Loss Strategies for Competing TCP/IP Connections*

Eitan Altman[1], Rachid El-Azouzi[2], David Ros[3], and Bruno Tuffin[4]

[1] INRIA, 2004 route des Lucioles, 06902 Sophia Antipolis, France
[2] Université d'Avignon et des Pays de Vaucluse (IUP), LIA-CERI, 339 chemin des Meinajariès, BP1228, 84911 Avignon Cedex 9, France
[3] GET/ENST Bretagne, Rue de la Châtaigneraie, CS 17607, 35576 Cesson Sévigné Cedex, France
[4] IRISA/INRIA, Campus Universitaire de Beaulieu, 35042 Rennes Cedex, France
Eitan.Altman@sophia.inria.fr, Rachid.Elazouzi@lia.univ-avignon.fr, David.Ros@enst-bretagne.fr, Bruno.Tuffin@irisa.fr

Abstract. We study in this paper two competing TCP connections that share a common bottleneck link. When congestion occurs, one (or both) connections will suffer a loss that will cause its throughput to decrease by a multiplicative factor. The identity of the connection that will suffer a loss is determined by a randomized "loss strategy" that may depend on the throughputs of the connections at the congestion instant. We analyze several such loss strategies. After deriving some results for the general asymmetric case, we focus in particular on the symmetric case and study the influence of the strategy on the average throughput and average utilization of the link. As the intuition says, a strategy that assigns a loss to a connection with a higher throughput is expected to give worse performance since the total instantaneous throughput after a loss is expected to be lower with such a strategy. We show that, surprisingly, the average throughput and average link utilizations are *invariant*: they are the same under any possible strategy; the link utilization is 6/7 of the link capacity. We show, in contrast, that the second moment of the throughput does depend on the strategy.

1 Introduction

The mathematical analysis of the performance of TCP has been a major research area in networking. Different types of approaches have been suggested and validated. On the one hand, there have been models focusing on a single connection that is subject to some exogenous loss process (which does not depend on that connection), see e.g. [1]. This approach is appealing when there is a large amount of traffic, so that we can neglect the effect of the single connection on events that cause losses. An alternative approach is necessary when the window increase of a connection is itself a central cause for losses. This occurs typically when a small number of connections compete over bandwidth, say, at a bottleneck link. A

* This work was partially supported by INRIA's "TCP" Cooperative Research Action.

main mathematical approach for studying this situation has been to study several connections sharing a bottleneck, and then make the simplifying assumption that all connections reduce their windows simultaneously upon congestion [2–4]. With this approach, it has been shown [4] that the throughput achieved by a TCP connection is inversely proportional to RTT^{α} with $1 < \alpha < 2$, where RTT is the two-way propagation delay of the connection. However, it turns out that in practice this assumption does not hold, except for drop tail buffers and connections with similar Round Trip Times (RTTs) [5]. Indeed, traces in [2] (e.g. Fig. 5) show that the synchronization assumption is invalid for asymmetric connections for a drop tail buffer.

Instead of considering synchronization, two modeling approaches have been developed for determining which connection will suffer a packet loss. In the model of Baccelli and Hong [6], the probability that a connection will lose a packet is a constant: it does not depend on its current throughput. As argued in [7], such an assumption is valid in describing AIMD protocols in which packet transmission rates are constant, and the throughput is varied by changing the packet size. An alternative model has been considered in [8] in which the probability that a connection loses a packet is proportional to the throughput at the congestion instant. This is called the "proportional strategy". As validated by simulations [9], this model is appropriate for standard TCP where packet size is constant.

Motivated by these two approaches, we raise the question of what is the throughput of an AIMD protocol as a function of the strategy that determines which connection loses a packet at a congestion instant. We focus on the simple scenario of two competing connections.

Our findings are as follows. We first study the constant probability model in [6]. In that paper, a linear set of stochastic recursive equations has been introduced for obtaining the throughput, in which the state variables correspond to the connections' throughputs after a loss. In this paper we present an alternative set of stochastic recursive equations in which the states correspond to the throughput *just before* the loss occurs. We show that our approach allows us to reduce the dimensionality of the system by one, so in particular, the case of two connections can be described by a one-dimensional state equation. This allows us to obtain an *explicit* expression for the throughput in the general asymmetric case for the constant probability model. As a corollary of this result, it is seen in the symmetric case that the link utilization is 6/7 of its capacity.

We then study a new strategy in which the connection with the larger instantaneous throughput is the one to lose a packet at congestion instants. Surprisingly, we obtain the same average throughput and link utilization in the symmetric case as for the constant probability model. Moreover, this is the same utilization also obtained for the proportional strategy. This motivated us to examine the behavior of an arbitrary strategy. Our main finding is that although the expectation of throughputs at loss instants depend on the strategy, the average throughput is an invariant quantity for the case of symmetric connections.

