams of an end system which have the same receiver or at least receivers in the same part of the network. These data streams form a (data stream) *ensemble*. The algorithms that determine which, how, and when network information among different data streams of an ensemble is shared form the actual controller of a common congestion control approach.

A common congestion controller's job can be divided into two main tasks: First, a common congestion controller has to manage the one-time network information exchange between *existing* (or *recently closed*) data streams of an ensemble and a *new* data stream joining this particular ensemble. This task is similar to the controller's job in existing pure network information reuse approaches like the ensemble or temporal TCP control block interdependence (TCBI) [11, 9, 10]. Second, a common congestion controller is responsible for the continuous network information exchange between *concurrent* data streams of an ensemble to reach a common congestion control for this ensemble.

These two tasks can be fulfilled in a number of different ways. The main ideas and methods of the four most relevant approaches have been described in reference [8]. One of these approaches, the ensemble TCP (E-TCP) [1], has been identified as a good basis for our new common congestion control approach called "ensemble flow congestion management" (EFCM). The E-TCP approach provides a common congestion control among TCP connections of an end system in a way that an ensemble of n TCP connections is no more aggressive to the network than a single TCP connection. Network information is shared between a new TCP connection and existing or recently closed TCP connections and between concurrent TCP connections. In addition, for every ensemble the E-TCP approach uses a scheduler that determines which TCP connection of an ensemble sends the next segment. A rate-based pacing mechanism can be optionally used for TCP connections of an ensemble. The assumption that an entire ensemble should be at most as aggressive as a single connection appears very conservative; it is here that EFCM and E-TCP differ most (details are presented in Chapter 2). We claim that the EFCM approach results in considerable performance gains at a moderate increase in implementation complexity and only limited adverse effects on the network.

Nevertheless, making common congestion control a practical solution is still faced with a practical challenge: In a simple common congestion controller the IP address of the receiver or the IP subnet address of the receivers are used to determine if some TCP connections can form an ensemble and use a common congestion control or not. If some mechanisms like NAT or Mobile IP are used, the IP addresses of the receivers can no longer form the criterion to build an ensemble, since with these mechanisms the IP addresses of the receivers do not reflect the current location of the receivers. Therefore, it might be sometimes difficult or even impossible to find out which TCP connections of an end system can form an ensemble and can use a common congestion control. Further research must be done on this topic to determine the constraints and possible solutions for using a common congestion control in the real Internet.

The remainder of this technical report is organized as follows: The design goals and algorithms of the EFCM approach are explained in Chapter 2. These algorithms were evaluated by simulations. The network topology and simulation scenarios are shown in Chapter 3, appropriate evaluation metrics are described in Chapter 4, and Chapter 5 presents and discusses the simulation results. Finally, Chapter 6 contains the conclusion of this technical report and gives an outlook on our future research activities. In Appendix A, a rough estimate of the additional time and space consumption of the EFCM controller compared to standard TCP is given. Appendix B contains both a description of the statistical evaluation method as well as a detailed exposition of the evaluation results. Since the EFCM controller is under development, Appendix C describes the additional algorithms of the latest version of the EFCM controller compared to the EFCM controller depicted in Chapter 2.

Chapter 2

Ensemble Flow Congestion Management (EFCM)

In Chapter 1, the general concept of sharing network information between some data streams has been presented, using the notion of a controller that manages the information exchange. A controller is an abstract entity which needs to be specified further to determine a concrete, implementable and testable functionality. One possibility of such a controller, the EFCM controller, is presented here.

2.1 The EFCM design constraints

The design constraints of the ensemble flow congestion management are:

- The control algorithms of the EFCM approach are confined to a sending end system, i.e., an end system where the senders of the data streams are located. This means that except for some code in the transport protocol no adaptations and additional changes neither in the network nor in the receiving end system(s) have to be done.
- The EFCM approach supports standard transport layer interfaces, i.e., sockets. Therefore, the EFCM is transparent for all applications and Internet services running on the end system.
- In contrast to the E-TCP approach, the algorithms of the EFCM controller ensure that an ensemble of n data streams must be no more aggressive to the network than n separate data streams of an end system.
- Another design constraint of the EFCM controller is a fair sharing of the available bandwidth among the data streams in an ensemble.

2.2 The EFCM controller

The current version of the EFCM controller performs one-time network information reuse for a new connection as well as common congestion control between concurrently existing connections of an ensemble. It does not, however, reuse network information obtained from recently closed TCP connections (as, e.g., temporal TCBI does [11]). It also does not control UDP data streams with network

information obtained from existing or recently closed TCP connections. No rate-based pacing mechanism is implemented in the current EFCM controller. Mechanisms to implement these functionalities are under development.

2.3 The EFCM jointly controlled TCP variables

TCP uses the following variables for congestion control: congestion window size (cwnd), slow start threshold (ssthresh), round trip time (rtt), smoothed round trip time (srtt), and round trip time variance (rttvar). The first two TCP control variables restrict the load a single TCP connection can send into the network, the last three TCP control variables lead to adequate timeout timer values for TCP segments send from a single TCP connection.

The EFCM controller jointly controls the congestion window size, the slow start threshold, the smoothed round trip time, and the round trip time variance of TCP connections in an ensemble. Hence, the EFCM controller restricts the load the TCP connections of an ensemble can send into the network and all TCP connections of an ensemble obtain the same adequate value for their timeout timer.

2.4 The EFCM control algorithms

In the following two paragraphs, we describe the proposed algorithms of the EFCM controller for both tasks, i.e., initializing new connections and updating concurrent connections, of a common congestion controller:

2.4.1 The network information reuse of the EFCM controller for a new TCP connection

If useful network information is available for a new TCP connection, the new TCP connection will reuse this network information and will start with more adequate values for the load the network can cope with and the timeout timer value. The algorithms of the EFCM controller for the network information reuse between existing TCP connections of an ensemble and a new TCP connection of the same ensemble are described in the following list:

Congestion window size: The EFCM controller computes the sum of all current congestion window sizes of the existing TCP connections of the ensemble plus the standard initial congestion window size 2, representing the new TCP connection. This value is used to calculate a fair share, i.e., an arithmetic mean value, of the congestion window size for all TCP connections in the ensemble. At the beginning of a new TCP connection, all TCP connections of the ensemble get this congestion window size fair share as their new congestion window size.

Slow start threshold: The EFCM controller computes the sum of all current slow start thresholds of the existing TCP connections of the ensemble plus the standard initial slow start threshold 64, representing the new TCP connection. This value is used to calculate a fair share of slow start threshold for all TCP connections in the ensemble. At the beginning of a new TCP connection, all TCP connections of the ensemble are assigned this slow start threshold fair share as their new slow start threshold.

Smoothed round trip time: The EFCM controller uses the current value of an aggregated smoothed round trip time of the existing TCP connections of an ensemble as the initial smoothed round trip time of the new TCP connection. If the new TCP connection is the only stream in its ensemble then the initial smoothed round trip time of the new TCP connection is set to the standard value.

Round trip time variance: The EFCM controller uses the current value of an aggregated round trip time variance of the existing TCP connections of an ensemble as the initial round trip time variance of the new TCP connection. If the new TCP connection is the only stream in its ensemble then the initial round trip time variance of the new TCP connection is set to the standard value.

Since in the current EFCM controller, no rate pacing for a new TCP connection is implemented, the first TCP segments of every new TCP connection will be sent in a burst.

2.4.2 The common congestion control of the EFCM controller for concurrent TCP connections

Whenever a standard TCP implementation would change the value of one of the commonly controlled variables of a connection, EFCM uses this change to trigger updates to these variables for all other connections within the same ensemble, according to the following rules:

Congestion window size: After every change of the congestion window size of one of the existing TCP connections in an ensemble, an aggregated congestion window size for this ensemble is computed by adding all current congestion window sizes of the TCP connections in the ensemble. This value is used to calculate a fair share of congestion window size. This congestion window size fair share is the new congestion window size of every TCP connection in an ensemble.

Slow start threshold: After every change of the slow start threshold of one of the existing TCP connections in an ensemble an aggregated slow start threshold for this ensemble is computed by adding all current slow start thresholds of the TCP connections in the ensemble. This value is used to calculate a fair share of slow start threshold. This value is the new slow start threshold of every TCP connection in an ensemble.

Smoothed round trip time: After every change of the smoothed round trip time of one of the n TCP connections in an ensemble an aggregated smoothed round trip time of this ensemble is updated by a weighted calculation of $(n - 1)/n$ times the last value of the aggregated smoothed round trip time plus $1/n$ times the new smoothed round trip time. All TCP connections in an ensemble get this calculation result of the smoothed round trip time as their new smoothed round trip time.

Round trip time variance: After every change of the round trip time variance of one of the n TCP connections in an ensemble an aggregated round trip time variance of this ensemble is updated by a weighted calculation of $(n - 1)/n$ times the last value of the aggregated round trip time variance plus $1/n$ times the new round trip time variance. All TCP connections in an ensemble get this calculation result of the round trip variance as their new round trip time variance.

If one of the TCP connections in an ensemble is affected by a packet loss, all TCP connections of the ensemble will fairly reduce their congestion window size and slow start threshold to the new calculated values. This maintains the congestion control of standard TCP if the packet loss is caused by congestion in the network.

If a TCP connection leaves the ensemble, i.e., the TCP connection has been closed, the current aggregated congestion window size and the current aggregated slow start threshold are fairly shared among the remaining TCP connections in the ensemble.

Since the current EFCM controller has no built-in rate pacing, some of the concurrent TCP connections of the ensemble can exhibit bursty sending behavior. For example, one TCP connection of an ensemble is waiting for acknowledgments while other connections of the same ensemble receive acknowledgments, consequently increasing the aggregated congestion window size and the fair share of the congestion window size. If the waiting TCP connection does receive an acknowledgment, it sends the now permitted, larger number of new TCP segments in a burst.

Evidently, these EFCM computations impose some overhead in time and space. A rough estimate of the additional time and space complexity of EFCM compared to standard TCP is given in Appendix A.

Chapter 3

Simulation Model

Evaluating the throughput and fairness gain of the EFCM approach is done by simulations. The network topology for these simulations is intended to reflect common client-server architectures in which the client end systems, i.e. the receivers, obtain information from applications running on the (optionally) EFCM-controlled server end system, i.e., the (EFCM) end system. In typical Internet setups, a server is connected to an ISP's backbone network with a high-bandwidth, reliable link. Typical client end systems are connected via networks with smaller bandwidths, e.g., DSL lines or low-speed LANs. These networks can be both reliable or, in case of modern wireless LAN installations, unreliable. The backbone network itself has a bandwidth-delay product varying over time depending on its current load. We chose our simulated network topology with regard to these properties of the Internet.

In addition to the network topology there are other factors that influence the performance of the EFCM approach, e.g., the number of TCP sender instances in the (EFCM) end system, the type of applications running on the (EFCM) end system, the round trip time between the (EFCM) end system and the receiver(s), the packet loss rate in the links between the (EFCM) end system and the receiver(s), the background traffic load in the path(s) from the (EFCM) end system to the receiver(s), and the maximum segment size (MSS) of TCP connections. For each of these factors a reasonable subset of values must be defined and for every combination of parameter settings standard TCP must be compared with the EFCM approach to reach a complete performance evaluation of the EFCM approach.

