

Parallel TCP Sockets:
Simple Model, Throughput and Validation

Eitan Altman (INRIA), Dhiman Barman (INRIA),
Bruno Tuffin (IRISA/INRIA), Milan Vojnović (MSRC)

October 2005

Technical Report
MSR-TR-2005-130

Microsoft Research
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052

<http://www.research.microsoft.com>

Abstract— We consider a simple model of parallel TCP connections defined as follows. There are N connections competing for a bottleneck of fixed capacity. Each connection is assumed to increase its send rate linearly in time in absence of congestion indication and otherwise decreases its rate to a fraction β of the current send rate. Whenever aggregate send rate of the connections hits the link capacity, a single connection is signalled a congestion indication. Under the prevailing assumptions, and assuming only in addition a mild stability condition, we obtain that the throughput is the factor of the link capacity, $1 - 1/(1 + \text{const } N)$, with $\text{const} = (1 + \beta)/(1 - \beta)$. This result appears to be previously unknown; despite simplicity of its final form, it is not immediate. The result is of practical importance as it elucidates the throughput of parallel TCP sockets, an approach used widely to improve throughput performance of bulk data transfers (e.g. GridFTP), in regimes when individual connections are none or weakly synchronized. We argue that it is important to distinguish two factors that contribute to TCP throughput deficiency (F1) TCP window synchronization and (F2) TCP window adaptation in congestion avoidance. Our result is a good news as it suggests that in regimes when (F1) does not hold, already a few sockets are enough to almost entirely eliminate the deficiency due to (F2). Specifically, the result suggests that already 3 TCP connections yield 90% link utilization and 95% is almost achieved by 6 connections. This analytically proven result should provide incentive to throughput-greedy users to limit the number of their parallel TCP sockets as a few connections already ensure effectively 100% utilization, and any additional connection would provide only a marginal throughput gain. Opening too many sockets is not desirable as such transfers may beat down other connections sharing a link on the path of this transfer. However, there still remains a throughput deficiency due to (F1), which may provide incentive to users to open more sockets. The result found in this paper suggests that throughput-deficiency of parallel TCP sockets would be largely attributed to the synchronization factor (F1) and not to window control (F2). This motivates intelligent queueing disciplines that help mitigating the synchronization. As a by-product, the result shows that emulation of parallel TCP connections by MultTCP protocol is a good approximation. The implication of the result is that aggregate throughput achieved by connections is insensitive to a choice of loss policy which connection is signalled a congestion indication at congestion events. This perhaps somewhat counter-intuitive throughput insensitivity property is showed not to hold for steady-state distribution of the aggregate send rate. We provide validation of our results by simulations and internet measurements.

I. INTRODUCTION

Parallel TCP sockets is a generic “hack” to improve throughput attained by TCP for bulk data transfers by opening

several TCP connections and striping the data file over them. There are two factors that cause TCP throughput deficiency (F1) TCP window synchronization and (F2) window adaptation in TCP congestion avoidance. The former is well known and fixes have been proposed a while ago (e.g. RED [12]). The throughput deficiency due to (F2) is intrinsic to congestion window control in TCP congestion avoidance and is also well known. It is perhaps best illustrated by admitting a simple model of a single TCP connection over a bottleneck whose steady-state window dynamics is deterministic and periodic described as follows. In a period, the window W increases linearly in time with rate 1 packet per round-trip time and is reduced to one half of its current level when it encounters congestion. It is readily computed that such idealized TCP connection attains 75% link utilization. The simple model is the underlying model of a particular TCP square-root loss-throughput formula, as found in [11], [18]. The concept of parallel TCP sockets is widely used by applications such as for example GridFTP*. In practice, it is often unclear how many sockets one needs to open in order to achieve satisfactory throughput. Opening too many connections may be undesirable for various reasons and it is thus of importance to understand opening how many connections is sufficient to yield good throughput performance. Relation of TCP throughput of a single connection to loss event rate and mean round-trip time is to date fairly understood (e.g. [4], [18], [19]), but the results do not readily extend to parallel TCP connections. These formulae only leave us with two other unknowns. The problem is that it is not well understood how these two unknown parameters depend on the number of TCP connections competing in a bottleneck. The approach leveraging on known TCP loss-throughput formulae was pursued by Hacker et al [14], which provides very informative measurement results, but no throughput formula for parallel TCP connections.

It is important to distinguish the factors (F1) and (F2) as they are different throughput-deficiency causes. Our analysis is concerned with factor (F2). It is a good news as it suggests that already a few connections fixes the problem, in operational regimes when there is no synchronization (i.e (F1) not true) or there is some but is weak.

In this paper, the send rate of each connection is assumed to increase linearly in time in absence of congestion indication to this connection, and otherwise decrease to a fraction β of the send rate just before a congestion indication (multiplicative decrease). We adopt a model of bottleneck originally introduced by Baccelli and Hong [9] where link congestion events occur whenever the aggregate rate of connections hits the link capacity. Strictly speaking, the model assumes zero-buffer; see [9] for discussion how this can be relaxed to accommodate a finite buffer. The zero-buffer assumption is made for tractability and would be justified in scenarios when propagation delay dominates the queueing delay. We provide validation of this in Section V. We assume

(ONE) At each congestion event, a single connection is

*www.globus.org/datagrid/gridftp.html




Fig. 1. Sample paths of 3 AIMD connections competing for a link.

signalled to reduce its send rate by one multiplicative decrease.

Our main result is a formula for aggregate throughput $\bar{x}(N)$ of N connections over a link of fixed capacity c ,

$$\bar{x}(N) = c \left(1 - \frac{1}{1 + \frac{1+\beta}{1-\beta}N} \right) \quad (1)$$

where $0 < \beta < 1$; for TCP connections $\beta = 1/2$. The result is *exact* under a mild “stability” condition (see Section II, Theorem 1). Note that for the special case of one connection, $N = 1$, the result reads $\bar{x}(N) = c(1 + \beta)/2$. Hence, for TCP-like setting $\bar{x}(N) = (3/4)c$, which recovers the 75% utilization pointed out earlier. Note that nothing is said on which connection is chosen at link congestion events. Figure 1 illustrates several system assumptions that we made so far: linear increase and multiplicative decrease, link congestion times, and assumption (ONE). The assumption (ONE) was made in some other works, e.g. [10]. It would be more general to assume that at each link congestion event a subset of connection is chosen to undergo a multiplicative decrease. On the other extreme, we may assume that at each link congestion event, ALL connections undergo multiplicative decrease. This corresponds to TCP window synchronization. Baccelli and Hong [9], in turn, assume each connection is chosen with some fixed probability, so that at each congestion event a fraction of connections is signalled a congestion indication. The reader should bear in mind that (ONE) and similar assumptions are mathematical abstractions. In practice, though the system dynamics is blurred by various factors such as packetization and feedback delays and there will always be some positive time between any two connections undergoing multiplicative decrease, so it may not be justified to aim at validating (ONE) in some strict sense. The definition of link congestion times would in practice accommodate TCP connections over a network interface with finite buffer and FIFO queueing discipline (DropTail), with buffer length sufficiently smaller than bandwidth-delay product. In practice, most network interfaces are DropTail. The linear increase of send rate would approximate well TCP [15] in regime where round-trip time is dominated by propagation delay and queueing delay can be neglected [9]. It is well known that with substantial queueing

delay, TCP window growth over time is sub-linear for large windows [6]. The assumption that send rate increase is linear over time, may be satisfied by transport control protocols other than standard TCP, e.g. [17] shows this to be true for Scalable TCP [16].

In our analysis we also assume:

- (R) TCP connections are not constrained by their receiver windows.

With TCP receiver window in effect, the throughput would increase roughly linearly with the number of parallel TCP sockets for a range from 1 to some number of TCP connections. TCP receiver window constraint is not further considered in this paper. In analysis we assume (R) holds and in experimental validations we set TCP receiver window larger than the bandwidth-delay-product whenever possible.

We argue that it is of interest to understand throughput efficiency of parallel TCP sockets for a small number of parallel sockets. Opening too many sockets may not be of practical interest as this imposes processing overhead on end-host operating systems and also leaves smaller bandwidth to other flows that may sporadically share the bottleneck during a parallel-socket data transfer. The result (1) suggests that the throughput gain of N parallel TCP sockets monotonically increases with N and this increase is such that already 3 sockets is sufficient to attain 90% link utilization. Setting $N = 6$ would yield almost 95%. This suggests that a few connections is already enough to compensate for the throughput deficiency of TCP. This is a good news as it would provide incentive to rational throughput-greedy users not to use an overly large number of parallel sockets per transfer, as this would yield only a marginal throughput gain and as argued earlier opening too many sockets may not be desirable.

There has been a recent surge of interest on TCP and small buffer sizes [8]. While results in this paper do have some connection to this line of the work, as our analysis would be accurate for appropriately chosen small buffer size, there are notable differences. The results in [8] suggest that buffers in routers can be set to smaller values than bandwidth-delay product (BDP) for a large number of TCP connections and still have full link utilization. More specifically, it is suggested that for a link shared by N connections and link capacity scaled linearly with N , it is sufficient to have buffer length of the

order BDP/\sqrt{N} , asymptotically for large N . The throughput formula (1) is for any finite N and as argued earlier it is our particular interest to address the case when the number of connections N is not too large. In our context, it may well happen that bottleneck link is in access network, not a backbone router, and it may simply turn out to be small relative to propagation round-trip time of a network path between end-points of a parallel transfer. This would hold in practice for many end-to-end paths. Consider a bottleneck link with buffer size L packets, packets of length MTU B , and link capacity c MB/s. Then, we may regard the buffer to be small for network paths of propagation round-trip time RTT such that $\text{RTT} \gg L \text{ MTU}/c$. For example, consider in practice typical values $c = 100$ Mb/s, $\text{MTU}=1500$ B, and $L=100$ packets. Then, the condition is $\text{RTT} \gg 12$ ms, which would be true for many wide-area data transfers.