We finally derive a general expression for the second moment of the throughput and compare the performance of the three strategies mentioned above in the symmetric case, in order to find out which one has the smallest throughput variability.

The structure of the paper is as follows. In Sect. 2 we study the throughput of the constant loss strategy, whereas the Larger Throughput Loss (LTL) strategy is analyzed in Sect. 3. Section 4 then presents some numerical experimentations and comparisons between the strategies. Section 5 studies the average throughputs in the symmetric setting under an arbitrary strategy and obtains the invariance property. Section 6 then provides an expression for the second moment of the throughput under an arbitrary strategy and a comparison for the three aforementioned strategies. We end with a concluding section.

2 Fixed Loss Probabilities: Model and Analysis

2.1 Basic Definitions and Assumptions

This model is based on [6] where an additive increase, multiplicative decrease (AIMD) model is used to describe the joint throughput evolution of a set of TCP sessions sharing a common router bottleneck.

In full generality, let N be the number of TCP sessions competing for bandwidth, and C the capacity of the bottleneck router. Let T_n be the n-th congestion epoch and $\tau_{n+1} = T_{n+1} - T_n$. Let also η_i be the additive increase rate for session i and $\beta^{(i)}$ be its multiplicative decrease rate. Usually, $\beta^{(i)} = 1/2\ \forall i$ and η_i is taken as the square inverse of the round trip time of session i. We consider here $Y_n^{(i)}$, the throughput of session i *before* the n-th congestion epoch, instead of $X_n^{(i)}$, the throughput after the n-th congestion epoch like in [6].

Denote by $\bar{Y}^{(i)}$ session i's mean throughput. As in [6], let $a_n^{(i)}$ be a Bernoulli random variable with value 1 if session i experiences a loss at the n-th congestion epoch, and 0 otherwise, so that $\mathbb{E}[a_n^{(i)}] = p^{(i)}$. Note that the $a_n^{(i)}$ $(1 \leq i \leq N)$ are correlated to make sure that at least one packet is lost at each congestion time. We have

$$Y_{n+1}^{(i)} = \gamma_n^{(i)} Y_n^{(i)} + \tau_{n+1}\eta_i \tag{1}$$

where $\gamma_n^{(i)} = (1 - a_n^{(i)}) + \beta^{(i)} a_n^{(i)}$. As in [6], we assume here that there is a loss as soon as the router capacity is reached, i.e., as soon as

$$\sum_{i=1}^{N} \gamma_n^{(i)} Y_n^{(i)} + \tau_{n+1} \sum_{i=1}^{N} \eta_i = C \ . \tag{2}$$

This assumption will allow us to derive the throughput at the different congestion epochs.

2.2 Computation of the Average Throughput

The goal of this subsection is to derive the average throughput of a session in terms of the loss probabilities when the number of sessions is $N = 2$.

First, using (2), we get the time between the n-th and $(n+1)$-th congestion epochs

$$\tau_{n+1} = \frac{C - \sum_{i=1}^{N} \gamma_n^{(i)} Y_n^{(i)}}{\sum_{i=1}^{N} \eta_i} \quad . \tag{3}$$

Using this relation we are able to derive a closed-form of the average throughput $\bar{Y}^{(1)}$ of session 1. The average throughput $\bar{Y}^{(2)}$ of session 2 can be obtained in the same way (or by switching the indexes 1 and 2 in the following formula).

Proposition 1. *Assume that* $N = 2$*. If we denote* $p^{(12)} = \mathbb{E}(a^{(1)}a^{(2)})$*, we obtain*