In this technical report, we consider one specific network model with a reasonable choice of parameters. This model is described in detail in the following Section 3.1; the load models are contained in Section 3.2.

3.1 Simulated network topology

The performance evaluation of the EFCM approach uses a simulated network topology shown in Figure 3.1. The simulation network topology consists of several TCP senders (S_1, \dots, S_3 and BS_1, BS_2) and TCP receivers (R_1, \dots, R_3 and BR_1, BR_2), two routers, and one Ethernet-type LAN on the sender side. The sender LAN is characterized by a bit rate of 100 Mbps with a propagation delay of $0.5 \mu\text{s}$. The routers are connected via links with a bit rate of 100 Mbps and a propagation delay of 10 ms, 30 ms, or 50 ms. For each incoming link the routers have a queuing capacity of 20 IP packets. Together with the given load in this tiny simulation model some packet losses in the routers can be

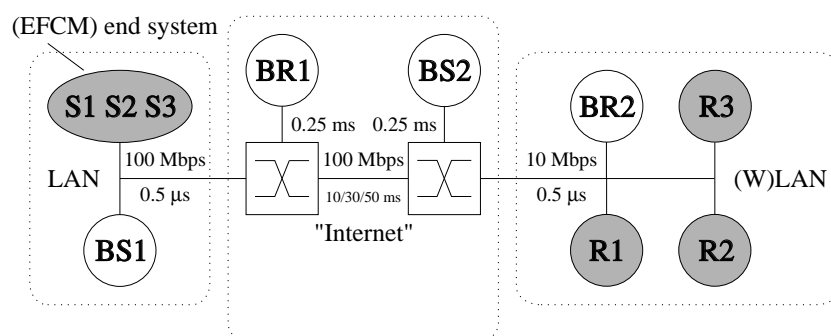


Figure 3.1: Structure of the simulated network topology

observed with this relatively small queuing capacity.

The TCP senders S1, S2, and S3 are located in an end system which is optionally equipped with an EFCM controller. The other TCP senders BS1 and BS2 are located in different end systems and generate background traffic.

The end system of the background traffic TCP sender BS2 is connected to the network via links with a bit rate of 100 Mbps and a propagation delay of 0.25 ms. The end system of the background traffic TCP sender BS1 and the EFCM end system are connected to the network via the sender LAN. The end systems of the TCP receivers R1, R2, and R3 and the background traffic TCP receiver BR2 are connected to the network via the receiver LAN.

The receiver LAN consists of an Ethernet-type shared medium with an overall bit rate of 10 Mbps and a propagation delay of 0.5 μs. The packet loss rate in the receiver LAN is adjustable to investigate the influence of different packet loss probabilities in the last hop of a TCP connection on the overall throughput of the TCP connections. Hence, with this receiver LAN either a reliable (wired) Ethernet or an unreliable wireless LAN can be modeled. In the wired last hop scenario, no errors occur in the receiver LAN; in the wireless LAN case, packets can be lost with a fixed packet loss rate of 5 percent.

In summary, the chosen parameter values should represent typical Internet scenarios fairly well. It would be particularly interesting consider the impact of various values for e.g. the round trip time, the maximum size of TCP segments, and the packet loss rate as important factors for the performance gain of EFCM. This technical report concentrates on varying the round trip time; other parameters are kept constant.

3.2 Traffic load models

Properly characterizing traffic loads for interactive Internet users is a difficult undertaking. In order to represent some mixture of different applications, we decided to use two different classes of TCP-based applications to generate traffic. The first class consists of short TCP connections with five segments to send. This class can be used to model, for example, short replies to search requests like a MP3 music piece search or the content of a file directory in an FTP server. The second class includes TCP connections whose number of segments to send is determined by a WWW traffic model [7]. This traffic model is derived from real HTTP traces in corporate and educational environments and uses three abstraction levels: The session level, the page level, and the packet level. Here, a

simplified version of this WWW traffic model is used which consists only of the first two levels. In every WWW session a log-normally distributed number of WWW pages with Pareto-distributed page sizes are sent. The time between the pages, i.e., the inter-connection or reading time, is gamma distributed. The load in the network can be easily adjusted by using the exponentially distributed session interarrival time with a different parameter. The inter-connection time of TCP connections derived from the first application class with short TCP connections is also gamma distributed. The distributions and parameters chosen for the stochastic variables of the simplified WWW traffic model are shown in Table 3.1.

Table 3.1: Distributions and parameters for the stochastic variables of the simplified WWW model

Stochastic variable	Distribution	Distribution parameter(s)
Inter-session time	Exponential	$\mu = 5.0$ s
Pages per session (pps)	Lognormal	$\mu = 25.807$ pps $\sigma = 78.752$ pps
Inter-page time (reading time)	Gamma	$\mu = 35.286$ s $\sigma = 147.390$ s
Page size	Pareto	$\alpha = 1.7584$ $\beta = 30458$ Bytes

The TCP senders of the (optionally) EFCM-equipped end system use both application classes: one TCP sender only opens short TCP connections, the other two TCP senders use WWW TCP connections.

The traffic of all background TCP senders is modeled by the second application class with modified inter-connection time and session interarrival time distributions, where TCP connections are immediately restarted once they have terminated. This is done to reach a higher load in the network with these few background traffic TCP senders.

The whole simulation model is implemented in ns-2 (version 2.1b8a) [2]. For all standard TCP connections, the ns-2 implementation of a TCP Newreno [3] sender or receiver, respectively, is used. The TCP connections of the (EFCM) end system are instances of a new TCP sender class (EFCM_TCP) derived from the ns-2 implementation of a TCP Newreno sender.¹² This new TCP sender class provides additional network information reuse and common congestion control mechanisms between the TCP connections of an EFCM ensemble and some statistical performance evaluation methods.

¹The current ns-2 implementation of the TCP Newreno sender and the TCP receiver does not perform the connection setup/teardown of a TCP connection. In future investigations of the EFCM approach also the influence of the TCP connection setup/teardown protocol mechanism on the overall throughput of a TCP connection will be considered.

²The page sizes in bytes of the senders is converted into page sizes in maximum segment sizes (MSSs) by using the ceil-function of C++. Due to a small bug in the simulation program, the page size in MSSs of background traffic sender 1 (BS 1) is equal or at most one smaller than the page sizes of the other senders.

Chapter 4

Evaluation Metric

In each simulated scenario and for all new TCP connections the mean throughput, the mean initial congestion window size, the mean initial slow start threshold, the mean initial smoothed round trip time, and the mean initial round trip time variance are compared between the standard TCP and the EFCM controller. In addition, also a fairness index between concurrent TCP connections of an EFCM ensemble is used to compare standard TCP with the EFCM approach: If n TCP connections are established in parallel for a period of time and reach the mean throughputs \bar{t}_i , $1 \leq i \leq n$, in this period of time, then for these TCP connections a fairness index for this period of time can be computed as follows [5]:

$$I_f = \frac{\left(\sum_{i=1}^n \bar{t}_i\right)^2}{n \cdot \sum_{i=1}^n \bar{t}_i^2} \quad \text{with} \quad \frac{1}{n} \text{ (very bad)} \leq I_f \leq 1 \text{ (excellent)}.$$

A fairness index of $1/n$ denotes that one of the n concurrent TCP connection gets the entire available bandwidth while a fairness index of 1 means that all n concurrent TCP connections get the same portion of the available bandwidth.

Only those TCP connections are considered in the comparison shown in the following tables which either are controlled by the EFCM approach or are not controlled but could be controlled by the EFCM approach (as at least one concurrent TCP connection to the same LAN is already established and useful information about the network is available). These TCP connections are called EFCM-capable TCP connections. In general, the percentage of EFCM-capable TCP connections depends on the type of the EFCM end system. For example, the numerous TCP connections of a large WWW or proxy server have a higher probability of using the EFCM approach than the few TCP connections of an ordinary end system. In order not to reflect this dependency in the evaluation, the metric computations were restricted to the EFCM-capable TCP connections and are hence independent of the end system's type. The overall performance of the EFCM approach on the throughput of all TCP connections can then be approximated by using the share of the EFCM TCP connections of all TCP connections of the EFCM end system.

To compare standard TCP with the EFCM controller, two different mean throughput computations for the EFCM-capable TCP connections of the two application classes are used. If one of these TCP connections has sent s segments in duration d , then the two mean throughput calculations work as follows:

- Computation of the overall mean throughput (\bar{T}_1): The sum of sent segments of all n TCP connections is divided by the overall duration of these TCP connections, i.e.:

$$\bar{T}_1 = \frac{\sum_{i=1}^n s_i}{\sum_{i=1}^n d_i}$$

- Computation of the connection-oriented mean throughput (\bar{T}_2): For each of the n TCP connections a mean throughput t is calculated. All these mean throughput values are then used to compute the overall mean throughput of the TCP connections by a normal non-weighted arithmetic mean calculation independent of the number of segments sent by each of the TCP connections, i.e.:

$$\bar{T}_2 = \frac{1}{n} \cdot \sum_{i=1}^n \frac{s_i}{d_i} = \frac{1}{n} \cdot \sum_{i=1}^n t_i$$

The former throughput calculation is the more important one, since with this throughput metric the overall throughput of the different TCP controllers can be evaluated. The latter calculation gives a connection-oriented mean throughput which can be understood as the mean throughput a single TCP connection can expect.

Chapter 5

Simulation Results

5.1 Overview

The same simulations were performed without and with the EFCM controller in a network scenario with 10 ms, 30 ms, or 50 ms delay between the two routers. The last hop was either reliable or had a packet loss rate (PLR) of 5 %. The simulation results for these simulation scenarios are shown in detail in the next sections and summarized and discussed at the end of this chapter.

In every following table, TCP 1 is a short TCP connection and TCP 2 and TCP 3 are WWW TCP connections of the (EFCM) end system. The simulation results of one simulation scenario shown in the tables are averages over ten independent simulation runs for the respective simulation scenario. Both stated mean throughput metrics ($\overline{T}_1, \overline{T}_2$) of the EFCM-capable TCP connections are measured in TCP segments per second; in every TCP segment, the payload length is set to 1000 bytes. For the new TCP connections entering an ensemble also the mean initial congestion window size (\overline{cwnd}), the mean initial slow start threshold ($\overline{ssthresh}$), the mean initial smoothed round trip time (\overline{srtt}), and the mean initial round trip time variance (\overline{rttvar}) are shown. For the concurrent TCP connections of an ensemble both mean throughput metrics ($\overline{T}_1, \overline{T}_2$) and the mean fairness index (\overline{I}_f) are shown.

The last column Δ in the following tables denotes whether the simulation results are statistically significantly different or not. A '+' or '-' denotes that the EFCM controller or standard TCP is significantly better. A '=' means that with the simulation results no significant difference between the EFCM and the standard TCP controller can be concluded. The details of the statistical evaluation of the simulation results can be found in Appendix B of this technical report.