The result (1) tells us more. It reveals that throughput is insensitive to a choice of loss policy that decides which connection is signalled per congestion event as long as the loss policy verifies (ONE) and stability conditions in Theorem 1. Note that the set of loss policies that verify (ONE) is large and includes many natural candidates for loss policies. This insensitivity result may appear counter-intuitive to some as one's intuition may be that loss policies that favour picking connections with larger rate at congestion times would yield smaller aggregate throughput. The result says that this is not true. The throughput insensitivity in the present form was suggested by a result in [5], which is under more restrictive assumptions of only 2 connections and special loss policies that render send rates a Markov process. One may wonder whether steady-state probability distribution, and in particular, higher-order moments, of aggregate send rate are insensitive to a choice of loss policy. We show that this is not true in Section III. The result in Section III is not new as it was already found in [5], however, Section III provides corrected versions with complete proofs. The results in Section III and Section IV are particular choices of loss policy: (i) rate-proportional (= a connection is chosen with probability proportional to its send rate just before the congestion event), (ii) largest-rate (= a connection with largest rate just before the congestion event is chosen), (iii) rate-independent (= a connection is chosen with fixed probability), and (iv) beatdown (= a connection is chosen as long as permitted by link capacity). The rate-proportional loss policy may be deemed a natural candidate and may be seen as randomized version of the largest-rate policy. We consider the rate-independent policy in view of past work [9]. The fancy beatdown loss policy is considered only to demonstrate that throughput insensitivity found in this paper holds for a versatile set of loss policies. Note that in the Internet either of these loss policies or some other policies may emerge as a result of the dynamics of TCP congestion window control and particular queue discipline such as for example DropTail. More radical approach would be to enforce assumption (ONE) and particular loss policy by an intelligent queue discipline. Analysis of benefits and performance of such intelligent queue disciplines is beyond the scope of this paper.

A. Organization of the paper

Section II contains our main result. In that section, we provide a sufficient set of assumptions and notations and then state the main result, along with discussion of its implications. Proofs are deferred to Appendix. Section III shows that steady-state distribution of the aggregate send rate of connections is not insensitive to loss policy of the bottleneck link. Several auxiliary results are provided in Section IV. Section V contains experimental validation, both simulations and internet measurements. The last section concludes the paper.

II. MAIN RESULT

A. Model and result

We consider N connections that compete for a link of capacity c by adapting their send rates as follows. Let $X_i(t)$ be the send rate of a connection i at time t . A time t' is said to be a link congestion time whenever it holds $\sum_{i=1}^N X_i(t') = c$, i.e. the aggregate send rate hits the link capacity. Let T_n be the n -th link congestion time enumerated such that $T_0 \leq 0 < T_1$, for some time instant labelled with 0. A connection i increases its send rate linearly in time with rate $\eta_i > 0$ and it is reduced at some link congestion times to a fraction $0 < \beta_i < 1$ of the current send rate. We call such an adaptive connection, additive-increase and multiplicative-decrease (AIMD) connection with additive-increase parameter η_i and multiplicative-decrease parameter β_i .

The foregoing verbal definitions can be formally rephrased as, for $i = 1, 2, \dots, N$,

$$X_i(t) = X_i(T_n) + \eta_i(t - T_n), \quad T_n \leq t < T_{n+1} \quad (2)$$

and

$$X_i(T_{n+1}) = (1 - (1 - \beta_i)z_i(n))X_i(T_n) + \eta_i S_n \quad (3)$$

with

$$S_n = \frac{1}{\bar{\eta}} \left(c - \sum_{j=1}^N (1 - \beta_j)z_j(n)X_j(T_n) \right), \quad (4)$$

where $S_n := T_{n+1} - T_n$ is obtained from $\sum_{i=1}^N X_i(T_{n+1}) = c$, $\bar{\eta} := \sum_{j=1}^N \eta_j$, and $z_j(n) = 1$ if the connection j is signalled a congestion indication at the n -th link congestion event, else, $z_j(n) = 0$. Under (ONE) $\sum_{j=1}^N z_j(n) = 1$. We are now ready to




Fig. 2. Throughput of N identical AIMD connections competing for a link of unit capacity with TCP-like setting ($\beta_1 = 1/2$).

state our main result.

Theorem 1: Consider the prevailing system of N AIMD connections with identical AIMD parameters, competing for a link of capacity c . Assume that (ONE) at each link congestion event exactly one connection is signalled a congestion indication. Assume in addition that the following stability conditions hold: the system has a unique time-stationary distribution and times between successive link congestion events, sampled at congestion events, have a strictly positive mean. Then, the throughput X^{ONE} has the time-stationary expected value

$$\frac{\mathbb{E}(X^{\text{ONE}}(0))}{c} = 1 - \frac{1}{1 + \frac{1+\beta_1}{1-\beta_1}N}. \quad (5)$$

The proof is given in Appendix A. We deliberately deferred the proof to the appendix, so as to discuss the implications of the result first and leave the proof isolated only to the interested reader.

B. Discussion of the result and its implications

We discuss several implications of Theorem 1.

1) *Throughput of parallel TCP sockets:* The throughput formula (5) suggests explicit dependence of the throughput on the number of parallel TCP sockets N . Indeed, the throughput monotonically increases from $(1 + \beta_1)/2c$ for $N = 1$ to 1 as N tends to infinity. The former case of one connection is precisely the known result $(3/4)c$ that hold for TCP-like setting $\beta_1 = 1/2$. The dependence on the number of sockets given by (5) is showed in Figure 2 for $\beta_1 = 1/2$. The result suggests that opening as few as 3 connections would already yield 90% link utilization. Further, opening 6 sockets yields almost 95% link utilization. Hence, opening a number of sockets beyond some relatively small value

would yield only a marginal throughput gain. Equation (5) gives explicit dependence of the throughput on the number of sockets N , and thus provides a guideline to choose the magic configuration parameter of the number of sockets for a data transfer. Choosing this configuration parameter in practice was a puzzle as throughput dependence on the number of sockets was largely not understood. The result also suggests incentive to throughput-greedy users not to open too many sockets, as opening a few is enough.

2) *Throughput is insensitive:* No matter which connection is signalled a congestion indication per link congestion event, the mean throughput remains the same, as long as the mild stability conditions of the theorem hold. This throughput insensitivity finding may appear counter-intuitive as one may expect that loss policies that favour signalling congestion indication to large-rate connections would have detrimental effect on the mean aggregate throughput. The result reveals this is not true. We show in Section III that the aggregate send rate higher-order statistics is not insensitive on the loss policy.

The throughput insensitivity in the general form found in this paper was motivated by original finding in [5]. Note that we generalize the finding of [5] in the following directions: (i) to any finite number of connections N ([5] is for $N = 2$) and (ii) a larger set of loss policies. We discuss (ii) in some more detail now. The loss policies in [5] are restricted to those that base decision which connection to signal a congestion indication only on the send rates just before a link congestion event. As an aside, note that this renders the send rates a Markov process. The result of Theorem 1 tells us that the throughput insensitivity holds for a much broader set of loss policies that can depend on any past system state as long as the mild stability conditions of the theorem hold.

Note that the throughput insensitivity allows to bias signalling congestion indication to some connections at link congestion times. The throughputs achieved by these special connections would indeed deteriorate by this bias, however, the aggregate throughput over all connections remains the same. This is illustrated by the following example.

Example 1: Consider 4 AIMD connections competing for a link of unit capacity. The connections have identical AIMD parameters, set as $\eta_i = 1$ and $\beta_i = 1/2$. The loss policy is rate-independent with the probability that a connection i is selected equal to $\sigma/2$, for $i = 1, 2$, and $(1 - \sigma)/2$, for $i = 3, 4$, for some $0 < \sigma < 1$. The parameter σ is the bias parameter that for $\sigma < 1/2$ favours picking either connection 1 or 2, $\sigma = 1/2$ favours no connection, and $\sigma > 1/2$ favours picking connections 3 or 4. Hence, as σ ranges from 0 to 1, the connections 1 and 2 (resp. 3 and 4) will receive equal throughput that will decrease (resp. increase) with σ . In the prevailing example, the system does have a unique time-stationary distribution, hence the hypotheses of Theorem 1 are satisfied. Theorem 1 implies that the mean aggregate throughput is insensitive to the preferential treatment of connections. The simulation results in Figures 3 and 4 indeed validate the claim.




Fig. 3. Throughputs of 4 AIMD connections in Example 1. There are two classes of connections, each class containing 2 connections. A connection selected at a link congestion time is from the first class with probability σ . The flows of the two classes attain throughputs that depend on the parameter σ . However, the aggregate throughput is *insensitive* (dots aligned with the analytical prediction shown by the solid line). This validates Theorem 1.

3) *Comparison with the synchronized case:* Consider now the loss policy we call ALL, which signals all connections a congestion indication at each link congestion event. Note that this loss policy satisfies (ONE) only if $N = 1$. The loss policy ALL models the special regime of synchronized TCP connections. A natural question is to compare the throughputs under loss policy ALL and loss policies that verify (ONE). The throughput under loss policy ALL is readily obtained to be[†]

$$\frac{\mathbb{E}(X^{\text{ALL}}(0))}{c} = \frac{1 + \beta_1}{2}. \quad (6)$$

It is easy to check that

$$\mathbb{E}(X^{\text{ONE}}(0)) \geq \mathbb{E}(X^{\text{ALL}}(0))$$

with equality only if $N = 1$. Note that, indeed, the throughput achieved under the loss policy ALL for $N > 1$ is exactly that achieved by a single connection under either ALL or ONE. Thus, the ratio of the throughputs achieved under a loss policy that satisfies (ONE) and under the loss policy ALL corresponds to the throughput gain achieved under the loss policy that satisfies (ONE) for N connections with respect to the throughput for $N = 1$. The loss policies under (ONE) provide no gain with respect to the loss policy ALL for $N = 1$, but they provide larger throughput for a factor which is exactly that given by (5) divided by $(1 + \beta_1)/2$. The last means that loss policies under (ONE) provide throughput gain with respect the loss policy ALL equal to $2/(1 + \beta_1)$, as the number of connections N tends to infinity. For TCP-like

[†]This follows by the fact that under ALL there is a unique limit point x^* of the send rates just before a link congestion event such that $x_i^* = c/N$, $i = 1, 2, \dots, N$.




Fig. 4. Sample-paths of the rates of the 4 AIMD connections "orchestrated" by the rate-independent policy (Example 1). A connection is chosen from the probability distribution $[\sigma, \sigma, 1 - \sigma, 1 - \sigma]/2$. (Left) $\sigma = 1/2$ so that each connection has equal chance to be selected. (Right) $\sigma = 1/8$ so that two connections are selected more often than the other two connections.

setting, this amounts to $4/3$, which is precisely compensating for the throughput deficiency of a single TCP connection.