$$\begin{aligned}
\bar{Y}^{(1)} = \frac{C}{2}\Bigg(\Bigg(&\frac{{\eta_1}^2\big(1-\beta^{(2)}\big)^2 p^{(2)}}{\big(\eta_1+\eta_2\big)^2} + 2{\eta_1}^2 p^{(2)}\big(1-\beta^{(2)}\big)\times \\
&\times\left(\big(1-\beta^{(2)}\big)p^{(2)} - \frac{\eta_1\big(1-\beta^{(2)}\big)^2 p^{(2)}}{\eta_1+\eta_2} - \frac{\eta_2\big(1-\beta^{(1)}\big)\big(1-\beta^{(2)}\big)p^{(12)}}{\eta_1+\eta_2}\right)\times \\
&\times\left(\Big(\big(1-\beta^{(2)}\big)\eta_1 p^{(2)} + \big(1-\beta^{(1)}\big)\eta_2 p^{(1)}\Big)^{-1}\big(\eta_1+\eta_2\big)^{-1}\right)\times \\
&\times\left(\frac{\eta_1\big(1-\beta^{(2)}\big)^2 p^{(2)}}{\eta_1+\eta_2} - \left(1+\frac{\eta_2}{\eta_1+\eta_2}\right)\big(1-\beta^{(1)}\big)^2 p^{(1)} + \right. \\
&\left. + 2\frac{\eta_2\big(1-\beta^{(1)}\big)\big(1-\beta^{(2)}\big)p^{(12)}}{\eta_1+\eta_2} - 2\big(1-\beta^{(2)}\big)p^{(2)} + 2\big(1-\beta^{(1)}\big)p^{(1)}\right)\times \\
&\times\left(2\frac{\big(1-\beta^{(1)}\big)\eta_2 p^{(1)}}{\eta_1+\eta_2} + 2\frac{\big(1-\beta^{(2)}\big)\eta_1 p^{(2)}}{\eta_1+\eta_2} - \frac{{\eta_1}^2\big(1-\beta^{(1)}\big)^2 p^{(1)}}{\big(\eta_1+\eta_2\big)^2}\right. \\
&\left. - \frac{{\eta_2}^2\big(1-\beta^{(2)}\big)^2 p^{(2)}}{\big(\eta_1+\eta_2\big)^2} - 2\frac{\eta_1\big(1-\beta^{(1)}\big)p^{(1)}\eta_2\big(1-\beta^{(2)}\big)p^{(12)}}{\big(\eta_1+\eta_2\big)^2}\right)^{-1} + \\
&+ 2\eta_1 p^{(2)}\left(\big(1-\beta^{(2)}\big)p^{(2)} - \frac{\eta_1\big(1-\beta^{(2)}\big)^2 p^{(2)}}{\eta_1+\eta_2} - \frac{\eta_2\big(1-\beta^{(1)}\big)\big(1-\beta^{(2)}\big)p^{(12)}}{\eta_1+\eta_2}\right)\times \\
&\times \frac{1-\beta^{(2)}}{\big(1-\beta^{(2)}\big)\eta_1 p^{(2)} + \big(1-\beta^{(1)}\big)\eta_2 p^{(1)}} + \frac{\eta_1\big(1-\beta^{(2)}\big)^2 p^{(2)}}{\eta_1+\eta_2}\Bigg)\times \\
&\times\left(\frac{\eta_1 p^{(2)} + \eta_2 p^{(1)} - \beta^{(2)}\eta_1 p^{(2)} - \beta^{(1)}\eta_2 p^{(1)}}{\big(\eta_1+\eta_2\big)\big(1-\beta^{(1)}\big)\big(1-\beta^{(2)}\big)p^{(1)}p^{(2)}}\right)
\end{aligned}$$

The proof of this proposition is provided in [10].

Corollary 1. *Still assuming* $N = 2$*, the symmetric case yields*

$$\begin{aligned}
\bar{Y}^{(1)} = \frac{C}{4} \times &\frac{2\,p^{(1)} - p^{(12)} + p^{(12)}\,\beta^{(1)}}{{p^{(1)}}^2\Big(3 + \beta^{(1)} - p^{(12)} + p^{(12)}\,\beta^{(1)}\Big)} \\
&\times\Big(p^{(12)}\,\beta^{(1)}\,p^{(1)} + 2\,p^{(1)} + 2\,p^{(1)}\,\beta^{(1)} - p^{(1)}\,p^{(12)} + p^{(12)} - p^{(12)}\,\beta^{(1)}\Big) \quad .
\end{aligned} \tag{4}$$

Proof. Just replace η_2 by η_1, $\beta^{(2)}$ by $\beta^{(1)}$ and $p^{(2)}$ by $p^{(1)}$ in Proposition 1. □

2.3 Sampling the Loss Probabilities

The previous expressions of the average throughput are general in the sense that no special sampling structure has been used for the losses. In this section, we aim at studying how the losses can be sampled and how it impacts on the average throughput formula.

Independent Sampling. As in [6], we can assume that the $a_n^{(i)}$ are at first generated independently, such that $\mathbb{P}[a_n^{(i)} = 1] = \pi^{(i)}$, with $\pi^{(i)}$ given, but that the samples are restricted to the domain where at least one loss is experienced. This requires a derivation of $\pi^{(i)}$ in terms of the $p^{(j)}$.