For standard TCP connections the mean initial smoothed round trip time and the mean initial round trip time variance are not applicable (N/A) for the computation of the initial timeout timer, i.e., the initial timeout timer is set to the fixed standard value.

For all simulation scenarios the different TCP controllers (standard TCP/no EFCM, EFCM) are investigated for a simulated time of 250000 s in each simulation run. In this simulated time approximately 22150 TCP connections starting at the (EFCM) end system can be observed. Only some of them, i.e., those TCP connections which have concurrent TCP connections, are controlled or could be controlled by the new network information reuse and common congestion control mechanisms provided by the EFCM controller. In the simulation model and with the chosen traffic load model the percentage of concurrent TCP connections is relatively low. For example, an average computation over all simulations with a router-router delay of 10 ms shows that approximately 4.3 % of the

new TCP connections in the reliable last hop scenarios and approximately 9.1 % of the new TCP connections in the unreliable last hop scenarios are controlled or could be controlled by the EFCM.

In all simulations with the EFCM controller and with the observed occurrence of having concurrent TCP connections the mean throughput of the background TCP connections is not negatively affected by these EFCM-controlled TCP connections. Therefore, the EFCM-controlled TCP connections are no more aggressive to the network than standard TCP connections.

5.2 Scenario 1: 10 ms router-router delay, reliable last hop

Table 5.1 shows the simulation results for new TCP connections entering an ensemble. Compared to standard TCP, new short TCP connections benefit from the EFCM controller with a large gain of approximately 51 % for the overall mean throughput and of approximately 28 % for the connection-oriented mean throughput. New WWW TCP connections slightly benefit from the EFCM controller with a gain of approximately 3 % for the overall mean throughput compared to standard TCP. For the connection-oriented mean throughput the standard TCP and the EFCM controller achieve almost the same results for new WWW TCP connections

Table 5.1: Simulation results of Scenario 1 (new TCP connections) — \bar{T}_1 is the overall mean throughput, \bar{T}_2 is the connection-oriented mean throughput, Δ denotes the statistical significance of the simulation results ('+' : EFCM controller is significantly better, '=' : no significant difference between standard TCP and EFCM controller, '-' : standard TCP is significantly better)

		no EFCM	EFCM	Δ
\bar{T}_1 of new TCP connections [segments/second]	TCP 1	11.41	17.26	+
	TCP 2	68.68	70.80	=
	TCP 3	70.80	72.96	=
\bar{T}_2 of a new TCP connection [segments/second]	TCP 1	25.03	31.94	+
	TCP 2	84.66	85.29	=
	TCP 3	85.23	86.00	=
cwnd of a new TCP connection [segments]	TCP 1	2.00	11.41	
	TCP 2	2.00	8.51	
	TCP 3	2.00	8.42	
ssthresh of a new TCP connection [segments]	TCP 1	64.00	63.16	
	TCP 2	64.00	63.23	
	TCP 3	64.00	63.49	
srtt of a new TCP connection [ns-2 time unit]	TCP 1	N/A	7.44	
	TCP 2	N/A	6.21	
	TCP 3	N/A	6.18	
rttvar of a new TCP connection [ns-2 time unit]	TCP 1	N/A	3.88	
	TCP 2	N/A	3.22	
	TCP 3	N/A	3.24	

Table 5.2 shows the simulation results for concurrent TCP connections of an ensemble. Compared to standard TCP the EFCM controller produces a slight degradation of approximately 4 % for both throughput metrics. These results show that concurrent EFCM-controlled TCP connections are no more aggressive to the network in terms of bandwidth consumption than concurrent standard TCP connections. The huge fairness improvement of the EFCM controller for concurrent TCP connections is remarkable.

Table 5.2: Simulation results of Scenario 1 (concurrent TCP connections) — \bar{T}_1 is the overall mean throughput, \bar{T}_2 is the connection-oriented mean throughput, \bar{I}_f is the mean fairness index, Δ denotes the statistical significance of the simulation results ('+': EFCM controller is significantly better, '=': no significant difference between standard TCP and EFCM controller, '-': standard TCP is significantly better)

	no EFCM	EFCM	Δ
\bar{T}_1 of concurrent TCP connections [segments/second]	60.29	58.08	=
\bar{T}_2 of a concurrent TCP connection [segments/second]	50.80	49.03	-
\bar{I}_f of concurrent TCP connections []	0.7937	0.9390	+

5.3 Scenario 2: 10 ms router-router delay, unreliable last hop

Table 5.3 shows the simulation results for new TCP connections entering an ensemble. Compared to standard TCP, the EFCM controller achieves a huge gain for the overall mean throughput for both new short (116 %) and new WWW TCP (57 %) connections. For the connection-oriented mean throughput the EFCM controller achieves a gain of approximately 9 % for new short TCP connections and of approximately 7 % for new WWW TCP connections.

Table 5.3: Simulation results of Scenario 2 (new TCP connections) — \bar{T}_1 is the overall mean throughput, \bar{T}_2 is the connection-oriented mean throughput, Δ denotes the statistical significance of the simulation results ('+': EFCM controller is significantly better, '=': no significant difference between standard TCP and EFCM controller, '-': standard TCP is significantly better)

		no EFCM	EFCM	Δ
\bar{T}_1 of new TCP connections [segments/second]	TCP 1	5.24	11.34	+
	TCP 2	33.95	53.45	+
	TCP 3	34.52	53.83	+
\bar{T}_2 of a new TCP connection [segments/second]	TCP 1	57.74	63.05	+
	TCP 2	93.41	99.51	+
	TCP 3	92.97	99.46	+
cwnd of a new TCP connection [segments]	TCP 1	2.00	4.31	
	TCP 2	2.00	3.85	
	TCP 3	2.00	3.91	
ssthresh of a new TCP connection [segments]	TCP 1	64.00	45.06	
	TCP 2	64.00	50.22	
	TCP 3	64.00	50.06	
$\overline{\text{srtt}}$ of a new TCP connection [ns-2 time unit]	TCP 1	N/A	2.62	
	TCP 2	N/A	2.25	
	TCP 3	N/A	2.27	
$\overline{\text{rttvar}}$ of a new TCP connection [ns-2 time unit]	TCP 1	N/A	1.26	
	TCP 2	N/A	1.07	
	TCP 3	N/A	1.09	

Table 5.4 shows the simulation results for concurrent TCP connections of an ensemble. Compared to standard TCP, the EFCM controller achieves a large gain of the overall mean throughput of approximately 31 % and a gain of the connection-oriented throughput of approximately 8 % for concurrent TCP connections. For concurrent TCP connections only a slight fairness improvement of the EFCM controller can be observed.

Table 5.4: Simulation results of Scenario 2 (concurrent TCP connections) — \bar{T}_1 is the overall mean throughput, \bar{T}_2 is the connection-oriented mean throughput, \bar{I}_f is the mean fairness index, Δ denotes the statistical significance of the simulation results ('+': EFCM controller is significantly better, '=': no significant difference between standard TCP and EFCM controller, '-': standard TCP is significantly better)

	no EFCM	EFCM	Δ
\bar{T}_1 of concurrent TCP connections [segments/second]	37.69	49.44	+
\bar{T}_2 of a concurrent TCP connection [segments/second]	70.83	76.69	+
\bar{I}_f of concurrent TCP connections []	0.8013	0.8229	+

5.4 Scenario 3: 30 ms router-router delay, reliable last hop

Table 5.5 shows the simulation results for new TCP connections entering an ensemble. Compared to standard TCP, new short TCP connections benefit from the EFCM controller with a large gain of approximately 66 % for the overall mean throughput and of approximately 45 % for the connection-oriented mean throughput. New WWW TCP connections also benefit from the EFCM controller with a gain of approximately 20 % for the overall mean throughput and of approximately 17 % for the connection-oriented mean throughput compared to standard TCP.

Table 5.5: Simulation results of Scenario 3 (new TCP connections) — \bar{T}_1 is the overall mean throughput, \bar{T}_2 is the connection-oriented mean throughput, Δ denotes the statistical significance of the simulation results ('+': EFCM controller is significantly better, '=': no significant difference between standard TCP and EFCM controller, '-': standard TCP is significantly better)

		no EFCM	EFCM	Δ
\bar{T}_1 of new TCP connections [segments/second]	TCP 1	8.90	14.81	+
	TCP 2	39.55	47.11	+
	TCP 3	38.89	47.14	+
\bar{T}_2 of a new TCP connection [segments/second]	TCP 1	16.56	23.94	+
	TCP 2	52.62	61.25	+
	TCP 3	51.47	61.00	+
\overline{cwnd} of a new TCP connection [segments]	TCP 1	2.00	9.16	
	TCP 2	2.00	7.17	
	TCP 3	2.00	7.08	
$\overline{ssthresh}$ of a new TCP connection [segments]	TCP 1	64.00	54.43	
	TCP 2	64.00	56.83	
	TCP 3	64.00	56.69	
\overline{srtt} of a new TCP connection [ns-2 time unit]	TCP 1	N/A	12.83	
	TCP 2	N/A	11.07	
	TCP 3	N/A	10.87	
\overline{rttvar} of a new TCP connection [ns-2 time unit]	TCP 1	N/A	6.47	
	TCP 2	N/A	5.66	
	TCP 3	N/A	5.55	

Table 5.6 shows the simulation results for concurrent TCP connections of an ensemble. Compared to standard TCP, the EFCM controller achieves a gain of the overall mean throughput of approximately 12 % and a gain of the connection-oriented throughput of approximately 11 % for concurrent TCP connections. The huge fairness improvement of the EFCM controller for concurrent TCP connections is remarkable.

Table 5.6: Simulation results of Scenario 3 (concurrent TCP connections) — \bar{T}_1 is the overall mean throughput, \bar{T}_2 is the connection-oriented mean throughput, \bar{I}_f is the mean fairness index, Δ denotes the statistical significance of the simulation results ('+': EFCM controller is significantly better, '=': no significant difference between standard TCP and EFCM controller, '-': standard TCP is significantly better)

	no EFCM	EFCM	Δ
\bar{T}_1 of concurrent TCP connections [segments/second]	32.23	36.10	+
\bar{T}_2 of a concurrent TCP connection [segments/second]	31.99	35.44	+
\bar{I}_f of concurrent TCP connections []	0.8375	0.9133	+

5.5 Scenario 4: 30 ms router-router delay, unreliable last hop

Table 5.7 shows the simulation results for new TCP connections entering an ensemble. Compared to standard TCP, new short TCP connections benefit from the EFCM controller with a large gain of approximately 87 % for the overall mean throughput and of approximately 18 % for the connection-oriented mean throughput. New WWW TCP connections also benefit from the EFCM controller with a large gain of approximately 55 % for the overall mean throughput and of approximately 13 % for the connection-oriented mean throughput compared to standard TCP.