4) *What this tells us about MultTCP?:* MultTCP, proposed by Oechlin and Crowcroft [10], is a window control protocol that aims at emulating a given number of N TCP connections. The design rationale is to scale the linear increase of the control with N as $N\eta_1$ and set the multiplicative decrease parameter to $1 - (1 - \beta_1)/N$. MultTCP can thus be seen as a single AIMD connection with appropriately set additive-increase and multiplicative-decrease parameters as a function of N . The choice of the multiplicative decrease parameter is based on an approximation by considering a virtual set of N parallel connections. It is assumed that at each congestion indication to these connections, only one connection undergoes a multiplicative decrease. Thus, it admits loss policies verifying (ONE). The send rate of this selected connection at a time just before a congestion indication is x/N , where x is the aggregate rate over all connections just before the congestion indication. The last is an approximation as, indeed, the selected connection may have the send rate other than the even share x/N . In the sequel, we consider MultTCP for TCP-like setting $\beta_1 = 1/2$. We already observed that MultTCP can be treated as a single AIMD connection with additive-increase parameter $N\eta_1$ and multiplicative-decrease parameter $1 - N/2$, hence, the throughput follows immediately from (6)

$$\frac{\mathbb{E}(X^{\text{MultTCP}}(0))}{c} = 1 - \frac{1}{4N}.$$

In turn, from (5), the throughput of N TCP-like connections is

$$\frac{\mathbb{E}(X^{\text{ONE}}(0))}{c} = 1 - \frac{1}{1 + 3N}.$$

We note that $\mathbb{E}(X^{\text{MultTCP}}(0)) \geq \mathbb{E}(X^{\text{ONE}}(0))$, for all $N \geq 1$, with equality only if $N = 1$. It is readily computed that the mean throughput of a MultTCP connection is less than 1.025

of the mean aggregate throughput of N TCP-like connections under a loss policy in ONE. It is hence, for any practical purpose, a good approximation.

The last observation does not provide much motivation to obtain a correction for MultTCP so that its throughput is exactly that of N TCP-like connections under a loss policy that verifies (ONE). Note, however, redefining only β_1 to a function of N provides no solution, except for the trivial case $N = 1$. Indeed, by equating (5) and (6) and a simple calculation one obtains $\beta_1(N-1) = N-1$, hence, a $0 < \beta_1 < 1$ exists only if $N = 1$.

5) *Round-trip time insensitivity*: The result (1) suggests insensitivity of the mean throughput to the round-trip time of the connections as (5) does not depend on η_1 . With TCP-like connections, the parameter η_1 will be defined as reciprocal of a fixed round-trip time to the power of 2.

III. STEADY-STATE DISTRIBUTION OF THE AGGREGATE SEND RATE IS NOT INSENSITIVE

The insensitivity property found in the earlier section is for the first-order moment of the time-stationary aggregate send rate. The time-stationary distribution is *not* insensitive to the loss policy. This was already showed in [5] by computation of the time-stationary second moments for particular policies. We provide a complete proof for the second moments, which yields corrected versions of the corresponding expressions in [5]. A new information is for the beatdown policy.

Theorem 2: Consider two identical AIMD connections with $\beta_1 = 1/2$. The time-stationary second moments for either flow are given by

- 1) **Beatdown**: $\frac{\mathbb{E}(X^{(1)}(0)^2)}{c^2} = \frac{1}{3} \approx 0.33333$
- 2) **Rate-independent**: $\frac{\mathbb{E}(X^{(1)}(0)^2)}{c^2} = \frac{5}{24} \approx 0.20833$
- 3) **Rate-proportional**: $\frac{\mathbb{E}(X^{(1)}(0)^2)}{c^2} = \frac{679}{3396} \approx 0.19994$
- 4) **Largest-rate**: $\frac{\mathbb{E}(X^{(1)}(0)^2)}{c^2} = \frac{4}{21} \approx 0.1905$.

The proof is a tedious exercise of elementary and Palm calculus and is deferred to Appendix B.

Note that the order of the loss policies with respect to the time-stationary second-order moment of the throughput is rather natural. Intuitively, one would expect that the largest-rate policy attains smallest second-order moment as it *balances* the rates at congestion times by picking a connection with largest rate. One also would expect that the rate-proportional policy has a larger second-order moment of the throughput, but not too much as it also favours selecting connections with larger rates at link congestion times. The rate-independent policy, with each connection selected equally likely, has a larger second-order moment of the throughput than both the largest-rate and rate-proportional. The beatdown policy attains the largest second-order moment of the throughput within the given set of policies.

IV. ADDITIONAL ANALYSIS FOR LARGEST-RATE AND BEATDOWN POLICIES

This section shows some auxiliary results on the dynamics for particular loss policies, the largest-rate and beatdown. The analysis provides some insight on the dynamics of the send rates and serve in the proof of the second-order moment results in the earlier section. It may be however skipped by a reader non interested in technical details, with no sacrifice of continuity.

A. Largest-rate policy

From now, we use the short-cut notation $x_i(n) := X_i(T_n)$. We show that for a homogeneous population of N AIMD connections, there exists a unique stationary point for the rate vector $x(n) = (x_1(n), \dots, x_N(n))$, under the largest-rate policy and identify this stationary point.

In view of (3) and (4), dynamics of the rates under the largest-rate policy is fully described by the following *deterministic* dynamical system:

$$x_i(n+1) = x_i(n) + \left(\frac{\eta_i}{\bar{\eta}} - 1_{\{M_n=i\}} \right) (1 - \beta_{M_n}) \bigvee_{j=1}^N \{x_j(n)\}$$

where $M_n \in \arg\max_j x_j(n)$. Throughout, we use the operators $x \vee y := \max\{x, y\}$ and $x \wedge y := \min\{x, y\}$, for $x, y \in \mathbb{R}$.

This can be further rewritten as follows. Let $\tilde{x}(n)$ be the vector $x(n)$ with coordinates ordered in decreasing order, for any given n . Thus, $\tilde{x}_1(n) \geq \tilde{x}_2(n) \geq \dots \geq \tilde{x}_N(n)$, for all n . We can write:

$$x_i(n+1) = \begin{cases} a_i \tilde{x}_1(n), & M_n = i \\ x_i(n) + b_i \tilde{x}_1(n), & M_n \neq i, \end{cases} \quad (7)$$

where $a_i := \beta_i + \eta_i(1 - \beta_i)$ and $b_i := \eta_i(1 - \beta_i)$, $i = 1, 2, \dots, N$.

In the remainder of this section, we confine to homogeneous flows, thus $a_i = a_1$ and $b_i = b_1$, for all $i = 2, 3, \dots, N$. Without loss of generality, we assume $c = 1$. Thus, for any n , $x(n)$ takes value on a $N-1$ dimensional simplex $\mathcal{S}_{N-1} = \{x \in \mathbb{R}_+^N : \sum_{i=1}^N x_i = 1\}$ and $\tilde{x}(n)$ on the subset of \mathcal{S}_{N-1} , $\mathcal{S}_{N-1}^0 = \{x \in \mathbb{R}_+^N : \sum_{i=1}^N x_i = 1, x_i \geq x_j, i \leq j\}$.

Lemma 1 (representation): Consider a homogeneous population of N AIMD users competing for a link under the largest-rate loss policy. The dynamics of the ordered vector of rates $\tilde{x}(n)$ is given by the recurrence

$$\tilde{x}(n+1) = g(\tilde{x}(n)), \quad (8)$$

$$\text{where } \begin{array}{ccc} \mathcal{S}_{N-1}^0 & \mapsto & \mathcal{S}_{N-1}^0 \\ x & \rightarrow & g(x) \end{array}$$

with $g(\cdot) = (g_1(\cdot), g_2(\cdot), \dots, g_N(\cdot))$, and

$$g_i(x) = \begin{cases} a_1 x_1 \vee (x_2 + b_1 x_1), & i = 1 \\ (a_1 x_1 \vee (x_{i+1} + b_1 x_1)) \wedge (x_i + b_1 x_1), & \text{else} \\ a_1 x_1 \wedge (x_N + b_1 x_1), & i = N. \end{cases}$$

This lemma is proved in Appendix C.

As an aside technical remark, note that the mapping $x \rightarrow g(x)$ involves min, max, and plus operators, but is *not* a




Fig. 5. Sample-paths of the rates of two AIMD flows "orchestrated" by the beatdown policy.

min-max mapping (see Gunawardena [13] and the references therein). In particular, the mapping is not *homogeneous*, i.e. $g(x+h) \neq g(x)+h$, for all $h \in \mathbb{R}$, a property enjoyed by min-max mappings.

Proposition 1: The map $g : \mathcal{S}_{N-1} \rightarrow \mathcal{S}_{N-1}$ of the recurrence (8) has a unique fixed point x^0 in the interior of \mathcal{S}_{N-1} (all coordinates strictly positive) given by

$$x_i^0 = \begin{cases} \frac{2}{(1+\beta_1)^{N+1} + (1-\beta_1)} & i = 1 \\ x_1^0 \left(1 - \frac{1-\beta_1}{N}(i-1)\right) & i = 2, 3, \dots, N. \end{cases} \quad (9)$$

The proof is deferred to Appendix D.

Remark 1: The fixed point uniquely identifies a N -periodic sequence $x^0(n)$ in the simplex \mathcal{S}_{N-1} . We readily construct a continuous-time function $x^0(t)$ with values in \mathcal{S}_{N-1} associated to the sequence $x^0(n)$ by the definition of the inter-congestion times (4), with $T_0 = 0$. As customary, a time-stationary sample-path is constructed by shifting the T -periodic function $x^0(t)$ for a time uniformly at random on the interval $[0, T]$.

Remark 2: The proposition says that the stationary point of the rates just before link congestion event is such that that connections are signalled a congestion indication in a round-robin fashion.

B. Beatdown

It is particular to the beatdown loss policy that congestion times accumulate, i.e. the number of congestion events on some finite time intervals is infinite (see Figure 5). This implies that $\mathbb{E}^0(T_1) = 0$. This prevents us from applying Theorem 1. In this section we circumvent this problem by looking at some special link congestion times, and show that indeed the throughput-insensitivity property holds.