Assuming $N = 2$, we have as in [6]

$$\begin{cases} p^{(1)} = \frac{\pi^{(1)}}{1-(1-\pi^{(1)})(1-\pi^{(2)})} \\ p^{(2)} = \frac{\pi^{(2)}}{1-(1-\pi^{(1)})(1-\pi^{(2)})} \end{cases}$$

where $\pi^{(i)}$ is for the loss probability for user i, sampled independently, but reduced to the domain such that a loss is actually experienced. This gives

$$\begin{cases} \pi^{(1)} = p^{(1)}(\pi^{(1)} + \pi^{(2)} - \pi^{(1)}\pi^{(2)}) \\ \pi^{(2)} = p^{(2)}(\pi^{(1)} + \pi^{(2)} - \pi^{(1)}\pi^{(2)}) \end{cases}$$

We obtain the relation

$$\pi^{(1)} = \frac{p^{(1)}}{p^{(2)}}\pi^{(2)}$$

which gives (assuming $\pi^{(2)} > 0$)

$$\pi^{(2)} = \frac{p^{(1)} + p^{(2)} - 1}{p^{(1)}} \quad \text{and then} \quad \pi^{(1)} = \frac{p^{(1)} + p^{(2)} - 1}{p^{(2)}} \ .$$

Then an assumption $p^{(1)} + p^{(2)} = 1$ cannot be used. Also, it seems difficult to make sure that $\pi^{(1)} \leq 1$ and $\pi^{(2)} \leq 1$ for every pair $(p^{(1)}, p^{(2)})$. Thus, this sampling procedure does not work in full generality.

A Single Loss at Congestion Epochs. The simplest way to sample is by using the relation

$$a_n^{(2)} = 1 - a_n^{(1)}$$

with $a_n^{(1)}$ Bernoulli random variable such that $\mathbb{P}[a_n^{(1)} = 1] = p^{(1)}$. This means that at each congestion epoch, one and only one session will see a decrease of its throughput. We then have

$$p^{(2)} = 1 - p^{(1)} \quad \text{and} \quad p^{(12)} = 0 \ .$$

Substituting these values of $p^{(2)}$ and $p^{(12)}$ in the equation given in Proposition 1, we can obtain a closed, explicit expression for $\bar{Y}^{(1)}$ (not shown for space reasons). The symmetric case (that is, taking $p^{(1)} = p^{(2)} = 1/2$ and $\beta^{(2)} = \beta^{(1)}$) yields

$$\bar{Y}^{(1)} = \frac{\left(1 + \beta^{(1)}\right) C}{3 + \beta^{(1)}} \ .$$

If $\beta^{(1)} = 1/2$, we obtain $\bar{Y}^{(1)} = \frac{3}{7}C$, like in [8] for the proportional loss strategy.

3 The Largest Throughput Loss (LTL) Strategy

Let us look at the case where the session that is penalized is systematically the one with the largest throughput. We call this the "Largest Throughput Loss" (LTL) strategy. Consider the n-th congestion epoch, with throughputs $Y_n^{(1)}$ and $Y_n^{(2)}$ such that $Y_n^{(1)} + Y_n^{(2)} = C$. Without loss of generality, assume $Y_n^{(1)} > Y_n^{(2)}$ and that the additive increase is 1.

3.1 The Symmetric Case: the Periodic Solution

We identify a periodic solution for the evolution of the system. In this regime, we assume (without loss of generality) that at time n, connection 1 has a larger throughput than connection 2. We seek for a regime in which at time $n+1$ the situation is reversed, and so on. This gives the following dynamics:

$$\begin{cases} Y_n^{(1)}/2 + \tau_{n+1} = Y_n^{(2)} \\ Y_n^{(2)} + \tau_{n+1} = Y_n^{(1)} \\ Y_n^{(1)} + Y_n^{(2)} = C \ , \end{cases}$$

leading to

$$\tau_{n+1} = \frac{1}{7}C, \quad Y_n^{(1)} = \frac{4}{7}C \quad \text{and} \quad Y_n^{(2)} = \frac{3}{7}C \ .$$

As in the proof of Proposition 1, but due to the periodicity of the system, the average throughput is given by $S/\mathbb{E}[2\tau]$ where S is the cumulative throughput of a session between congestion epochs n and $n+2$ (in one period, the throughput is going from $2C/7$ to $3C/7$ and in the other one from $3C/7$ to $4C/7$). This gives $S = \frac{12}{98}C^2$, leading again to $\bar{Y}^{(1)} = \bar{Y}^{(2)} = \frac{3}{7}C$ and an average utilization of $\frac{6}{7}$ as we obtained in the previous section and as is the case in the model in [8]. Obviously, $\mathbb{E}[Y_n^{(i)}]$ are also the same in all three cases (and equal to $C/2$). One could wonder whether in fact the distribution of the rates is independent of the way one chooses the connection to decrease the rate at T_n. Note however, that $\mathbb{E}[(Y_n^{(i)})^2] = 25C^2/98$ in our example, which is different than the value of $7C^2/26$ obtained in the regime considered in [8].