Table 5.7: Simulation results of Scenario 4 (new TCP connections) — \bar{T}_1 is the overall mean throughput, \bar{T}_2 is the connection-oriented mean throughput, Δ denotes the statistical significance of the simulation results ('+': EFCM controller is significantly better, '=': no significant difference between standard TCP and EFCM controller, '-': standard TCP is significantly better)

		no EFCM	EFCM	Δ
\bar{T}_1 of new TCP connections [segments/second]	TCP 1	5.18	9.68	+
	TCP 2	26.70	41.09	+
	TCP 3	26.25	40.79	+
\bar{T}_2 of a new TCP connection [segments/second]	TCP 1	25.31	29.98	+
	TCP 2	53.47	61.23	+
	TCP 3	54.16	60.48	+
\overline{cwnd} of a new TCP connection [segments]	TCP 1	2.00	4.28	
	TCP 2	2.00	3.86	
	TCP 3	2.00	3.82	
$\overline{ssthresh}$ of a new TCP connection [segments]	TCP 1	64.00	44.06	
	TCP 2	64.00	49.20	
	TCP 3	64.00	48.71	
\overline{srtt} of a new TCP connection [ns-2 time unit]	TCP 1	N/A	5.91	
	TCP 2	N/A	5.17	
	TCP 3	N/A	5.22	
\overline{rttvar} of a new TCP connection [ns-2 time unit]	TCP 1	N/A	1.87	
	TCP 2	N/A	1.76	
	TCP 3	N/A	1.76	

Table 5.8 shows the simulation results for concurrent TCP connections of an ensemble. Compared to standard TCP, the EFCM controller achieves a large gain of the overall mean throughput of approximately 35 % and a gain of the connection-oriented throughput of approximately 16 % for concurrent TCP connections. The EFCM controller significantly improves the fairness between concurrent TCP connections.

Table 5.8: Simulation results of Scenario 4 (concurrent TCP connections) — \bar{T}_1 is the overall mean throughput, \bar{T}_2 is the connection-oriented mean throughput, \bar{I}_f is the mean fairness index, Δ denotes the statistical significance of the simulation results ('+': EFCM controller is significantly better, '=': no significant difference between standard TCP and EFCM controller, '-': standard TCP is significantly better)

	no EFCM	EFCM	Δ
\bar{T}_1 of concurrent TCP connections [segments/second]	27.61	37.16	+
\bar{T}_2 of a concurrent TCP connection [segments/second]	37.53	43.42	+
\bar{I}_f of concurrent TCP connections []	0.8178	0.8515	+

5.6 Scenario 5: 50 ms router-router delay, reliable last hop

Table 5.9 shows the simulation results for new TCP connections entering an ensemble. Compared to standard TCP, new short TCP connections benefit from the EFCM controller with a large gain of approximately 70 % for the overall mean throughput and of approximately 50 % for the connection-oriented mean throughput. New WWW TCP connections also benefit from the EFCM controller with a large gain of approximately 29 % for the overall mean throughput and of approximately 24 % for the connection-oriented mean throughput compared to standard TCP.

Table 5.9: Simulation results of Scenario 5 (new TCP connections) — \bar{T}_1 is the overall mean throughput, \bar{T}_2 is the connection-oriented mean throughput, Δ denotes the statistical significance of the simulation results ('+' : EFCM controller is significantly better, '=' : no significant difference between standard TCP and EFCM controller, '-' : standard TCP is significantly better)

		no EFCM	EFCM	Δ
\bar{T}_1 of new TCP connections [segments/second]	TCP 1	7.55	12.82	+
	TCP 2	29.83	39.59	+
	TCP 3	30.78	38.52	+
\bar{T}_2 of a new TCP connection [segments/second]	TCP 1	12.60	18.96	+
	TCP 2	40.30	50.45	+
	TCP 3	41.19	50.64	+
\overline{cwnd} of a new TCP connection [segments]	TCP 1	2.00	8.53	
	TCP 2	2.00	7.07	
	TCP 3	2.00	7.07	
$\overline{ssthresh}$ of a new TCP connection [segments]	TCP 1	64.00	54.46	
	TCP 2	64.00	56.39	
	TCP 3	64.00	56.49	
\overline{srtt} of a new TCP connection [ns-2 time unit]	TCP 1	N/A	16.15	
	TCP 2	N/A	14.50	
	TCP 3	N/A	14.16	
\overline{rttvar} of a new TCP connection [ns-2 time unit]	TCP 1	N/A	7.07	
	TCP 2	N/A	6.44	
	TCP 3	N/A	6.29	

Table 5.10 shows the simulation results for concurrent TCP connections of an ensemble. Compared to standard TCP, the EFCM controller achieves a gain of the overall mean throughput of approximately 13 % and a gain of the connection-oriented throughput of approximately 12 % for concurrent TCP connections. The huge fairness improvement of the EFCM controller for concurrent TCP connections is remarkable.

Table 5.10: Simulation results of Scenario 5 (concurrent TCP connections) — \bar{T}_1 is the overall mean throughput, \bar{T}_2 is the connection-oriented mean throughput, \bar{I}_f is the mean fairness index, Δ denotes the statistical significance of the simulation results ('+': EFCM controller is significantly better, '=': no significant difference between standard TCP and EFCM controller, '-': standard TCP is significantly better)

	no EFCM	EFCM	Δ
\bar{T}_1 of concurrent TCP connections [segments/second]	25.50	28.70	+
\bar{T}_2 of a concurrent TCP connection [segments/second]	25.04	28.16	+
\bar{I}_f of concurrent TCP connections []	0.8416	0.9207	+

5.7 Scenario 6: 50 ms router-router delay, unreliable last hop

Table 5.11 shows the simulation results for new TCP connections entering an ensemble. Compared to standard TCP, new short TCP connections benefit from the EFCM controller with a large gain of approximately 81 % for the overall mean throughput and of approximately 18 % for the connection-oriented mean throughput. New WWW TCP connections also benefit from the EFCM controller with a large gain of approximately 54 % for the overall mean throughput and of approximately 19 % for the connection-oriented mean throughput compared to standard TCP.

Table 5.11: Simulation results of Scenario 6 (new TCP connections) — \bar{T}_1 is the overall mean throughput, \bar{T}_2 is the connection-oriented mean throughput, Δ denotes the statistical significance of the simulation results ('+': EFCM controller is significantly better, '=': no significant difference between standard TCP and EFCM controller, '-': standard TCP is significantly better)

		no EFCM	EFCM	Δ
\bar{T}_1 of new TCP connections [segments/second]	TCP 1	4.79	8.65	+
	TCP 2	21.20	32.33	+
	TCP 3	20.95	32.71	+
\bar{T}_2 of a new TCP connection [segments/second]	TCP 1	16.66	19.71	+
	TCP 2	36.88	43.72	+
	TCP 3	36.90	43.87	+
\overline{cwnd} of a new TCP connection [segments]	TCP 1	2.00	4.17	
	TCP 2	2.00	3.85	
	TCP 3	2.00	3.83	
$\overline{ssthresh}$ of a new TCP connection [segments]	TCP 1	64.00	43.68	
	TCP 2	64.00	48.27	
	TCP 3	64.00	48.78	
\overline{srtt} of a new TCP connection [ns-2 time unit]	TCP 1	N/A	9.41	
	TCP 2	N/A	8.38	
	TCP 3	N/A	8.36	
\overline{rttvar} of a new TCP connection [ns-2 time unit]	TCP 1	N/A	2.58	
	TCP 2	N/A	2.51	
	TCP 3	N/A	2.56	

Table 5.12 shows the simulation results for concurrent TCP connections of an ensemble. Compared to standard TCP, the EFCM controller achieves a large gain of the overall mean throughput of approximately 34 % and a gain of the connection-oriented throughput of approximately 19 % for concurrent TCP connections. The EFCM controller significantly improves the fairness between concurrent TCP connections.

Table 5.12: Simulation results of Scenario 6 (concurrent TCP connections) — \bar{T}_1 is the overall mean throughput, \bar{T}_2 is the connection-oriented mean throughput, \bar{I}_f is the mean fairness index, Δ denotes the statistical significance of the simulation results ('+': EFCM controller is significantly better, '=': no significant difference between standard TCP and EFCM controller, '-': standard TCP is significantly better)

	no EFCM	EFCM	Δ
\bar{T}_1 of concurrent TCP connections [segments/second]	21.18	28.45	+
\bar{T}_2 of a concurrent TCP connection [segments/second]	24.95	29.78	+
\bar{I}_f of concurrent TCP connections []	0.8288	0.8599	+

5.8 Summary and discussion

Evidently, the benefits and disadvantages of a joint congestion control depend to a large degree on the system scenario. To summarize:

- **Reliable last hop scenario:** In all considered scenarios with a reliable last hop, for new short TCP connections the EFCM controller achieves a considerable gain in the overall mean throughput over standard TCP. In Scenario 1, for new WWW TCP connections the EFCM controller slightly increases the overall mean throughput compared to standard TCP. In Scenario 3 and 5, the EFCM controller noticeably increases the overall mean throughput of new WWW TCP connections compared to standard TCP. For concurrent TCP connections no real difference can be observed for the overall mean throughput between standard TCP and EFCM-controlled TCP in Scenario 1. In Scenario 3 and 5, the overall mean throughput of concurrent TCP connections is higher if they are EFCM-controlled. A huge fairness gain for concurrent TCP connections can be observed if the EFCM controller is used.

For new short and for new WWW TCP connections the performance gain of the EFCM controller increases with the round trip time between the TCP senders and the TCP receivers. This expected result is based on the network information reuse algorithms of the EFCM controller that lower the negative influence of standard TCP slow start mechanism on the throughput of TCP connections over paths with large round trip times. The same correlation between the throughput and the round trip time can be observed for concurrent TCP connections.

Figures 5.1 and 5.2 show some typical processes of the congestion window size and the sequence number for either two standard TCP connections or two EFCM-controlled TCP connections of one ensemble over a reliable last hop (Scenario 1). In Figure 5.1, the concurrent TCP connections are separately controlled and obtain different congestion window sizes. In Figure 5.2, the concurrent TCP connections are jointly controlled and obtain the same congestion window sizes. It can be seen that one of the concurrent TCP connections has sometimes a bursty sending behavior, since the current EFCM controller has no built-in rate pacing mechanism.

- **Unreliable last hop scenario:** In all considered scenarios and for both new short and new WWW TCP connections and for concurrent TCP connections, the EFCM controller achieves a large gain in the overall mean throughput compared to standard TCP. Therefore, a common congestion control for TCP connections can partly prevent the negative influence of packet losses in the last hop of the network on the throughput of TCP connections. The only slight fairness gain of the EFCM controller for concurrent TCP connections can be explained as follows: If two or more TCP connections of one ensemble are established in parallel during one measurement period for the fairness index, they have equal values for their jointly controlled TCP variables. This is ensured by the EFCM controller. But the EFCM controller can not prevent that these TCP connections observe different segment loss frequencies during this measurement period. Therefore, some TCP senders have to wait for acknowledgments while other TCP senders can send their (new) segments. As a result, the concurrent TCP connections of an ensemble can have (very) different throughputs in a measurement period. Hence, the fairness index can be low.