Define τ_n as a link congestion time at which a connection becomes tagged according to the beatdown policy. Let $y_i(n)$ be the rate of a flow i at time τ_n (subsequence of $x_i(n)$). With an abuse of notation, let $z_i(n)$ indicates whether a flow i was selected at time τ_n . The rates embedded at the selected link congestion times evolve as follows:

$$y(n+1) = (1 - z_i(n))(y_i(n) + \eta_i(\tau_{n+1} - \tau_n)). \quad (10)$$

Indeed, if a flow i is tagged at τ_n , i.e. $z_i(n) = 1$, then $y_i(n+1) = 0$. Else, $z_i(n) = 0$, and thus the flow continues increasing its rate with rate η_i from the initial rate $y_i(n)$.

Summing $\sum_{j=1}^N y_j(n) = c$, yields

$$\tau_{n+1} - \tau_n = \frac{\sum_{j=1}^N z_j(n)y_j(n)}{\sum_{j=1}^N (1 - z_j(n))\eta_j}. \quad (11)$$

Finally, we have:

$$y_i(n+1) = (1 - z_i(n)) \left(y_i(n) + \frac{\eta_i}{\bar{\eta} - \eta_{M_n}} y_{M_n}(n) \right),$$

where, recall, $M_n = i$ for i such that $z_i(n) = 1$.

We have the following result whose proof is in Appendix E.

Theorem 3: For a finite and homogeneous population of AIMD connections, any time-stationary beatdown loss policy is throughput-insensitive.

In the sequel of this section, we give the stationary point for the flow rates at selected congestion times, for a homogeneous population of AIMD connections. Let $\tilde{y}(n)$ be the vector of the rates $y(n)$ sorted in decreasing order. The dynamics is as follows:

$$y_i(n+1) = \begin{cases} 0 & M_n = i \\ y_i(n) + \frac{1}{N-1} \tilde{y}_1(n) & \text{otherwise.} \end{cases}$$

The stationary point satisfies the following identities: $\tilde{y}_k = \tilde{y}_{k+1} + \frac{1}{N-1} \tilde{y}_1$, $k = 1, 2, \dots, N-1$. It follows

$$\tilde{y}_k = \frac{N-k}{N-1} \tilde{y}_1.$$

Now, $\sum_{i=1}^N \tilde{y}_i = c$, from which it follows $\tilde{y}_1 = 2c/N$. In summary,

$$\frac{\tilde{y}_i}{c} = \begin{cases} \frac{2}{N} & i = 1 \\ \frac{1}{N-1} \left(1 - \frac{i}{N}\right) & i = 2, 3, \dots, N-1 \\ 0 & i = N. \end{cases}$$

V. EXPERIMENTAL RESULTS

The results in earlier sections are exact for AIMD connections and the prevailing assumptions therein. In this section, we provide experimental validation of our analysis by simulations and internet measurements. Simulation and measurements scripts and data are available on-line at [1].

We performed ns-2 simulations as this allows us to easily configure queueing disciplines in order to conduct controlled experiments. Our internet measurements provide validation in real-world scenarios.

A. Setup of ns-2 experiments

We describe first our ns-2 experiments. We consider a system of N identical TCP-SACK connections that traverse a bottleneck as showed in Figure 6. The bottleneck link implements either DropTail or RED queueing discipline. Payload of a TCP packet is set to 1000 bytes. Maximum window size is set to 20000 packets. Each simulation is run for 2000 seconds and initial 1000 seconds are truncated to reduce the effect of initial transient. Start times of connections are randomized by uniformly drawing the starting times within interval $[0, 4]$ seconds.

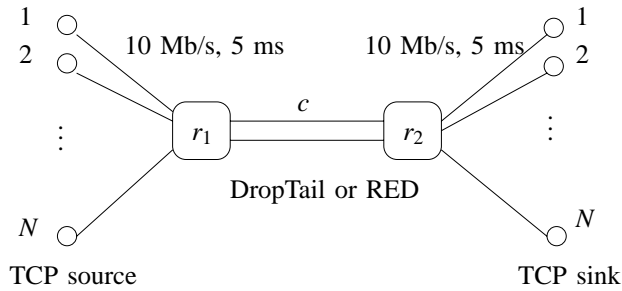


Fig. 6. ns-2 network configuration.




Fig. 7. Time-average send rate normalized with link capacity versus the number of TCP connections. Results are obtained in ns-2 for DropTail with varying buffer length b as indicated in the figure.

In some simulations we use RED bottleneck. Original design goal of RED was to mitigate TCP window synchronization [12]. RED parameters are set as follows. Minimum and maximum threshold sizes are set to 0.6 and 0.8 factor of the buffer length. The RED packet drop probability function is chosen with the value equal to 0.1 at the maximum buffer threshold. The queue size averaging factor is set to 0.001. The RED option wait was enabled; this would enforce some waiting between successive packet drops.

Theorem 1 suggests the following claim: (H) a system of N identical TCP connections that compete for a bottleneck with a small buffer attains link utilization $1 - 1/(1 + 3N)$.

B. DropTail

The results are shown in Figure 7. We observe that typically, for small buffer lengths, the achieved throughput is smaller than predicted by the analysis. Moreover, we note that the larger the buffer is, the larger the throughput is. If buffer length is sufficiently large, throughput is larger than predicted by the analysis. [TBD: discuss bandwidth-delay product!] We verified by inspection that in many cases in Figure 7, TCP




Fig. 8. Sum of congestion windows over all TCP connections for (top) $N = 5$ and (bottom) $N = 10$, competing for bottleneck link with propagation round-trip time 120 ms and buffer lengths as indicated in the figure. The results clearly indicate synchronization of TCP connections.




Fig. 9. Synchronization in DropTail bottleneck with $N = 10$ and $b = 5$.

windows are synchronized. This is exhibited for $N = 5$ and 10 in Figure 8. The figure shows sum of congestion windows over all connections versus time. The sum of congestion windows essentially evolves over time as congestion window of a single TCP connection, which indicates synchronization. This is further reflected in Figure 9 that shows congestion windows of individual connections.

C. RED

We show results in Figure 10. The results exhibit reasonable conformance to Theorem 1.

D. Loss polices that conform to ONE

We observed TCP window synchronization in experiments with DropTail. In order to validate Theorem 1, our aim is to design experiments in which assumption ONE holds, i.e. at




Fig. 10. Time-average send rate normalized with link capacity versus the number of TCP connections. Results are obtained in ns-2 for RED with varying buffer length as indicated in the figure.

each congestion event exactly one connection is signalled congestion. To that end, we implement a loss-policy—threshold-dropper—described as follows. The queue buffer length is set to virtually infinite value. The dropping is based on a finite positive threshold T_h and loss policy described as follows. Consider the system from a time at which there are at most T_h queued packets. The first packet arrival that exceeds threshold is dropped and the corresponding connection is declared selected. The selected connection remains selected for an interval of τ time units, where τ is set to a factor of the mean round-trip time. This process repeats. No packets are dropped during this period. It is clear from the description that threshold-dropper enforces assumption ONE. It also permits implementation of various loss policies.

Figure 11–left shows throughput achieved with threshold-dropper that implements rate-independent loss policy. We note a good qualitative conformance with Theorem 1, in particular when the number of connections is not too large. For sufficiently large number of connections, the throughput is essentially equal to the link capacity. In these case, the queue does not deplete, which explains larger throughput than predicted by the analysis. Analogous results for largest-rate loss policy are showed in Figure 11–right. The experimental results obtained for rate-independent and largest-rate policy match well, which indicates throughput insensitivity. This is in conformance to the analysis.

We further provide results that indicate threshold-dropper mitigates TCP window synchronization. Figure 12–top shows the congestion window evolution of 10 individual TCP connections. This is substantiated by looking at the sum of congestion windows in Figure 12–bottom. The alert reader will compare this with Figure 8.

1) *What loss policies emerge with DropTail and RED?:* It is often assumed that a connection observes loss events with




Fig. 12. (Top) Congestion windows of $N = 10$ TCP connections, (bottom) sum of all congestion windows. The loss policy is threshold-dropper with threshold $T_h = 3$. The results suggest absence of synchronization.

intensity in time proportional to send rate. This corresponds to rate-proportional loss policy introduced earlier. We investigate over a limited set of experiments whether or not rate-proportional loss policy emerges with DropTail and RED. To that end, we proceed with the following test. Time is binned into consecutive intervals of length T time units. A connection is assumed to receive a congestion indication in a bin n , if there is at least one packet drop for this connection in the n th bin. The window size of the connection i in bin n is denoted by $w_i(n)$. Our goal is to estimate $\mathbb{P}(z_i(n) = 1 | w_i(n) = w)$, where $z_i(n) = 1$ if the connection i is indicated congestion in bin n , else $z_i(n) = 0$. We estimate this conditional probability by the empirical mean

$$\bar{z}(w) := \frac{\sum_n \sum_{i=1}^N 1_{\{z_i(n)=1, w_i(n)=w\}}}{\sum_n \sum_{i=1}^N 1_{\{w_i(n)=w\}}}$$

which under stationary and ergodicity assumptions converges with probability 1 to $\mathbb{P}(z_i(n) = 1 | w_i(n) = w)$, as number of time bins tends to infinity. In Figure 13, we show the empirical mean $\bar{z}(w)$ versus w for systems of $N = 10$ TCP connections with propagation round-trip time of 120 ms and DropTail link for buffer sizes = 5, 10, 15, 20 packets. Same is showed in Figure 14, but for RED. Some of the results suggest the larger the congestion window of a connection just before a link loss event is, the more likely the connection is signalled a congestion indication.





Fig. 11. The ratio of the throughput and link capacity versus the number of connections N with threshold-dropper: (left) rate-independent, (right) largest-rate. The results conform well with Theorem 1 (solid line).

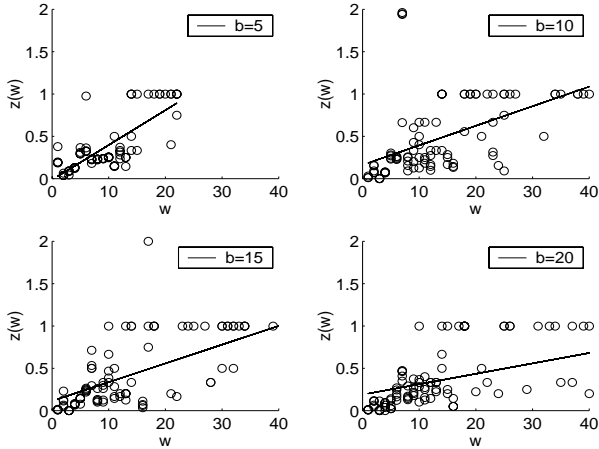


Fig. 13. Empirical estimate of $\mathbb{P}(z_i(n)|w_i(n) = w)$ versus the window size w for DropTail link and 10 TCP connections, for buffer length as indicated in the graphs.