3.2 The Dynamic Equations for the Asymmetric Case

For each connection $i = 1, 2$ we have

$$Y_{n+1}^{(i)} = \begin{cases} Y_n^{(i)}/2 + \tau_{n+1}\eta_i & \text{if } Y_n^{(i)} > C/2, \\ Y_n^{(i)} + \tau_{n+1}\eta_i & \text{if } Y_n^{(i)} < C/2. \end{cases} \tag{5}$$

For the case that $Y_n^{(i)} = C/2$ any tie-breaking rule can be considered. Combining this with the relation $Y_n^{(2)} = C - Y_n^{(1)}$ as well as $Y_{n+1}^{(2)} = C - Y_{n+1}^{(1)}$ gives

$$\tau_{n+1} = \begin{cases} \frac{Y_n^{(i)}}{2(\eta_1+\eta_2)} & \text{if } Y_n^{(i)} > C/2, \\ \frac{C-Y_n^{(i)}}{2(\eta_1+\eta_2)} & \text{if } Y_n^{(i)} < C/2. \end{cases}$$

Substituting in (5) gives

$$Y_{n+1}^{(i)} = \begin{cases} \frac{1}{2}\left(1 + \frac{\eta_i}{\eta_1+\eta_2}\right) Y_n^{(i)} & \text{if } Y_n^{(i)} > C/2, \\ \left(1 - \frac{\eta_i}{2(\eta_1+\eta_2)}\right) Y_n^{(i)} + \frac{C\eta_i}{2(\eta_1+\eta_2)} & \text{if } Y_n^{(i)} < C/2. \end{cases}$$

These equations can be used to obtain the exact transient behavior of the system. The average throughput can then be computed by

$$\bar{Y}^{(i)} = \lim_{n\to\infty} \frac{\sum_{k=1}^{n} \tau_{k+1}(Y_{k+1}^{(i)} + \gamma_k^{(i)} Y_k^{(i)})/2}{\sum_{k=1}^{n} \tau_{k+1}} \quad .$$

3.3 The Case $\eta_2/\eta_1 \to 0$

We consider here the case of $x \to 0$ where $x := \eta_2/\eta_1$ and assume for simplicity that $\beta^{(i)} = 1/2$. We present a heuristic argument to compute the bandwidth sharing.

Connection 2 will increase its rate until it reaches $C/2$, so its trajectory at steady state will be periodic (with a period of duration of $C/(4\eta_2)$), linearly increasing between $C/4$ to $C/2$. Its average throughput is 3C/8.

Connection 1 Fix $\Delta = \sqrt{x}/\eta_2$. We can view the problem as one with two time scales: connection 1 is much faster than connection 2, so during the interval $[n\Delta, (n+1)\Delta)$, the throughput of connection 2 can be approximated by a constant which we denote by $Y^{(2)}(n)$; assume that this constant is smaller than $C/2$. During that interval, the throughput of connection 1 will oscillate very quickly (between half of the remaining and all the remaining bandwidth) so that it will use in average over that interval 3/4 of the remaining bandwidth. Thus its average bandwidth during the interval is $(3/4)(C - Y^{(2)}(n))$, and during the whole period of $C/(4\eta_2)$ it will be $(3/4)(C - 3C/8) = 15C/32$.

Thus as $x \to 0$ we see that the fast connection will get 5/4 of the throughput of the slow connection under the LTL strategy.

4 Numerical Results for the Fairness in Bandwidth Sharing

We study in this section the fairness in throughput as a function of the round trip times. We recall that the square root formula of TCP as well as its refinements (see [1, 11]) predict that the throughput of a connection should be inversely proportional to its RTT. We will compare this with the fairness obtained under our model of interacting connections.