For new short TCP connections the performance gain of the EFCM controller decreases but remains on a high level with the round trip time between the TCP senders and the TCP receivers. For new WWW TCP connections the performance gain of the EFCM controller is nearly independent from the round trip time. For concurrent TCP connections the performance gain of the EFCM controller slightly increases with the round trip time.

Figures 5.3 and 5.4 show some typical processes of the congestion window size and the sequence number for either two standard TCP connections or two EFCM-controlled TCP connections over an unreliable last hop. In Figure 5.3, the concurrent TCP connections are separately controlled and use different congestion window sizes. In Figure 5.4, the concurrent TCP connections are jointly controlled and use the same congestion window sizes. Due to packet losses some concurrent TCP connections have to wait for a longer period of time to send some new segments. But the negative influence of packet losses on the throughput of TCP connections is noticeably absorbed by a common congestion control. Due to the smaller increase of the aggregated congestion window size compared to the scenarios with a reliable last hop a remarkable decrease of the bursty sending behavior of concurrent TCP connections can be observed.

In the unreliable last hop scenario the EFCM controller significantly improves both throughput metrics compared to standard TCP. In this scenario the negative influence of packet losses on the throughput of TCP connections can be compensated for by sharing the decreased congestion window and slow start threshold of the TCP connection affected by packet losses among all ensemble members.

Figures 5.5 and 5.6 show the histograms of the fairness index for concurrent standard or EFCM-controlled TCP connections for both last hop scenarios in the simulations with a router-router delay of 10 ms. Especially for the reliable last hop scenario the positive influence of the EFCM controller on the fairness of concurrent TCP connections is obvious. Figures 5.7 and 5.8 show the histograms of the throughput for concurrent standard or EFCM-controlled TCP connections for both last hop scenarios in the simulations with a router-router delay of 10 ms.

It is remarkable that in the simulation Scenario 2 with an unreliable last hop the connection-oriented mean throughput (not the overall mean throughput!) of most TCP connections which are controlled or could be controlled by the EFCM approach is higher than in the simulation Scenario 1 with a reliable last hop. But this—at first astonishing—result can be easily explained: The background traffic TCP connections with receivers in the receiver LAN are often affected by packet losses in the unreliable last hop. Due to the TCP congestion control algorithms these background traffic TCP connections reach a smaller overall sending rate, allocate less bandwidth, and produce a substantially lower load in the receiver LAN. The other TCP connections with receivers in the receiver LAN are also affected by packet losses in the unreliable last hop. But from a single TCP connection's point of view, the lower load in the shared medium of the receiver LAN can sometimes compensate for and even overcompensate for the in general negative influence of packet losses in an unreliable last hop on the mean throughput of a TCP connection. For example, most of the short TCP connections and some of the WWW TCP connections are not affected at all by packet losses during their whole lifetime. These TCP connections can highly benefit from the lower load in the receiver LAN and reach a higher connection-oriented mean throughput. Therefore, the observed higher connection-oriented mean throughput in the unreliable last hop scenario is only based on the lower load in the receiver LAN.

In Table 5.13 the performance of the EFCM controller derived from the simulation results is summarized.

Table 5.13: Performance of the EFCM controller for the different scenarios — $\bar{T}_{\text{new,short}}$ or $\bar{T}_{\text{new,www}}$ are the overall mean throughput for short or WWW TCP connections, $\bar{T}_{\text{concurrent}}$ is the overall mean throughput for concurrent TCP connections, $\bar{I}_{f,\text{concurrent}}$ is the mean fairness index for concurrent TCP connections

	Reliable last hop	Unreliable last hop
$\bar{T}_{\text{new,short}}$	++	++
$\bar{T}_{\text{new,www}}$	=, +, +	++
$\bar{T}_{\text{concurrent}}$	=, +, +	++
$\bar{I}_{f,\text{concurrent}}$	++	+

A '++' or '+' denotes that the EFCM controller achieves a large gain or a gain (with respect to the performance metric shown in this row) compared to the standard TCP controller; a '=' denotes that no significant gain between the standard TCP and the EFCM controller can be observed. If the performance of the EFCM controller for a specific value is different in the scenarios then the performance result of each scenario is stated in the cell of the table which belongs to this value.

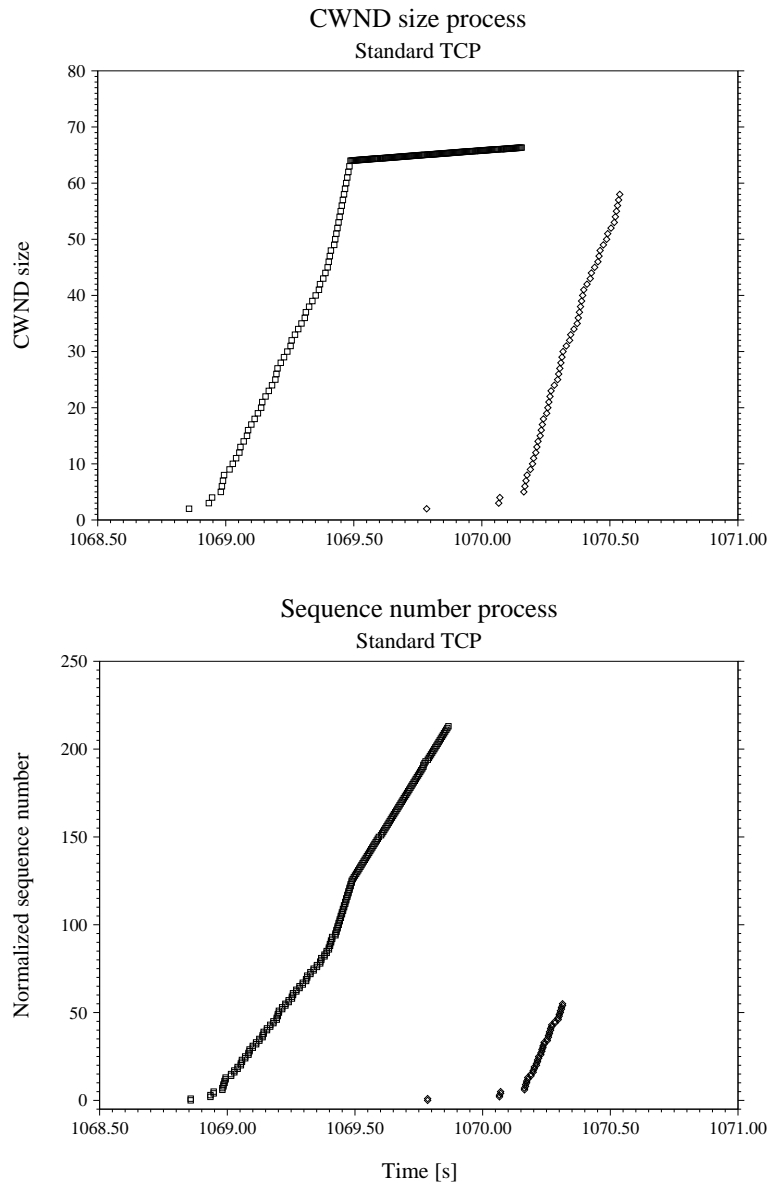


Figure 5.1: Standard TCP connections over a reliable last hop (Scenario 1) — \square and \diamond represent the first and second TCP connection of an ensemble, respectively

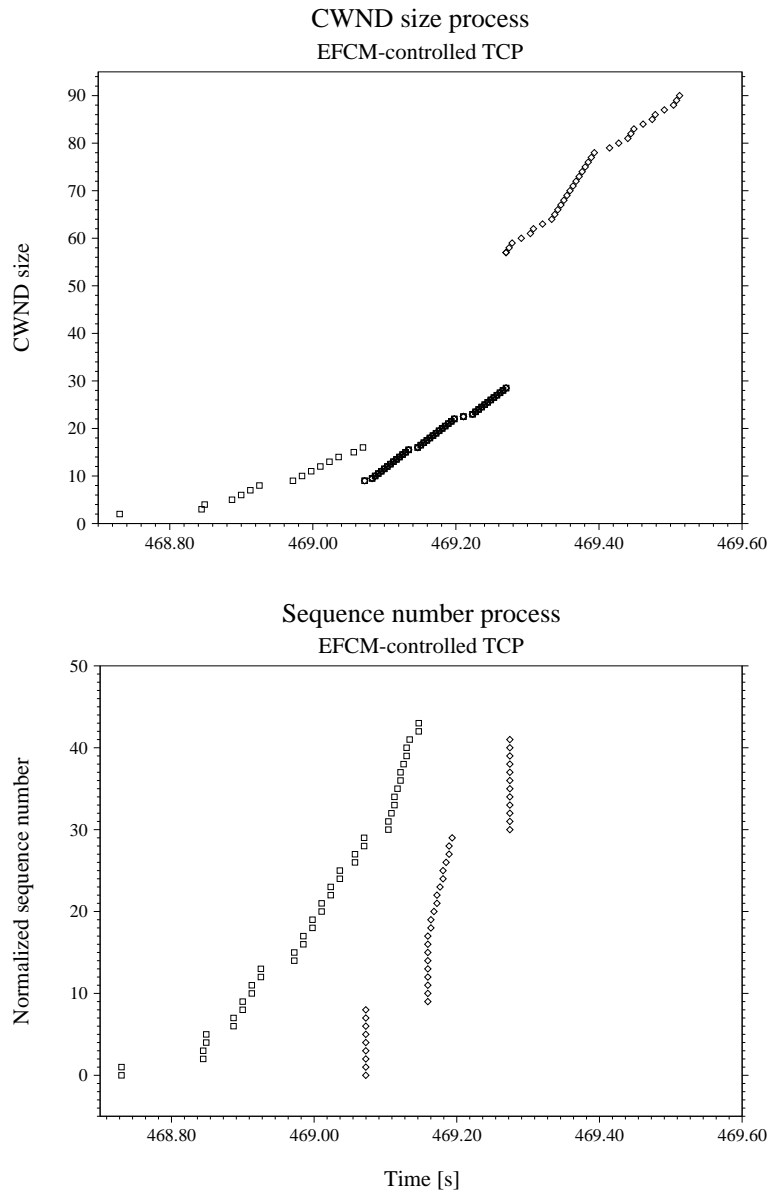


Figure 5.2: EFCM-controlled TCP connections over a reliable last hop (Scenario 1) — □ and ◇ represent the first and second TCP connection of an ensemble, respectively