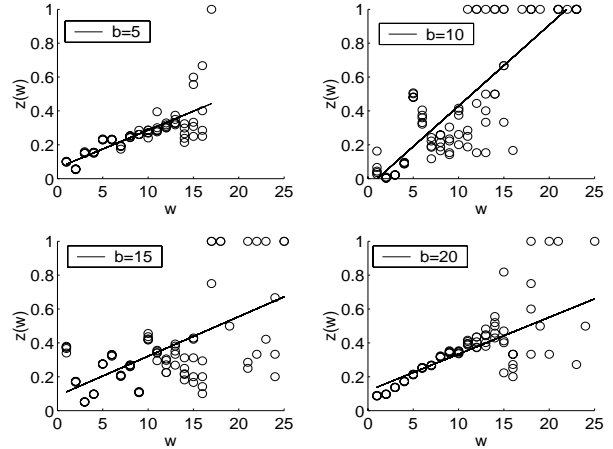


Fig. 14. Same as in Figure 13, but for RED.

E. Internet measurements

We now proceed with more realistic experiments by measurements on the internet planet-scale testbed, PlanetLab [2]. We choose randomly a set of distinct end-to-ends path from different continents such as Europe and North America. For each path, we conduct 10 repeated experiments. Each set of measurement starts with parallel TCP transfers with up to 50 parallel TCP connections. We measure the aggregate throughput obtained using the tool `iperf` [3] and round-trip time using `ping`. Some basic information about end-to-end paths is showed in Table I.

The measured throughputs are showed in Figure 15 and 16. In Figure 15, we note a very good match of measurements `inria-stanford` with Theorem 1. The measurements `fortiche-titeuf` suggests buffer saturation. All results

exhibit qualitative conformance with the throughput increase with the number of connections predicted by the analysis.

We next investigate whether connections are synchronized. On distinct end-to-end paths, we run `iperf` tool with 10 parallel TCP connections. We collect packets transmitted in both direction of a connection with `tcpdump` and then infer loss events. We clump loss events to congestion events based on a time threshold (a factor of the round-trip time) and count the number of distinct connections, M_n that suffer at least one loss event in a bin n . We plot the empirical complementary distribution function of M_n in Figure 17. We observe that M_n is a small fraction of $N = 10$ on almost all paths. Furthermore, we plot loss events over time in Figure 18-19 and observe that losses do not appear being clustered together, which suggests absence or weak TCP window synchronization.

TABLE I
SET OF END-TO-END PATHS USED IN MEASUREMENTS

path	Source host	Destination Host	RTT
umass-berkeley	planetlab1.cs.umass.edu	planetlab2.millennium.berkeley.edu	97ms
inria-stanford	planetlab3.inria.fr	planetlab-2.stanford.edu	194.7ms
cornell-greece	planetlab4-dsl.cs.cornell.edu	planet2.ics.forth.gr	337.9ms
LAN	einat.inria.fr	titeuf.inria.fr	1.6 ms




Fig. 15. Throughput measured on the internet with `iperf` with N parallel TCP connections. For each N , 10 distinct measurements were conducted. The figures show boxplots that indicate the measurement samples are concentrated around corresponding average values. Solid line is the result of Theorem 1 interpolated by fitting the measured value for $N = 50$ with the corresponding analytical value. The measurement results exhibit qualitative growth of the throughput with N as predicted by the analysis. (left) `umass-berkeley` (right) `inria-stanford`. The latter shows particularly good conformance with Theorem 1.




Fig. 16. Same as in Figure 15, but (left) `cornell-greece` (right) `fortiche-titeuf`.

VI. CONCLUSION

We obtained a simple and new throughput formula for parallel adaptive connections that adapt their send rates according to widely-known additive-increase and multiplicative decrease. The connections compete for a link of fixed capacity and link congestion times are defined by instants when the aggregate send rate hits the link capacity. The result is under assumption that at each link congestion event, exactly one connection undergoes a multiplicative decrease. This result

albeit being very simple in its final form is obtained by a non-trivial proof. Its practical implications are versatile. It suggests an explicit dependence of the throughput gain on the number of parallel TCP sockets. In particular, the result predicts that 90% utilization would already be achieved by as few as 3 sockets. It thus provides incentives to limit the number of parallel TCP sockets per data transfers, as beyond a relatively small number of sockets, any throughput gain becomes marginal. The result also reveals a somewhat counter-

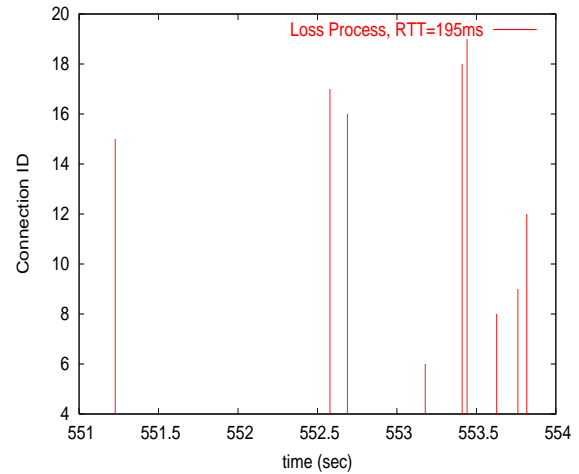
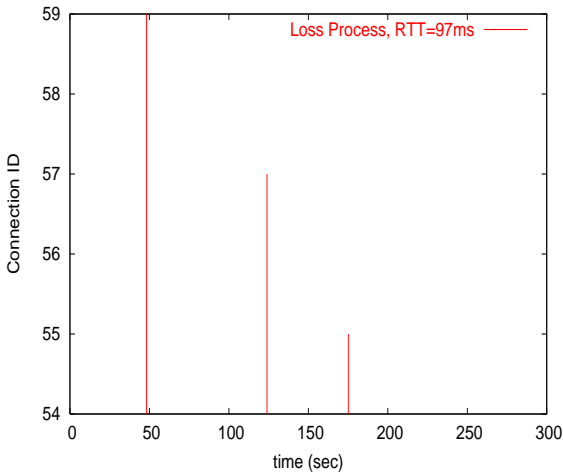


Fig. 18. Loss events: connection identifiers suffering loss are shown versus time where a connection identifier are least two significant digits of port number. (left) umass-berkeley (right) inria-stanford.

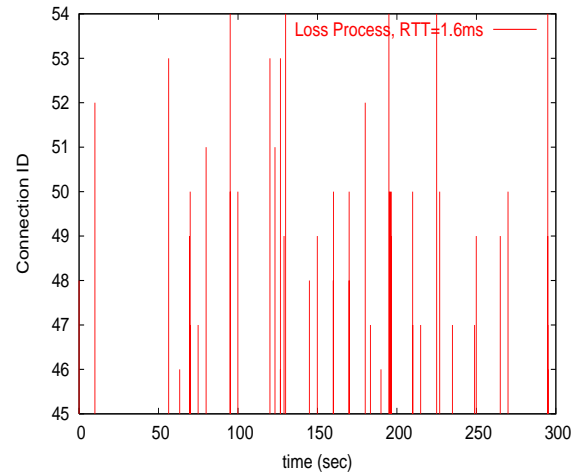



Fig. 19. Same as in Figure 18, but (left) cornell-greece (right) einat-titeuf.

intuitive throughput insensitivity property: no matter what loss policy that assigns congestion indications to connections is over a broad set of loss policies that select one connection per link congestion time, and a mild stability condition holds, the mean aggregate throughput remains the same. This is showed not to hold for the steady-state distribution of the aggregate send rate. The throughput formula is validated by simulations and internet measurements.

ACKNOWLEDGMENTS

E. A and B. T were in part supported by EuroNGI NoE and INRIA TCP cooperative research action. M. V. is grateful to Laurent Massoulié and Don Towsley for discussions of some technical points. He also thanks Dinan Gunawardena for his valuable suggestions on practical implications of the throughput-insensitivity found in this paper. Jon Crowcroft and Karl Jeacle pointed to the Grid application.

REFERENCES

[1] <http://research.microsoft.com/~milanv/socks.htm>.
 [2] <http://www.planet-lab.org>.

[3] <http://dast.nlanr.net/Projects/Iperf/>.
 [4] E. Altman, K. Avrachenkov, and C. Barakat. A stochastic model of tcp/ip with stationary random losses. In *Proc. of the Sigcomm'00*, pages 231–242, 2000.
 [5] E. Altman, R. El Azouzi, D. Ros, and B. Tuffin. Loss Strategies for Competing TCP/IP Connections. In *Proc. of the Networking 2004*, Athens, Greece, May 2004.
 [6] E. Altman, F. Baccara, J. Bolot, P. Nain, P. Brown, D. Collange, and C. Fenzy. Analysis of the tcp/ip flow control mechanism in high-speed wide-area networks. In *34th IEEE Conference on Decision and Control*, pages 368–373, New Orleans, Louisiana, Dec 1995.
 [7] E. Altman, T. Jiménez, and R. Núñez-Queija. Analysis of two competing tcp/ip connections. *Performance Evaluation*, 49(1):43–56, 2002.
 [8] G. Appenzeller, I. Keslassy, and N. McKeown. Sizing Router Buffers. In *Proc. of ACM Sigcomm 2004*, Portland, Oregon, USA, Aug/Sept 2004.
 [9] F. Baccelli and D. Hong. AIMD, Fairness and Fractal Scaling of TCP Traffic. In *Proc. of IEEE Infocom 2002*, New York, NY, <http://www.di.ens.fr/~trec/aimd>, 2002.
 [10] J. Crowcroft and P. Oechslin. Differentiated End-to-End Internet Services using a Weighted Proportional Fair Sharing TCP. *ACM Computer Communications Review*, 47(4):275–303, 2004.
 [11] S. Floyd and K. Fall. Promoting the use of end-to-end congestion control in the Internet. *IEEE/ACM Trans. on Networking*, 7(4):458–472, August 1999.
 [12] S. Floyd and Van Jacobson. Random early detection gateways for




Fig. 17. Complementary distribution function of the number of connection assigned a packet drop in an interval length of const \times RTT for 10 parallel TCP connections.

congestion avoidance. *IEEE/ACM Trans. on Networking*, 1(4):397–413, August 1993.