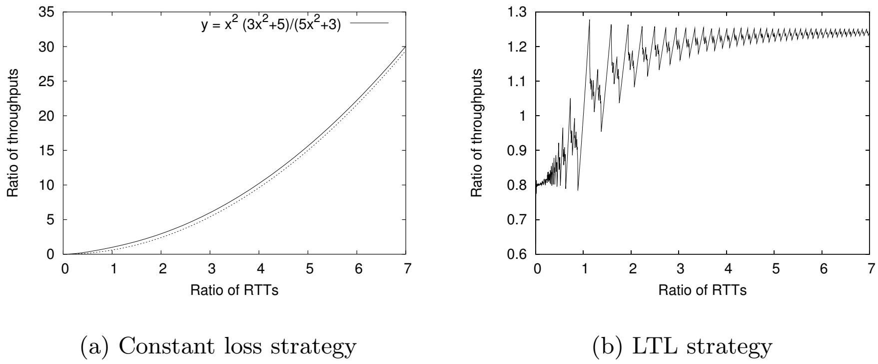

Fig. 1. The ratio $\bar{Y}_1/\bar{Y}_2$ as a function of the ratio $R^{(2)}/R^{(1)}$

4.1 Constant Loss Strategy

We now look at the ratio $\bar{Y}_1/\bar{Y}_2$ of average throughputs. To simplify the expressions, let us assume that $\beta^{(1)} = \beta^{(2)} = 1/2$ and that $p^{(1)} = p^{(2)} = p \geq 1/2$. We also assume that the linear growth rates are inversely proportional to the square of the round trip times, i.e., $\eta_i = 1/(R^{(i)})^2$ for $i = 1, 2$. (Indeed, the window increases by one each RTT, and since the throughput is given by the window size divided by the RTT, the increase rate of the throughput is $1/RTT^2$.)

We then obtain from (4) that

$$\frac{\bar{Y}_1}{\bar{Y}_2} = \left(\frac{R^{(2)}}{R^{(1)}}\right)^2 \frac{3p\frac{R^{(2)}}{R^{(1)}} + 5p - 2p^{(12)}}{5p\left(\frac{R^{(2)}}{R^{(1)}}\right)^2 + 3p - 2p^{(12)}\left(\frac{R^{(2)}}{R^{(1)}}\right)^2} \ .$$

If we further assume that exactly one flow will experience a loss, then we have $p^{(12)} = 0$ and $p = 1/2$ giving

$$\frac{\bar{Y}_1}{\bar{Y}_2} = \left(\frac{R^{(2)}}{R^{(1)}}\right)^2 \frac{3\left(\frac{R^{(2)}}{R^{(1)}}\right)^2 + 5}{5\left(\frac{R^{(2)}}{R^{(1)}}\right)^2 + 3} \ .$$

We show the fairness in throughputs for the fixed loss strategy in Fig. 1(a). Note that the ratio of average throughputs is very close to be linear in the square of the ratio of round trip times (the dotted line depicts the function $y = 3x^2/5$).

4.2 The LTL Strategy

In Fig. 1(b) we depict the throughput ratios as a function of the ratio of the inverse of the square of RTTs for the LTL strategy. The values are obtained by computing the throughput as in Sect. 3.2.

We observe that although in general the throughput has a tendency to increase as the corresponding RTT decreases, we see that the throughput curve is quite irregular and fractal, and locally there are many points where the opposite behavior is observed: increasing the RTT of a connection results in increasing its throughput. This can perhaps be explained in part by changes in the periodicity of the steady-state behavior and in other discrete nature behavior. The analysis of this phenomenon is beyond the scope of this paper. We note that other fractal aspects of AIMD connections in networks with several nodes have already been reported in [12]. We finally observe that as the RTT of a connection becomes negligible with respect to the other, its share of the throughput converges to $5/4$ of the throughput of the other connection, as predicted in Sect. 3.3.

4.3 Comparisons

We first observe that the throughput sharing in the LTL strategy is much more fair than in the probabilistic sharing: it is much less sensible to the differences in RTT. Indeed, a connection with 3 times smaller RTT gets only 1.21 times more throughput in the LTL strategy, whereas it gets 6 times more throughput in the case of the constant probabilities strategy.

The fairness behavior of the proportional drop strategy has already appeared in [9], where the connection with 3 times smaller RTT gets 2.75 times more throughput. Comparing to these results we see that, in terms of fairness, the LTL strategy gives the best results whereas the worse performance is provided by the fixed loss probabilities strategy.

The behavior of the throughput as the ratio of RTTs goes to zero is in particular interesting. The throughput of the long connection and its share of the throughput tend to zero in the constant loss strategy, as well as with the proportional strategy [8, Sect. 7-8], whereas it tends to a positive constant under the LTL strategy.

Note that the fact that we obtain different average throughput sharing under different policies reflects the fact that, in contrast to the symmetric case, the throughput is not invariant with respect to the strategy in the general asymmetric case.