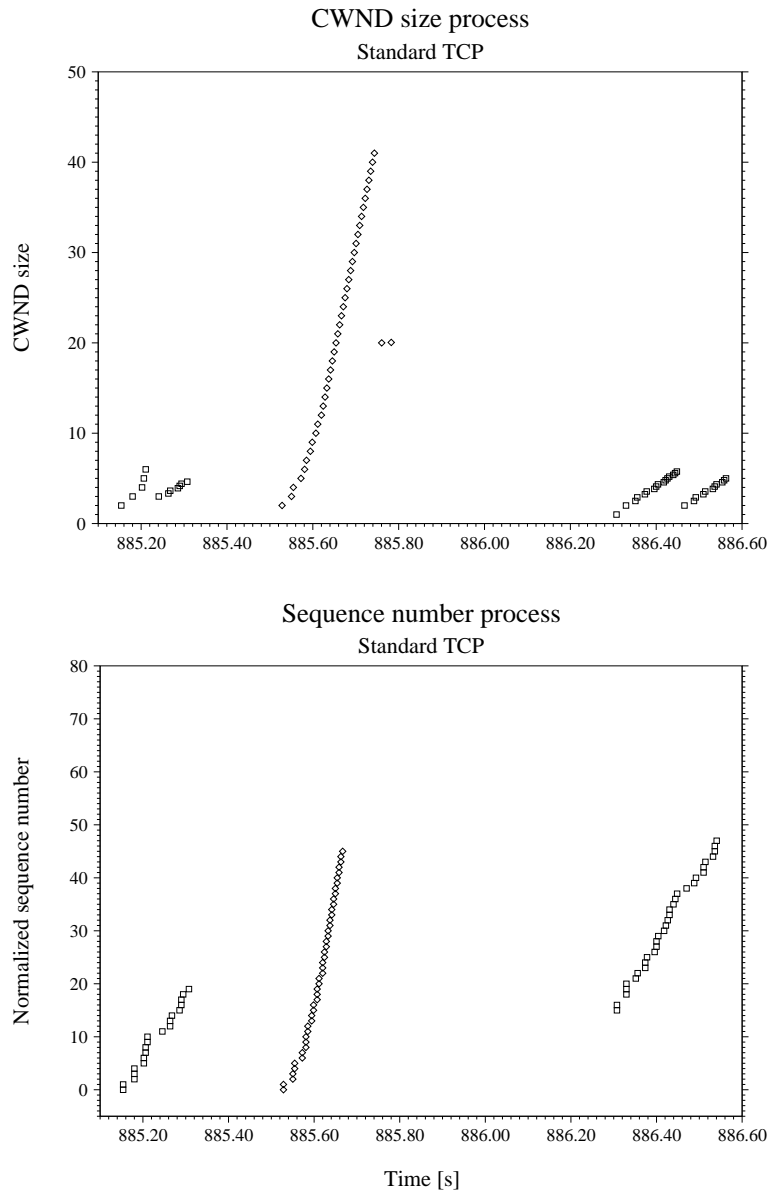


Figure 5.3: Standard TCP connections over an unreliable last hop (Scenario 2) — □ and ◇ represent the first and second TCP connection of an ensemble, respectively

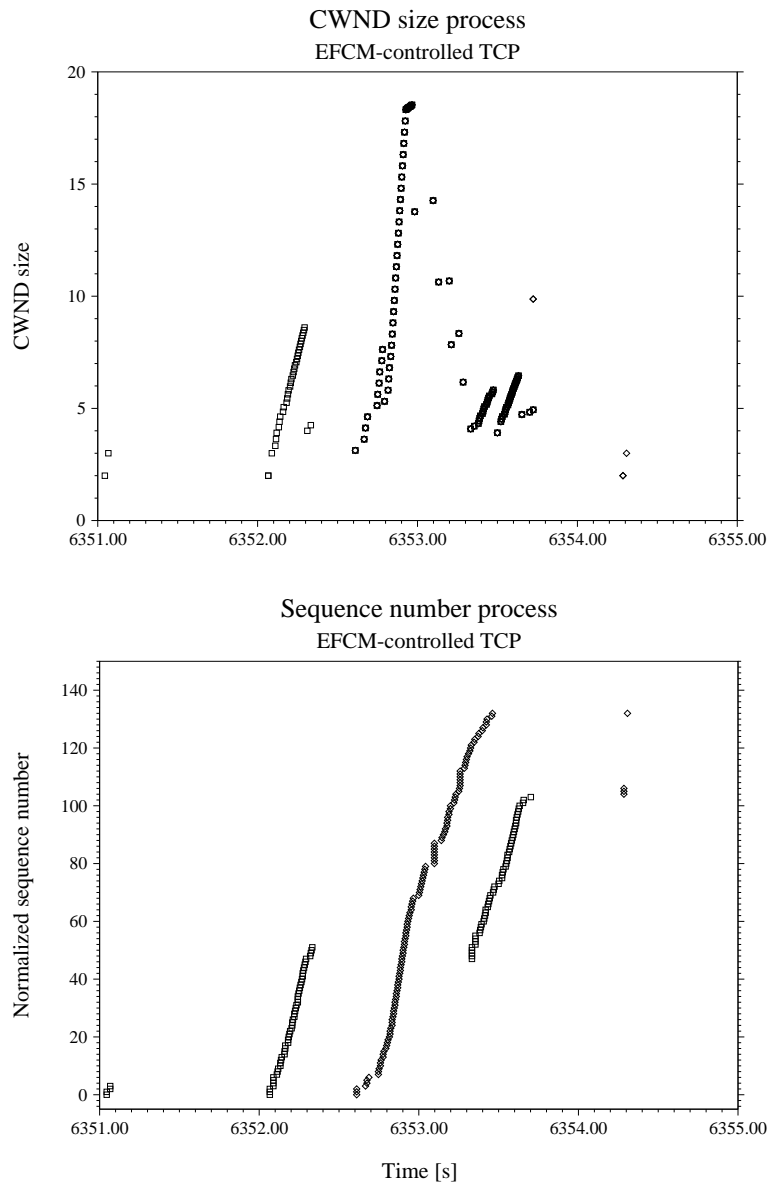


Figure 5.4: EFCM-controlled TCP connections over an unreliable last hop (Scenario 2) — \square and \diamond represent the first and second TCP connection of an ensemble, respectively

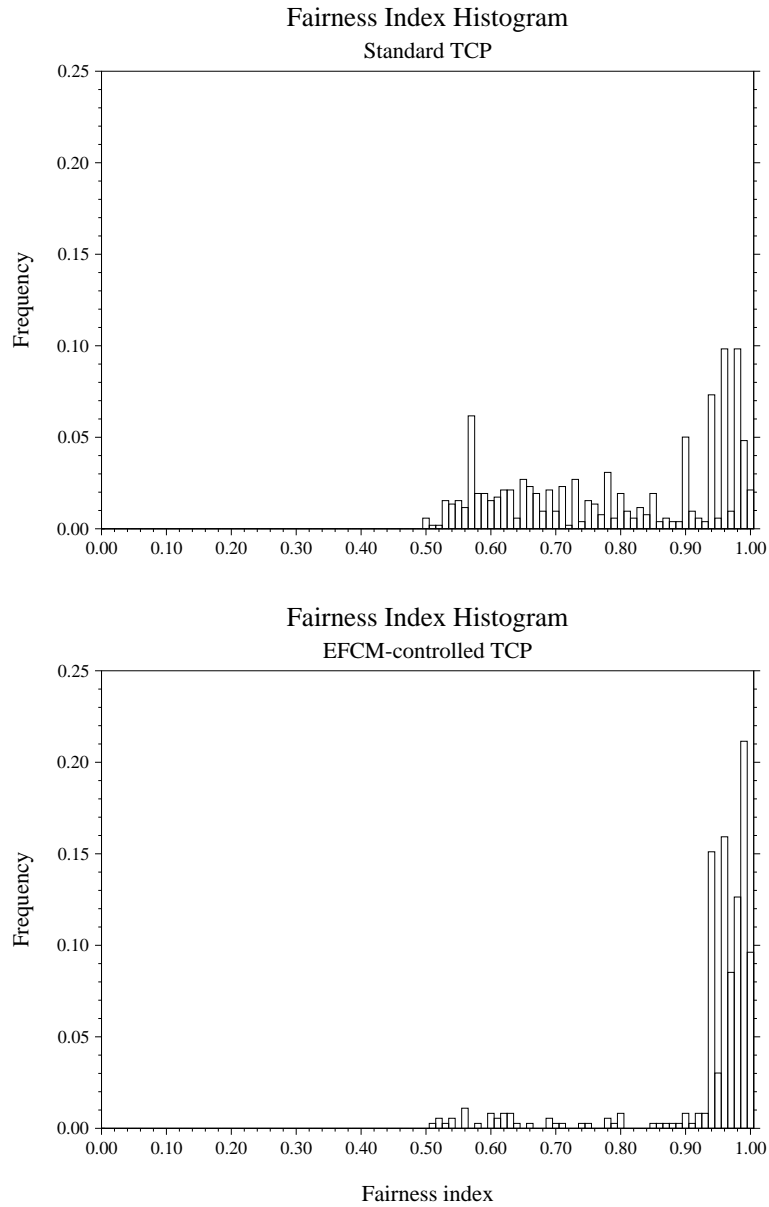


Figure 5.5: Fairness index histogram for concurrent standard or EFCM-controlled TCP connections over a reliable last hop (Scenario 1)

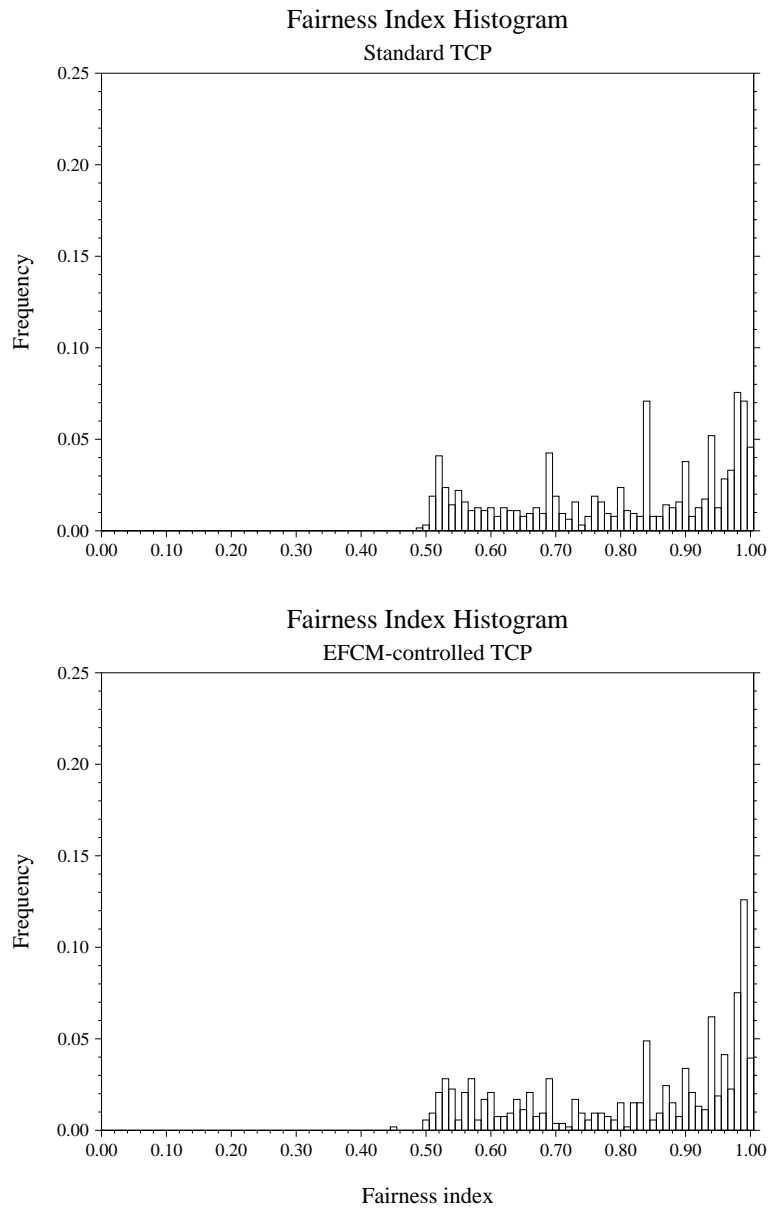


Figure 5.6: Fairness index histogram for concurrent standard or EFCM-controlled TCP connections over an unreliable last hop (Scenario 2)