- [13] J. Gunawardena. From max-plus algebra to nonexpansive mappings: a nonlinear theory for discrete event systems. *Theoretical Computer Science*, 293:141–167, 2003.
- [14] T. Hacker, B. Athey, and B. Noble. The end-to-end performance effects of parallel TCP sockets on a lossy wide-area network. In *16th IEEE-CS/ACM International Parallel and Distributed Processing Symposium*, Ft. Lauderdale, FL, April 2002.
- [15] Van Jacobson and M.J. Karels. Congestion avoidance and control. In *Proc. of the ACM SIGCOMM'88*, pages 314–329, Stanford, August 1988.
- [16] T. Kelly. Scalable TCP Improving Performance in HighSpeed Wide Area Networks. In *PFLDNet'03*, Geneva, Switzerland, February 2003.
- [17] R. El Khoury and E. Altman. Analysis of Scalable TCP. In *In Proc. of Het-Nets'04*, West Yorkshire, United Kingdom, July 2004.
- [18] J. Mahdavi and S. Floyd. TCP-Friendly Unicast Rate-Based Flow Control. Technical note sent to end-2-end interest mailing list, http://www.psc.edu/networking/papers/tcp_friendly.html, January 1997.
- [19] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose. Modeling TCP Reno Performance: A Simple Model and its Empirical Validation. *IEEE/ACM Trans. on Networking*, 8(2):133–145, 2000.

APPENDIX

A. Proof of Theorem 1

We first exhibit a few auxiliary results and then provide the proof the theorem. Let us begin with a characterization lemma:

Lemma 2: Let $\varphi(\cdot)$ be a positive-valued function, and $\psi(\cdot)$ be a primitive function of $\varphi(\cdot)$, that is $\psi'(\cdot) = \varphi(\cdot)$. The aggregate send rate X is characterized as

$$\mathbb{E}(\varphi(X(0))) = \frac{\psi(c) - \mathbb{E}^0(\psi(c - \bar{\eta}T_1))}{\bar{\eta}\mathbb{E}^0(T_1)}. \quad (12)$$

Proof. By Palm inversion formula,

$$\mathbb{E}(\varphi(X(0))) = \frac{\mathbb{E}^0\left(\int_0^{T_1} \varphi(X(s))ds\right)}{\mathbb{E}^0(T_1)}.$$

We have $X(t) = c - \bar{\eta}(T_{n+1} - t)$, $T_n \leq t < T_{n+1}$. Combining the last two displays, we have

$$\mathbb{E}(\varphi(X(0))) = \frac{\mathbb{E}^0\left(\int_0^{T_1} \varphi(c - \bar{\eta}(T_1 - s))ds\right)}{\mathbb{E}^0(T_1)}.$$

Now, use the change of variable $c - \bar{\eta}(T_1 - s) = v$ and rewrite the last display as

$$\mathbb{E}(\varphi(X(0))) = \frac{\mathbb{E}^0\left(\int_{c-\bar{\eta}T_1}^c \varphi(v)dv\right)}{\bar{\eta}\mathbb{E}^0(T_1)}.$$

The proof follows by the definition of $\psi(\cdot)$. \square

The next corollary follows immediately from the lemma by defining $\varphi(x) = x^m$ and invoking the binomial theorem.

Corollary 1: The moments of the aggregate rate are characterized as

$$\mathbb{E}(X(0)^m) = c^m - \frac{\sum_{k=2}^{m+1} \binom{m+1}{k} (-1)^k c^{m+1-k} \bar{\eta}^k \mathbb{E}^0(T_1^k)}{(m+1)\bar{\eta}\mathbb{E}^0(T_1)}.$$

As a remark, note that the $m+1$ th time-stationary moment of the aggregate send rate is entirely determined by the event-stationary moments of the inter link congestion times of the orders 1 to $m+1$. The following easy observation shall be useful in the sequel.

Lemma 3: Under a loss policy that verifies (ONE), the inter link congestion times satisfy, for any m ,

$$S_n^m = \frac{1}{\bar{\eta}^m} \sum_{i=1}^N (1 - \beta_i)^m z_i(n) x_i(n)^m, \quad \text{all } n. \quad (13)$$

Proof. The asserted equality follows directly by summing both sides of the equation (3) over i and invoking the definition of the link congestion times. \square

The remainder of this section is the proof of Theorem 1. Recall the short-cut $x_i(n) := X_i(T_n)$ and that $\mathbb{E}^0(T_1^m) = \mathbb{E}(S_n^m)$. From Corollary 1, we obtain

$$\mathbb{E}(X(0)) = c - \frac{\bar{\eta} \mathbb{E}^0(T_1^2)}{2 \mathbb{E}^0(T_1)}. \quad (14)$$

The remainder of the proof shows that

$$\frac{\mathbb{E}^0(T_1^2)}{\mathbb{E}^0(T_1)} = K$$

where K is a constant that depends on the parameters (η_1, β_1, N) . From (13), it follows that under the homogeneity assumption, we have

$$\frac{\mathbb{E}^0(T_1^2)}{\mathbb{E}^0(T_1)} = \frac{1 - \beta_1}{\eta_1 N} \frac{\sum_{i=1}^N \mathbb{E}^0(z_i(0) x_i(0)^2)}{\sum_{i=1}^N \mathbb{E}^0(z_i(0) x_i(0))}. \quad (15)$$

Now, rise to the power 2 both sides of the relation (3) using (13). After a few simplifications, we obtain

$$\begin{aligned} x_i(n+1)^2 &= x_i(n)^2 - \left(1 - \beta_i^2 + 2\frac{\eta_i}{\bar{\eta}}(1 - \beta_i)^2\right) (z_i(n)x_i(n))^2 \\ &\quad + \frac{2\eta_i}{\bar{\eta}}x_i(n) \sum_{j=1}^N (1 - \beta_j)z_j(n)x_j(n) \\ &\quad + \left(\frac{\eta_i}{\bar{\eta}}\right)^2 \sum_{j=1}^N (1 - \beta_j)^2 (z_j(n)x_j(n))^2. \end{aligned}$$

To arrive to the last display we use the property $z_i(n)z_j(n) = 0$, whenever $i \neq j$, implied by the ONE assumption. Next, we sum both sides of the last display over i and take expectation, to obtain

$$\begin{aligned} \sum_{i=1}^N \left(\frac{1 + \beta_i}{1 - \beta_i} + \frac{2\eta_i}{\bar{\eta}} - \sum_{j=1}^N \left(\frac{\eta_j}{\bar{\eta}} \right)^2 \right) (1 - \beta_i)^2 \mathbb{E}^0((z_i(0)x_i(0))^2) \\ = 2\mathbb{E}^0 \left(\sum_{i=1}^N \frac{\eta_i}{\bar{\eta}} x_i(0) \sum_{j=1}^N (1 - \beta_j) z_j(0) x_j(0) \right). \end{aligned}$$

For identical AIMD parameters, the last relation boils down to

$$\frac{\sum_{i=1}^N \mathbb{E}^0((z_i(0)x_i(0))^2)}{\sum_{i=1}^N \mathbb{E}^0(z_i(0)x_i(0))} = \frac{2c}{(1 + \beta_1)N + (1 - \beta_1)}.$$

It remains only to plug-in the last identity into (15) and then into (14) to retrieve the asserted identity of the theorem.

B. Proof of Theorem 2

Lemma 4: The time-stationary moments for two AIMD flows $i = 1$ and 2 that compete for a link of a capacity c are given by:

$$\begin{aligned} \frac{\mathbb{E}(X_1(0)^m)}{c^m} &= \frac{c}{(m+1)\eta_1 \mathbb{E}^0(T_1)} \\ &\times \left(a_{m+2} \frac{\mathbb{E}^0(z_1(0)x_1(0)^{m+1})}{c^{m+1}} - \frac{\mathbb{E}^0((1 - z_1(0))x_1(0)^{m+1})}{c^{m+1}} \right. \\ &\quad \left. + \sum_{k=0}^{m+1} a_k \frac{\mathbb{E}^0((1 - z_1(0))x_1(0)^k)}{c^k} \right) \end{aligned}$$

where for $k = 0, 1, \dots, m+1$,

$$\begin{aligned} a_k &= \binom{m+1}{m+1-k} \left(\frac{\eta_1}{\bar{\eta}}(1 - \beta_2) \right)^{m+1-k} \left(1 - \frac{\eta_1}{\bar{\eta}}(1 - \beta_2) \right)^k \\ a_{m+2} &= \left(\beta_1 + \frac{\eta_1}{\bar{\eta}}(1 - \beta_1) \right)^{m+1} - \beta_1^{m+1} \end{aligned}$$

and

$$\begin{aligned} \mathbb{E}^0(T_1) &= \frac{1}{\bar{\eta}} \left[(2 - \beta_1 - \beta_2) \mathbb{E}^0(z_1(0)x_1(0)) \right. \\ &\quad \left. + c(1 - \beta_2)(1 - \mathbb{E}^0(z_1(0)))(1 - \beta_2) \mathbb{E}^0(x_1(0)) \right]. \end{aligned}$$

Proof. The proof follows from Palm inversion formula, after some tedious computations. By the Palm inversion formula,

we have

$$\mathbb{E}(X_1(0)^m) = \frac{1}{\mathbb{E}^0(T_1)} \mathbb{E}^0 \left(\int_0^{T_1} X_1(s)^m ds \right). \quad (16)$$

We continue with computing the integral, given $(x_i(0), z_i(0))$, $i = 1, 2, \dots, N$,

$$\begin{aligned} \int_0^{T_1} X_1(s)^m ds &= \int_0^{T_1} [(1 - (1 - \beta_1)z_1(0))x_1(0) + \eta_1 s]^m ds \\ &= \frac{1}{\eta_1} \int_{(1 - (1 - \beta_1)z_1(0))x_1(0)}^{(1 - (1 - \beta_1)z_1(0))x_1(0) + \eta_1 S_0} u^m du \\ &= \frac{1}{(m+1)\eta_1} \left([(1 - (1 - \beta_1)z_1(0))x_1(0) \right. \\ &\quad \left. + \eta_1 S_0]^{m+1} - (1 - (1 - \beta_1^{m+1})z_1(0))x_1(0)^{m+1} \right) \\ &= \frac{1}{(m+1)\eta_1} \left(\sum_{j=0}^{m+1} \binom{m+1}{j} [1 - (1 - \beta_1^{m+1-j})z_1(0)] \right. \\ &\quad \left. \eta_1^j x_1(0)^{m+1-j} S_0^j - (1 - (1 - \beta_1^{m+1})z_1(0))x_1(0)^{m+1} \right). \end{aligned}$$

Note that the equality $((1 - (1 - \beta_1)z_1(0))x_1(0))^{m+1} = (1 - (1 - \beta_1^{m+1})z_1(0))x_1(0)^{m+1}$ comes from the fact that $z_1(0)$ is 0 or 1.