5 The Symmetric Case: Invariance of the Throughput for a General Strategy

Consider now a general strategy for deciding which connection will decrease its rate when capacity is reached. The decrease is by a constant β and the increase rate is η. We still restrict ourselves to the symmetric case of two connections, and assume that one and only one connection decreases its rate when the capacity is reached. At such a moment, connection 1 that transmits at a rate of y will decrease its rate with probability $f(y)$ and connection 2 will decrease its rate with probability $1 - f(y)$. We assume that the rate process of both connections is in a stationary ergodic regime. In particular we shall focus again on $Y_n^{(1)}$, the rate of connection 1 just before a rate decrease occurs.

Let us state one of the main results of the paper in the following proposition.

Proposition 2. *The average throughput $\bar{Y}$ of a connection in a symmetric network with two connections is given by*

$$\bar{Y} = \frac{1+\beta}{3+\beta} C \ , \tag{6}$$

independent of the sampling function f.

Proof. The proof is quite involved, so for space reasons we will roughly sketch its different steps; the full proof can be found in [10].

First, we focus in the throughput process $Y_n := Y_n^{(1)}$. In particular, we compute the cumulative throughput S between congestion epochs for two cases (y denotes the state of Y_n at time T_n):

- Connection 1 is the one to decrease its rate (this happens with probability $f(y)$), in which case we have: $S = \frac{1}{2}(\beta Y_n + Y_{n+1})\tau_{n+1} = y^2 \frac{(1+3\beta)(1-\beta)}{8\eta}$.
- Connection 2 is the one to decrease its rate (this occurs with probability $1 - f(y)$). In this case, $S = \frac{1}{2}(Y_n + Y_{n+1})\tau_{n+1} = \frac{1-\beta}{8\eta}(-(3+\beta)y^2 + 2C(1+\beta)y + C^2(1-\beta))$.

Next, we compute the expected time interval between congestion epochs: $\mathbb{E}[\tau] = \mathbb{E}\left[Y\frac{1-\beta}{2\eta}f(Y) + \frac{(C-Y)(1-\beta)}{2\eta}(1-f(Y))\right] = (1-\beta)\mathbb{E}[Yf(Y)]/\eta$, where Y denotes a random variable distributed like Y_n at steady-state. Remark that $\mathbb{E}[\tau]$ depends on the expectation $\mathbb{E}[Yf(Y)]$.

The average throughput $\bar{Y}$ is given by: $\bar{Y} = \mathbb{E}[S]/\mathbb{E}[\tau]$. Hence, we need to compute $\mathbb{E}[S]$, which can be expressed as: $\mathbb{E}[S] = \frac{1-\beta}{8\eta}\Big(-(3+\beta)\mathbb{E}[Y^2] + 4(1+\beta)\mathbb{E}[Y^2 f(Y)] - 2C(1+\beta)\mathbb{E}[Yf(Y)] + C^2\frac{3+\beta}{2}\Big)$.

We then obtain three expressions relating both the three unknowns that appear in the formula of $\mathbb{E}[S]$, that is: $\mathbb{E}[Yf(Y)]$, $\mathbb{E}[Y^2]$, $\mathbb{E}[Y^2 f(Y)]$, and the quantity $\mathbb{E}[Y^3]$ (hence, we have four unknowns and three equations). It happens that $\mathbb{E}[S]$ can be expressed as a function only of $\mathbb{E}[Yf(Y)]$, that is: $\mathbb{E}[S] = \frac{(1-\beta)(1+\beta)\mathbb{E}[Y(f(Y)]C}{\eta(3+\beta)}$. Therefore, the term $\mathbb{E}[Yf(Y)]$ cancels out when dividing $\mathbb{E}[S]$ by $\mathbb{E}[\tau]$, which gives (6).

□

6 The Symmetric Case: Second Moment of the Throughput

Even if all possible loss strategies provide the same average throughput in steady-state in the symmetric case we can wonder about the variability of the throughput. In real-time applications that may use AIMD protocols in order to be TCP-friendly, it is clearly advantageous to have the lowest possible throughput variability.

The following Proposition gives a general expression for the second moment of the throughput. As will be seen, this expression is not invariant any more, in contrast to the first moment.

Proposition 3. *Let* $\mathbb{E}[S_2]$ *denote the mean cumulative of the square throughput between two loss epochs. The (average) second moment of throughput is*

$$\frac{\mathbb{E}[S_2]}{\mathbb{E}[\tau]} = \frac{1}{8}(3 + 2\beta + 3\beta^2)\frac{\mathbb{E}[Y^3 f(Y)]}{\mathbb{E}[Y f(Y)]} \ .$$

The proof of this proposition, which follows along the same lines as that of Proposition 2, is given in [10].

Since we still have two unknowns, one could argue that their ratio is constant. Actually, it is not the case from the following proposition where we compare the second order moment for the three loss strategies (constant, proportional or largest flow).