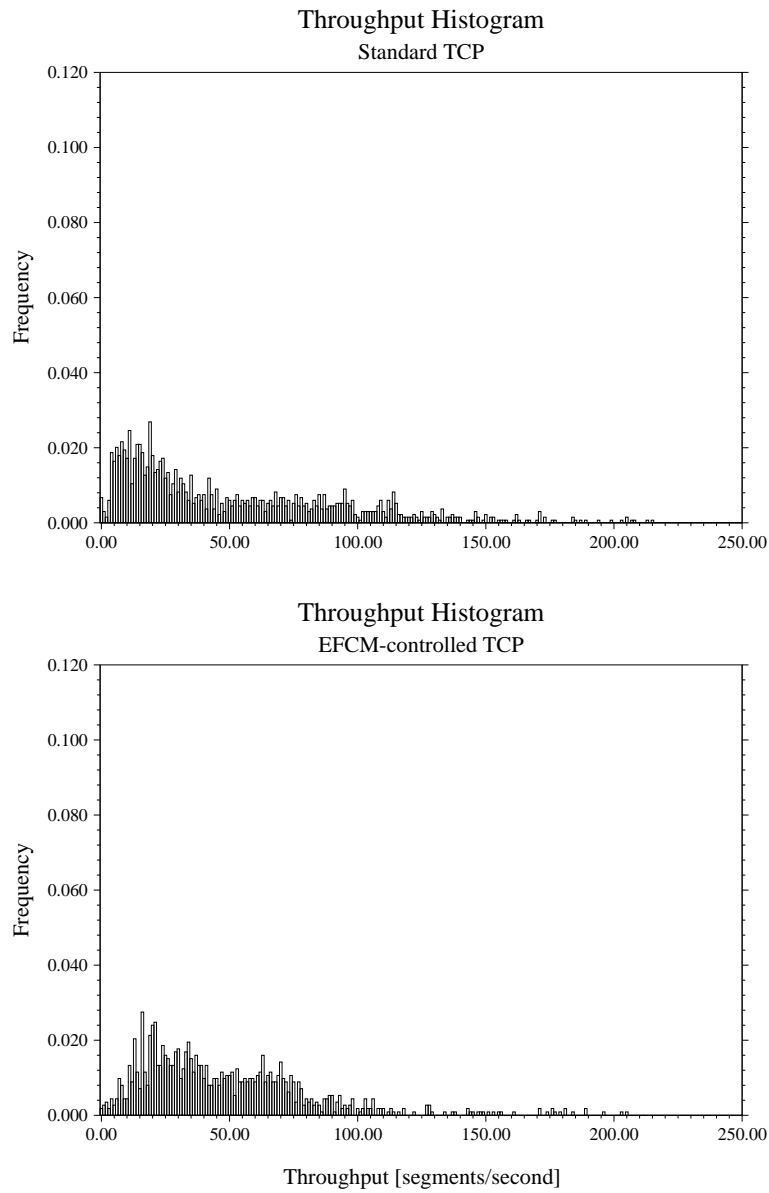


Figure 5.7: Throughput histogram for concurrent standard or EFCM-controlled TCP connections over a reliable last hop (Scenario 1)

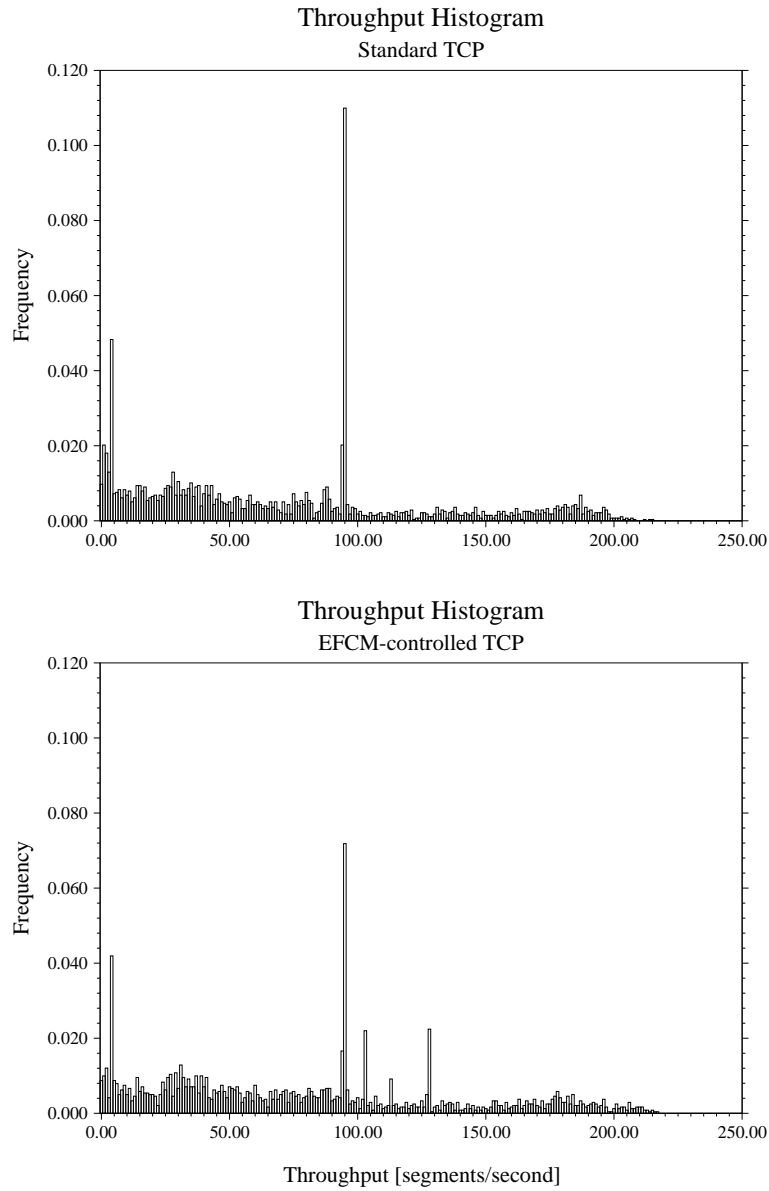


Figure 5.8: Throughput histogram for concurrent standard or EFCM-controlled TCP connections over an unreliable last hop (Scenario 2)

Chapter 6

Conclusion and Outlook

The simulation results show that the common congestion control approach EFCM improves the performance of standard TCP. If the EFCM controller is used in a scenario with a reliable last hop, a remarkably improved fairness can be observed between concurrent TCP connections of an ensemble. Also, short connections enjoy a much improved throughput. In the unreliable last hop scenario, the EFCM approach significantly improves the throughput of concurrent TCP connections of an ensemble. Particularly for such unreliable last hop scenarios the usage of the EFCM approach is highly recommended.

These benefits are achieved despite the fact that the EFCM controller investigated here uses relatively simple algorithms. In future investigations of the EFCM approach, some more complex algorithms will be considered, in the given simulation model with the here described parameter setting as well as in extended simulation models with the same and other parameter settings. For example, if one TCP connection of an ensemble does not use its whole fair share of the aggregated sending rate in a given time interval then the other TCP connections of the ensemble should increase their sending rate to reach the allowed aggregated sending rate. The intermittently silent TCP connection should get a credit for this unused sending rate to be allowed to have a sending rate higher than the fair share in the future. This new controller mechanism can be implemented by using a utilization factor for each TCP connection in an ensemble. More generally, allowing the controller to adapt sending rates within an ensemble enables new form of application support. As an example, the transport layer could support interactive applications by assigning them a greater share of an ensemble's total bandwidth budget and reducing this share again when the application has no data to send (compensating other applications for their backlog).

Another important extension of the current EFCM controller is the implementation of a rate pacing mechanism to prevent a bursty sending behavior, which can occur for new and, in some cases, concurrent TCP connections of an ensemble. The scope of an ensemble can be extended by including information from recently closed connections, assuming that the network conditions are slowly varying at best; the speed with which such information becomes inaccurate is an important parameter for such an undertaking. In a similar vein, the EFCM approach could also be used for handover TCP connections, i.e., TCP connections whose mobile receivers perform a handover. In this case, the handover must be non-transparent for the transport layer of the EFCM end system.

The mid-term objective will be the extension of the EFCM approach to jointly control TCP and UDP data streams, i.e., to reach a TCP-friendly behavior of UDP data streams in an EFCM end system. In addition, it is planned to implement a framework for a common congestion controller into

the linux kernel to evaluate the performance of the EFCM controller by measurements in the Internet environment.

Appendix A

Complexity of the current EFCM controller

The complexity of a controller can be expressed using the O-calculus [6] that gives an asymptotic upper bound of a complexity metric, e.g., the time or space consumption of an algorithm or a controller. Here, the time and space complexity of the EFCM controller is compared with standard TCP by using an implementation of the EFCM controller which is optimized for time consumption.

Let v be the number of jointly controlled TCP variables of the EFCM controller. For v_{fs} of these TCP control variables a fair share calculation, e.g., for the congestion window size or the slow start threshold, and for the remaining v_w TCP control variables a weighted calculation, e.g., for the smoothed round trip time and the round trip time variance, is used. Let n be the number of concurrent TCP connections which form an ensemble and are jointly controlled by the EFCM controller. For every change of one of the jointly controlled TCP variables of the EFCM controller, the additional time consumption of the EFCM controller compared with standard TCP can be estimated as:

$$\begin{aligned} \# \text{ Additions } (+, -) &= O(1) \\ \# \text{ Multiplications } (\cdot, /) &= O(1) \\ \# \text{ Assignments} &= O(n) \end{aligned}$$

The considered implementation of the EFCM controller has an additional space consumption compared to standard TCP whose upper bound is in $O(n \cdot v_{fs} + v_w)$ per ensemble. If an EFCM controller has to manage e ensembles at the same time the upper bound of the additional space consumption is in $O(e + \sum_{i=1}^e n_i \cdot v_{fs,i} + v_{w,i})$ where the additional e in the equation represents the space consumption of the search list that is used to determine whether a new TCP connection can join an existing ensemble or not. The overall additional time and space complexity of the EFCM controller compared to standard TCP is low. And this additional complexity, except for the search list, is only needed if the EFCM controller has to manage TCP connections in an ensemble.