It follows

$$\begin{aligned} (m+1)\eta_1 \int_0^{T_1} X_1(s)^m ds \\ = \sum_{j=1}^{m+1} \binom{m+1}{j} [1 - (1 - \beta_1^{m+1-j})z_1(0)] \eta_1^j x_1(0)^{m+1-j} S_0^j. \end{aligned} \quad (17)$$

By (13) and $z_i(n)z_j(n) = 0$ whenever $i \neq j$, for all n , we can write

$$\begin{aligned} [1 - (1 - \beta_1^{m+1-j})z_1(0)] S_0^j \\ = \frac{1}{\bar{\eta}} \sum_{k=1}^N [1 - (1 - \beta_1^{m+1-j})z_1(0)] (1 - \beta_k)^j z_k(0) x_k(0)^j \\ = \frac{1}{\bar{\eta}} \left[\beta_1^{m+1-j} (1 - \beta_1)^j z_1(0) x_1(0)^j + \sum_{k \neq 1} (1 - \beta_k)^j z_k(0) x_k(0)^j \right]. \end{aligned}$$

Plug-in the last identity in (17) to obtain

$$\begin{aligned} (m+1)\eta_1 \int_0^{T_1} X_1(s)^m ds \\ = \sum_{j=1}^{m+1} \binom{m+1}{j} \left(\frac{\eta_1}{\bar{\eta}} \right)^j (1 - \beta_1)^j \beta_1^{m+1-j} z_1(0) x_1(0)^{m+1} + \\ + \sum_{k \neq 1} \sum_{j=1}^{m+1} \binom{m+1}{j} \left(\frac{\eta_1}{\bar{\eta}} \right)^j (1 - \beta_k)^j z_k(0) x_1(0)^{m+1-j} x_k(0)^j. \end{aligned}$$

The first summation element on the right-hand side of the equality simplifies immediately, so that we have

$$\begin{aligned} (m+1)\eta_1 \int_0^{T_1} X_1(s)^m ds \\ = \left[\left(\beta_1 + \frac{\eta_1}{\bar{\eta}}(1 - \beta_1) \right)^{m+1} - \beta_1^{m+1} \right] z_1(0) x_1(0)^{m+1} + \\ + \sum_{k \neq 1} \sum_{j=1}^{m+1} \binom{m+1}{j} \left(\frac{\eta_1}{\bar{\eta}} \right)^j (1 - \beta_k)^j z_k(0) x_1(0)^{m+1-j} x_k(0)^j. \end{aligned}$$

Denote with s_N the second summation element in the right-hand side of the last equality. We now specialize to $N = 2$ and proceed with computing s_2 . We thus have

$$(m+1)\eta_1 \int_0^{T_1} X_1(s)^m ds = \left[\left(\beta_1 + \frac{\eta_1}{\bar{\eta}}(1-\beta_1) \right)^{m+1} - \beta_1^{m+1} \right] z_1(0)_{X_1(0)}^{m+1} + s_2 \quad (18)$$

with s_2 given by

$$\sum_{j=1}^{m+1} \binom{m+1}{j} \left(\frac{\eta_1}{\bar{\eta}}(1-\beta_2) \right)^j (1-z_1(0))_{X_1(0)}^{m+1-j} (c-x_1(0))^j.$$

Note that

$$\begin{aligned} & (1-z_1(0))_{X_1(0)}^{m+1-j} (c-x_1(0))^j \\ &= \sum_{k=0}^j \binom{j}{k} (-1)^{j-k} c^k (1-z_1(0))_{X_1(0)}^{m+1-k} \\ &= (-1)^j (1-z_1(0))_{X_1(0)}^{m+1} \\ & \quad + \sum_{k=1}^j \binom{j}{k} (-1)^{j-k} c^k (1-z_1(0))_{X_1(0)}^{m+1-k}. \end{aligned}$$

Substituting the last identity into s_2 , we have

$$s_2 = \sum_{j=1}^{m+1} \binom{m+1}{j} \left(\frac{\eta_1}{\bar{\eta}}(1-\beta_2) \right)^j \times \left[(-1)^j (1-z_1(0))_{X_1(0)}^{m+1} + \sum_{k=1}^j \binom{j}{k} (-1)^{j-k} c^k (1-z_1(0))_{X_1(0)}^{m+1-k} \right].$$

Computing the sum for the first element in brackets, we have

$$s_2 = \left[\left(1 - \frac{\eta_1}{\bar{\eta}}(1-\beta_2) \right)^{m+1} - 1 \right] (1-z_1(0))_{X_1(0)}^{m+1} + s'_2 \quad (19)$$

with

$$\begin{aligned} s'_2 &= c^{m+1} \sum_{j=1}^{m+1} \binom{m+1}{j} \left(\frac{\eta_1}{\bar{\eta}}(1-\beta_2) \right)^j \\ & \quad \times \sum_{k=1}^j \binom{j}{k} (-1)^{j-k} (1-z_1(0)) \frac{x_1(0)^{m+1-k}}{c^{m+1-k}}. \end{aligned}$$

We can write $s'_2 = c^{m+1} \sum_{j=1}^{m+1} \sum_{k=m+1-j}^m d_{j,k} (1-z_1(0)) \frac{x_1(0)^k}{c^k}$, where

$$\begin{aligned} d_{j,k} &= \binom{m+1}{j} \binom{j}{m+1-k} \left(\frac{\eta_1}{\bar{\eta}}(1-\beta_2) \right)^j (-1)^{-j+m+1-k} = \\ & \binom{m+1}{m+1-k} \binom{k}{j-m-1+k} \left(\frac{\eta_1}{\bar{\eta}}(1-\beta_2) \right)^j (-1)^{-j+m+1-k} \end{aligned}$$

Then,

$$s'_2 = c^{m+1} \sum_{k=0}^m a_k (1-z_1(0)) \frac{x_1(0)^k}{c^k} \quad (20)$$

with

$$a_k = \sum_{j=m+1-k}^{m+1} d_{j,k}, \quad k = 0, 1, \dots, m.$$

The coefficients a_k admit a simple form: $a_k =$

$$\begin{aligned} & \binom{m+1}{m+1-k} \sum_{j=m+1-k}^{m+1} \binom{k}{j-m-1+k} \left(\frac{\eta_1}{\bar{\eta}}(1-\beta_2) \right)^j (-1)^{-j+m+1-k} = \\ & \binom{m+1}{m+1-k} \sum_{i=0}^k \binom{k}{i} \left(\frac{\eta_1}{\bar{\eta}}(1-\beta_2) \right)^{i+(m+1-k)} (-1)^i = \\ & \binom{m+1}{m+1-k} \left(\frac{\eta_1}{\bar{\eta}}(1-\beta_2) \right)^{m+1-k} \left(1 - \frac{\eta_1}{\bar{\eta}}(1-\beta_2) \right)^k. \end{aligned}$$

This recovers the asserted definition of the coefficients a_k . It remains to see that the asserted identity for the moments is true. To that end, plug-in (20) into (19) and the result into (18). Define a_{m+2} as asserted. Take expectation on the identity obtained for $\int_0^{T_1} X_1(s)^m ds$. $\mathbb{E}^0(T_1)$ follows readily by taking expectation in (13). The result follows by (16). \square

Proof of Theorem 2

Take $m = 2$ in the identity for $\mathbb{E}(X_1(0)^m)$ given by Lemma 4. The setting is for symmetric TCP flows, we thus can take $\eta_1 = \eta_2 = 1$ and $\beta_1 = \beta_2 = 1/2$. By direct computation

$$(a_1, a_2, a_3, a_4) = \left(\frac{1}{64}, \frac{9}{64}, \frac{37}{64}, \frac{56}{64} \right).$$

We now proceed in four steps, each step showing the respective asserted results for the beatdown, rate-independent, rate-proportional, and largest-rate policy. We use the simplifying notation $x_k = \mathbb{E}^0(x_1(0)^k)/c^k$. Our two flows are with identical AIMD parameters, hence $x_i = 1/2$ and $\mathbb{E}^0(z_i(0)) = 1/2$, $i = 1, 2$.

Step 1 (Beatdown) A stationary sample-path is a uniform-phase shift of the periodic solution in Figure 4. Considering 1 as the tagged flow, under $\tau_0 = 0$, $x_1(0) = 1$, $x_2(0) = 0$, and plugging these values in (3) and (4), we have $x_1(n) = \frac{3}{4}x_1(n-1) = \left(\frac{3}{4}\right)^n$ and $S_n = \frac{1}{2\bar{\eta}} \left(\frac{3}{4}\right)^n$, $n = 1, 2, \dots$ (see also Step 1. of the proof of Theorem 3 for further details). It follows that the surface under the curve of the square throughput is

$$\begin{aligned} & \int_0^{\tau_1} X_1(s)^2 ds \\ &= \sum_{n=0}^{\infty} \frac{1}{3} \left(\frac{x_1(n+1)^3 - x_1(n)^3}{x_1(n+1) - x_1(n)} S_n \right) \\ &= \frac{1}{6\bar{\eta}} \sum_{n=0}^{\infty} \left(\left(\frac{3}{4} \right)^{2n+2} + \left(\frac{3}{4} \right)^{2n+1} \left(\frac{3}{4} \right)^{2n} \right) \left(\frac{3}{4} \right)^n \\ &= \frac{2}{3\bar{\eta}}. \end{aligned}$$

Note that, $\tau_1 = \sum_{n=0}^{\infty} S_n = \frac{2c}{\bar{\eta}}$, which can also be directly obtained as the time for the other flow to increase its rate from 0 to c . If flow 1 is not the one to experience the loss during the period, we similarly compute $\int_0^{\tau_1} X(s)^2 ds = \frac{c^2}{3} \tau = \frac{2c^3}{3\bar{\eta}}$. We have that either of the last two cases occur with Palm probability 1/2, hence by Palm inversion formula, we get $\frac{\mathbb{E}(x_1(0)^2)}{c^2} = \frac{1}{3}$, as asserted.