Proposition 4. *Let* $\beta = 1/2$*. Using the constant loss probability scheme, we get*

$$Q_{cst} = \frac{\mathbb{E}[S_2]}{\mathbb{E}[\tau]} = \frac{95}{448}C^2 \approx 0.21C^2$$

whereas when the loss is applied to the largest flow (LTL strategy), we have

$$Q_{ltl} = \frac{\mathbb{E}[S_2]}{\mathbb{E}[\tau]} = \frac{4}{21}C^2 \approx 0.19C^2$$

and the scheme with proportional losses gives

$$Q_{pro} = \frac{\mathbb{E}[S_2]}{\mathbb{E}[\tau]} = \frac{469373}{1467072}C^2 \approx 0.32C^2 \ .$$

We see from the proposition that in the symmetric case, the LTL strategy is to be preferred (in terms of lower second moment), whereas the strategy of losses proportional to the throughput has the worse performance. The proof of this proposition, omitted for space reasons, is given in [10].

7 Discussion and Future Research

We have introduced in this paper various loss strategies that determine which connection will lose a packet when a congestion occurs. We have shown that such loss strategies may have a considerable impact on the throughput variability (which may be an important performance measure in real-time applications that use AIMD protocols to be TCP-friendly) but that they all lead to the same average throughput in the special case of a symmetric network with two connections. Among three specific strategies that we introduced, we have shown in the above setting that the LTL strategy (i.e., the strategy that drops a packet from the connection with highest throughput) has the best performance in terms of throughput variability, and moreover, it guarantees a positive share of the throughput even when the RTT of one of the connections becomes arbitrarily large.

The mathematical study of the sharing of bandwidth under various loss strategies turns out to be quite involved. So far we have not been able to get explicit expressions for the asymmetric network with two connections when the LTL or the proportional loss strategies are used. We have provided however an (involved) explicit expression for the throughput for the case of constant loss strategy. For the symmetric case, however, we have obtained an explicit expression for the throughput under an arbitrary loss strategy.

Many open problems remain: 1. Is there any probabilistic argument that can explain the invariance of the average throughput in the loss strategy phenomenon in the case of two connections? 2. Does the invariance of the throughput holds for the case of more than two competing symmetric connections? 3. What is the reason for the fractal behavior of the throughput sharing under LTL? 4. How to implement LTL? Note that a desirable way of implementation should be stateless, and it should make use only of local information available at the bottleneck element.

References

1. Padhye, J., Firoiu, V., Kurose, J., Towsley, D.: Modeling TCP Throughput: A Simple Model and its Empirical Validation. In: Proceedings of ACM SIGCOMM 98. (1998) 303–314
2. Ait-Hellal, O., Altman, E., Elouadghiri, D., Erramdani, M., Mikou, N.: Performance of TCP/IP: The Case of Two Controlled Sources. In: Proceedings of ICCC'97, Cannes (1997) 469–477
3. Brown, P.: Resource Sharing of TCP Connections with Different Round Trip Times. In: IEEE Infocom. (2000)
4. Lakshman, T., Madhow, U.: The Performance of TCP/IP for Networks with High Bandwidth-Delay Products and Random Loss. IEEE/ACM Transactions on Networking (1997)
5. Zhang, L., Shenker, S., Clark, D.: Observations on the Dynamics of a Congestion Control Algorithm: The Effects of Two-Way Traffic. In: ACM SIGCOMM. (1991)
6. Baccelli, F., Hong, D.: AIMD, Fairness and Fractal Scaling of TCP Traffic. In: Proceedings of IEEE INFOCOM 02. (2002)
7. Vojnovic, M., Le-Boudec, J.Y.: Some Observations on Equation-Based Rate Control. In: Proceedings of ITC-17, Salvador da Bahia (2001)
8. Altman, E., Jiménez, T., Núñez-Queija, R.: Analysis of Two Competing TCP/IP Connections. Performance Evaluation **49** (2002) 43–56
9. Altman, E., Barakat, C., Laborde, E., Brown, P., Collange, D.: Fairness Analysis of TCP/IP. In: Proceedings of IEEE CDC'00, Sydney (2000)
10. Altman, E., El Azouzi, R., Ros, D., Tuffin, B.: Loss Strategies for Competing TCP/IP Connections. Research report RR-5100, INRIA (2004)
11. Altman, E., Avratchenkov, K., Barakat, C.: A Stochastic Model of TCP/IP with Stationary Random Losses. In: Proceedings of ACM SIGCOMM 2000. (2000)
12. Baccelli, F., Hong, D.: Interaction of TCP Flows as Billiards. Research Report RR-4437, INRIA (2002)