Appendix B

Statistical Evaluation

B.1 Statistical evaluation method

For the simulated scenarios with either a reliable or an unreliable last hop the results of the standard TCP (no EFCM) controller are statistically compared with the results of the EFCM controller.

The statistical evaluation method used for this comparison is called the t-test for unpaired observations of two alternatives and is described in detail in [4]. The main idea of this method is to compute a confidence interval for the difference of the mean values of both alternatives for a given confidence level. Then the decision criterion is:

- If the confidence interval includes zero, then the two alternatives can not be distinguished.
- If the confidence interval is above/below zero, then the first/second alternative is the better one.

Tests with confidence intervals give not only a yes-no answer like other hypothesis tests, they also give an answer to the question how precise the decision is. A narrow confidence interval indicates that the precision of the decision is high whereas a wide confidence interval indicates that the precision of the decision is rather low.

This t-test for unpaired observations of two alternatives is used for the statistical evaluation of the simulation results for both the overall mean throughput (\bar{T}_1) and the connection-oriented mean throughput (\bar{T}_2) of new or concurrent TCP connections controlled by the standard TCP or the EFCM controller and the mean fairness index (\bar{I}_f) of concurrent TCP connections.

The values for this statistical evaluation are produced by five independent simulation runs for every last hop scenario in the simulation model.

B.2 Statistical evaluation results

In the following Tables B.1 to B.6 the statistical evaluation of the simulation results are shown. For each confidence interval also the confidence level (0.90, 0.95 or 0.99) is depicted. If the simulation results for the standard TCP and the EFCM controller are not significantly different even for the confidence level 0.90, then the confidence interval for the confidence level 0.90 is stated.

Table B.1: Statistical evaluation of the simulation results of Scenario 1 (10 ms router-router delay, reliable last hop)

		no EFCM ↔ EFCM
\bar{T}_1 of new TCP connections [1 - α : conf()]	TCP 1	0.99 : (- 8.32, - 3.39)
	TCP 2	0.90 : (- 7.37, + 3.14)
	TCP 3	0.90 : (- 6.37, + 2.06)
\bar{T}_2 of a new TCP connection [1 - α : conf()]	TCP 1	0.99 : (- 8.11, - 5.71)
	TCP 2	0.90 : (- 2.66, + 1.41)
	TCP 3	0.90 : (- 2.68, + 1.14)
\bar{T}_1 of concurrent TCP connections [1 - α : conf()]		0.90 : (- 0.88, + 5.29)
\bar{T}_2 of a concurrent TCP connection [1 - α : conf()]		0.99 : (+ 0.32, + 3.21)
\bar{T}_f of concurrent TCP connections [1 - α : conf()]		0.99 : (- 0.16, - 0.14)

Table B.2: Statistical evaluation of the simulation results of Scenario 2 (10 ms router-router delay, unreliable last hop)

		no EFCM ↔ EFCM
\bar{T}_1 of new TCP connections [1 - α : conf()]	TCP 1	0.99 : (- 7.62, - 4.58)
	TCP 2	0.99 : (-23.25, -15.76)
	TCP 3	0.99 : (-23.18, -15.45)
\bar{T}_2 of a new TCP connection [1 - α : conf()]	TCP 1	0.99 : (- 7.60, - 3.02)
	TCP 2	0.99 : (- 9.09, - 3.11)
	TCP 3	0.99 : (- 9.60, - 3.39)
\bar{T}_1 of concurrent TCP connections [1 - α : conf()]		0.99 : (-13.88, - 9.63)
\bar{T}_2 of a concurrent TCP connection [1 - α : conf()]		0.99 : (- 7.52, - 4.20)
\bar{T}_f of concurrent TCP connections [1 - α : conf()]		0.99 : (- 0.03, - 0.02)

Table B.3: Statistical evaluation of the simulation results of Scenario 3 (30 ms router-router delay, reliable last hop)

		no EFCM ↔ EFCM
\bar{T}_1 of new TCP connections [1 - α : conf()]	TCP 1	0.99 : (- 6.99, - 4.83)
	TCP 2	0.99 : (-10.93, - 4.20)
	TCP 3	0.99 : (-11.88, - 4.62)
\bar{T}_2 of a new TCP connection [1 - α : conf()]	TCP 1	0.99 : (- 7.76, - 7.00)
	TCP 2	0.99 : (-10.59, - 6.66)
	TCP 3	0.99 : (-11.28, - 7.79)
\bar{T}_1 of concurrent TCP connections [1 - α : conf()]		0.99 : (- 5.15, - 2.59)
\bar{T}_2 of a concurrent TCP connection [1 - α : conf()]		0.99 : (- 4.25, - 2.64)
\bar{T}_f of concurrent TCP connections [1 - α : conf()]		0.99 : (- 0.08, - 0.07)

Table B.4: Statistical evaluation of the simulation results of Scenario 4 (30 ms router-router delay, unreliable last hop)

		no EFCM ↔ EFCM
\bar{T}_1 of new TCP connections [1 - α : conf()]	TCP 1	0.99 : (- 5.31, - 3.71)
	TCP 2	0.99 : (-16.15, -12.65)
	TCP 3	0.99 : (-16.37, -12.70)
\bar{T}_2 of a new TCP connection [1 - α : conf()]	TCP 1	0.99 : (- 5.36, - 3.98)
	TCP 2	0.99 : (- 8.70, - 6.83)
	TCP 3	0.99 : (- 7.90, - 4.73)
\bar{T}_1 of concurrent TCP connections [1 - α : conf()]		0.99 : (-10.43, - 8.67)
\bar{T}_2 of a concurrent TCP connection [1 - α : conf()]		0.99 : (- 6.52, - 5.27)
\bar{T}_f of concurrent TCP connections [1 - α : conf()]		0.99 : (- 0.04, - 0.03)

Table B.5: Statistical evaluation of the simulation results of Scenario 5 (50 ms router-router delay, reliable last hop)

		no EFCM ↔ EFCM
\bar{T}_1 of new TCP connections [1 - α : conf()]	TCP 1	0.99 : (- 6.04, - 4.50)
	TCP 2	0.99 : (-11.93, - 7.58)
	TCP 3	0.99 : (- 9.88, - 5.60)
\bar{T}_2 of a new TCP connection [1 - α : conf()]	TCP 1	0.99 : (- 6.61, - 6.12)
	TCP 2	0.99 : (-11.15, - 9.13)
	TCP 3	0.99 : (-11.25, - 7.67)
\bar{T}_1 of concurrent TCP connections [1 - α : conf()]		0.99 : (- 4.42, - 1.97)
\bar{T}_2 of a concurrent TCP connection [1 - α : conf()]		0.99 : (- 3.84, - 2.41)
\bar{T}_f of concurrent TCP connections [1 - α : conf()]		0.99 : (- 0.09, - 0.07)

Table B.6: Statistical evaluation of the simulation results of Scenario 6 (50 ms router-router delay, unreliable last hop)

		no EFCM ↔ EFCM
\bar{T}_1 of new TCP connections [1 - α : conf()]	TCP 1	0.99 : (- 4.46, - 3.27)
	TCP 2	0.99 : (-12.63, - 9.64)
	TCP 3	0.99 : (-13.09, -10.42)
\bar{T}_2 of a new TCP connection [1 - α : conf()]	TCP 1	0.99 : (- 3.42, - 2.69)
	TCP 2	0.99 : (- 7.55, - 6.14)
	TCP 3	0.99 : (- 8.00, - 5.95)
\bar{T}_1 of concurrent TCP connections [1 - α : conf()]		0.99 : (- 7.88, - 6.66)
\bar{T}_2 of a concurrent TCP connection [1 - α : conf()]		0.99 : (- 5.39, - 4.27)
\bar{T}_f of concurrent TCP connections [1 - α : conf()]		0.99 : (- 0.04, - 0.03)

B.3 Summary of the statistical evaluation

The statistical evaluation of the simulation results show for the throughput metrics rather wide confidence intervals. For new short TCP connections the EFCM controller reaches significantly higher throughputs in all considered simulation scenarios. With one exception for the less important connection-oriented mean throughput in Scenario 1 (this might be only a random effect), for new WWW or concurrent TCP connections the EFCM controller reaches a significantly higher or at least not a significantly lower throughput compared to standard TCP. And even for some results where no statistically significant difference between the standard TCP and the EFCM controller can be observed, a trend to the EFCM controller can be noticed.

In all simulation scenarios the fairness index of concurrent EFCM-controlled TCP connections is significantly higher than the fairness index of concurrent standard TCP connections.

Appendix C

Current developments of the EFCM controller

The EFCM controller is under development. The current version of the EFCM controller (date: 10. June 2002) has a built-in rate pacing algorithm that prevents the bursty sending behavior of new and concurrent TCP connections in an ensemble. In addition, a joint ack-clocking mechanism for all TCP connections of an ensemble is implemented that allows a more fair proportioning of the number of sent but currently not acknowledged segments (i.e., the in-flight segments of an ensemble) on the TCP connections in an ensemble.

Acknowledges

The work presented in this technical report has been partially supported by a research contract with Ericsson Eurolab, Aachen, Germany.

Bibliography

- [1] L. Eggert, T. Henderson, and J. Touch. Effects of ensemble TCP. *ACM SIGCOMM Computer Communication Review*, 30(January):15–29, 2000.
- [2] K. Fall and K. Varadhan. *The ns Manual*. VINT Project, 2001.
- [3] S. Floyd and T. Henderson. The new reno modification to TCP’s fast recovery algorithm. RFC 2582, 1999.
- [4] R. Jain. *The Art of Computer Systems Performance Analysis*. Wiley & Sons, 1991.
- [5] R. Jain. Congestion control and traffic management in ATM networks: Recent advances and a survey. *Computer Networks and ISDN Systems*, 28(13):1723–1738, 1996.
- [6] K. Mehlhorn. *Data Structures and Efficient Algorithms*. EATCS Monographs. Springer Verlag, 1984.
- [7] A. Reyes-Lecuona, E. González-Parada, E. Casilari, and J. Casasola. A page-oriented www traffic model for wireless system simulations. In *Proceedings ITC 16*, pages 1271–1280, 1999.
- [8] M. Savorić. Identifying and evaluating the potential of reusing network information from different flows. Technical report, TKN, 2001.
- [9] M. Savorić. The TCP control block interdependence in fixed networks — some performance results. In *Proceedings QOFIS 2001*, LNCS 2156, pages 261–272, 2001.
- [10] M. Savorić, H. Karl, and A. Wolisz. The TCP control block interdependence in fixed networks — new performance results. *to be published in Computer Communications*, 2002.
- [11] J. Touch. TCP control block interdependence. RFC 2140, 1997.