Step 2 (rate independent). A flow is selected independently of the rate process, hence $\mathbb{E}^0(z_1(0)_{X_1(0)}^m) =$

$\mathbb{E}^0(z_1(0))\mathbb{E}^0(x_1(0)^m)$, for any real m . Under the prevailing assumptions, the equation in Lemma 4 reads as

$$\frac{\mathbb{E}(X_1(0)^2)}{c^2} = \frac{2}{3}(a_4x_3 - 2a_3x_3 + a_2x_2 + a_1x_1 + a_0)$$

with

$$(x_1, x_2, x_3) = \left(\frac{1}{2}, \frac{2}{7}, \frac{5}{28}\right).$$

Indeed, $\frac{\mathbb{E}(X_1(0))}{c} = \frac{1}{2}$. Also, from the square of the equation (3), using (13) and the symmetry, we obtain

$$\begin{aligned} \mathbb{E}^0(x_1(0)^2) &= \mathbb{E}^0\left(\left((1 - \frac{1}{2}z_1(0))x_1 + \frac{z_1(0)x_1(0) + (1 - z_1(0))(c - x_1(0))}{4}\right)^2\right) \\ &= \frac{1}{16}\left(9\mathbb{E}^0(x_1(0)^2) + \frac{1}{2}c^2 + 3c\mathbb{E}^0(x_1(0))\right), \end{aligned}$$

yielding the asserted value for x_2 . Also, using the symmetry,

$$\begin{aligned} \mathbb{E}^0(x_1(0)^3) &= \mathbb{E}^0((c - x_1(0))^3) \\ &= c^3 - 3c^2\mathbb{E}^0(x_1(0)) + 3c\mathbb{E}^0(x_1(0)^2) - \mathbb{E}^0(x_1(0)^3), \end{aligned}$$

that results in the asserted value for x_3 , given x_1 and x_2 .

Step 3 (rate-proportional) It is readily checked that for an arbitrary positive integer m , for the prevailing rate-proportional loss policy, the expression in Lemma 4 reads as ($a_{-1} := 0$)

$$\frac{\mathbb{E}(X(0)^m)}{c^m} = \frac{a_{m+2}x_{m+2} - (a_{m+1} + a_m)x_{m+1} + \sum_{k=0}^m (a_k - a_{k-1})x_k}{(m+1)\eta_1\mathbb{E}(T_1)}$$

since $\mathbb{E}^0(z_1(0)x_1(0)^m) = \frac{\mathbb{E}^0(x_1(0)^{m+1})}{c}$. From [7], we have

$$(x_1, x_2, x_3, x_4) = \left(\frac{1}{2}, \frac{7}{26}, \frac{2}{13}, \frac{679}{7358}\right).$$

Plugging these values yields the asserted result.

Step 4 (largest-rate). From Proposition 1, a stationary sample-path is a uniform-phase shift of a 2-periodic function with values $\frac{2}{7}c$ and $\frac{3}{7}c$. It follows $\mathbb{E}^0(T_1) = c/7$ and $\int_0^{T_2} X(s)^2 ds = \int_0^{2c/7} (x + 2c/7)^2 dx = \frac{8}{147}c^3$. Dividing this value by $2\mathbb{E}^0(T_1)$, we obtain $\mathbb{E}(X^{(1)}(0)^2) = \frac{4}{21}c^2$.

C. Proof of Lemma 1

Consider (7) with $a_i = a_1$ and $b_i = b_1$. We have,

$$\begin{aligned} \tilde{x}_1(n+1) &= a_1\tilde{x}_1(n) \vee (\vee_{j \neq M_n} \{x_j(n)\} + b_1\tilde{x}_1(n)) \\ &= a_1\tilde{x}_1(n) \vee (\tilde{x}_2(n) + b_1\tilde{x}_1(n)). \end{aligned}$$

This shows $g_1(\cdot)$. By the same argument, but performing the minimum over the right-hand sides in (7), $g_N(\cdot)$ follows. Now, the expression for $i = 2, 3, \dots, N$, is merely inserting $a_1\tilde{x}_1(n)$ into the ordered sequence $\tilde{x}_j(n) + b_1\tilde{x}_1(n)$, $j = 2, 3, \dots, K$, so that the resulting sequence is sorted in a non-increasing order.

D. Proof of Proposition 1

The map $g : \mathcal{S}_{N-1} \rightarrow \mathcal{S}_{N-1}$ is continuous and \mathcal{S}_{N-1} is a subset of $\{x : \mathbb{R}^N : \sum_{i=1}^N x_i^2 \leq 1\}$, so that existence follows from Brouwer fixed point theorem. We next show that x^0 given by (9) is a fixed point of (8). Let $x^0 \in \mathcal{S}_{N-1}$ be such that $x^0 = g(x^0)$. Let $a_1x_1^0 < x_{i+1} + b_1x_1^0$ for all $i = 1, 2, \dots, N$. It follows that

$$\begin{aligned} x_i^0 &= x_{i+1}^0 + b_1x_1^0, \quad i = 1, 2, \dots, N-1 \\ x_N^0 &= a_1x_1^0. \end{aligned}$$

The last identities are equivalent to saying

$$\begin{aligned} x_i^0 &= (1 - (i-1)b_1)x_1^0, \quad i = 1, 2, \dots, N-1, \\ x_N^0 &= a_1x_1^0. \end{aligned}$$

The result follows from $\sum_{i=1}^N x_i^0 = 1$, and the definitions of a_1 and b_1 .

We next show that x^0 as defined above is a unique fixed point in the interior of \mathcal{S}_{N-1} . Recall that x^0 is such that $a_1x_1^0 < x_j^0 + b_1x_1^0$, for all $j = 1, 2, \dots, N$. To contradict, suppose y^0 is a fixed point of (8) in the interior of \mathcal{S}_{N-1} such that for some $i \in \{1, 2, \dots, N-1\}$, $a_1y_1^0 \geq y_i^0 + b_1y_1^0$. From the definition of $g(\cdot)$, it follows that $y_i^0 \geq y_i^0 + b_1y_1^0$, thus, $y_1^0 \leq 0$, a contradiction. This completes the proof.

Remark 3: The Jacobian $J(x^0)$ of the map $g(\cdot)$ at the point x^0 is equal to $(J(x^0))_{i,1} = b_1$, $i = 1, 2, \dots, N-1$, $(J(x^0))_{N,1} = a_1$, $(J(x^0))_{i,i+1} = 1$, $i = 1, 2, \dots, N-1$, with other elements equal to 0. Note that the Jacobian $J(x^0)$ is a companion matrix, so that its eigenvalues are the solutions of the characteristic polynomial $\lambda^{N-1} - b_1\lambda^{N-2} - \dots - b_1\lambda - a_1 = 0$, which can be rewritten as

$$\lambda^{N-1} = \eta_1(1 - \beta_1)(\lambda^{N-2} + \lambda^{N-3} + \dots + \lambda + 1) + \beta_1.$$

It is readily checked that $\lambda = 1$ is an eigenvalue.

E. Proof of Theorem 3

The proof proceeds in two steps:

Step 1 We first show that

$$\frac{\mathbb{E}(X(0))}{c} = 1 - \frac{1}{1 + \frac{1+\beta_1}{1-\beta_1}N} \frac{N \mathbb{E}^0(\sum_{j=1}^N z_j(0)x_j(0)^2)}{2 \mathbb{E}^0(\sum_{j=1}^N z_j(0)x_j(0))c}.$$

Suppose that at a selected congestion time τ_n , a flow i is tagged. Let $x_0 := x_i(n)$. The spare capacity just after τ_n is $(1 - \beta_i)x_0$. The flow i undergoes a multiplicative decrease at each congestion time in $[\tau_n, \tau_{n+1})$. With no loss of generality, let $\tau_n = 0$. Let s_n be the time between the congestion times $n-1$ and n as seen from 0, and x_n be the rate of the flow i at the n th congestion time. We have

$$x_{n+1} = \beta_i x_n + \eta_i s_n \quad (21)$$

The dynamics entails that

$$\beta_i x_n + \eta_i s_n = x_n - \eta_{-i} s_n$$

where $\eta_{-i} := \sum_{j=1}^N \eta_j 1_{\{j \neq i\}}$.

Hence,

$$s_n = \frac{1 - \beta_i}{\bar{\eta}} x_n.$$

Plugging the last identity into (21), we have $x_{n+1} = \rho_i x_n$ where we denote $\rho_i = \beta_i + \frac{\bar{\eta}_i}{\bar{\eta}}(1 - \beta_i)$. It follows that

$$x_n = x_0 \rho_i^n \quad \text{and} \quad s_n = \frac{1 - \beta_i}{\bar{\eta}} x_0 \rho_i^n.$$

The amount of units sent by the flow i in $[0, \tau_1)$ is equal to

$$\begin{aligned} \int_0^{\tau_1} X_i(u) du &= \sum_{n=0}^{\infty} \frac{1}{2} (x_{n+1} + \beta_i x_n) s_n \\ &= \frac{1}{2\bar{\eta}} \frac{(1 + \beta_i) \frac{\bar{\eta}}{\bar{\eta}_i} - (1 - \beta_i)}{1 + \rho_i} x_0^2. \end{aligned}$$

Subsequently, the amount of units sent on $[0, \tau_1)$ by flows other than i is

$$\begin{aligned} \int_0^{\tau_1} \sum_{j \neq i} X_j(u) du &= (c - x_0) \tau_1 + \frac{1}{2} x_0 \tau_1 \\ &= \frac{1}{\bar{\eta}_i} \left(c x_0 - \frac{1}{2} x_0^2 \right). \end{aligned}$$

Putting the last two displays into the Palm inversion formula and after some straightforward manipulations, we retrieve the asserted identity.

Step 2 By applying the same trick as in the proof of Theorem 1 to the identities (10)–(11), we obtain

$$\frac{\mathbb{E}^0(\sum_{j=1}^N z_j(0) \times_j(0)^2)}{\mathbb{E}^0(\sum_{j=1}^N z_j(0) \times_j(0))c} = \frac{2}{N}.$$

The proof follows.