

TCP/IP modeling and validation

Chadi Barakat

INRIA, 2004 route des Lucioles, 06902 Sophia Antipolis, France

Email : cbarakat@sophia.inria.fr

Abstract

We discuss in this paper the different issues to be considered when modeling the TCP protocol in a real environment. The discussion is based on measurements we made over the Internet. We show that the Internet is so heterogeneous that a simplistic assumption on TCP congestion control or on the network may lead to erroneous results. We outline some of our results in this field and we present a novel approach for a correct validation of a model for TCP.

I. INTRODUCTION

The modeling of TCP congestion control is an important task for improving the service provided to Internet users and the efficiency of network resource utilization. This importance is justified by the large amount of TCP traffic in the Internet: 95% of total bytes and 85%-95% of total packets are of TCP type. A TCP source controls the rate of application packets as a function of the way the network treats or reacts to these packets [13]. The main objective of modeling is to come with simple expressions of TCP throughput for a certain network reaction. This has two main advantages. First, it determines the factors impacting the performance of the protocol which gives people working on TCP insights on how the congestion control has to be improved. For example, the modeling has shown that in case of drop-tail buffers and synchronized TCP flows, the throughput of a TCP connection is inversely proportional to the square of the average round-trip time (RTT) [14]. This has given an explanation to and an evaluation of the bias of TCP against connections with long round-trip times. Using this result, [9] has proposed a modification to the standard window increase algorithm of TCP in order to improve the fairness of the protocol. The modeling has also shown that dropping packets randomly in network routers as done by RED (Random Early Detection) [12] improves the fairness of TCP by making the throughput of a connection inversely proportional to the average round-trip time instead of its square [3], [14], [15], [21]. Another possible use of an expression of TCP throughput could be the study of a change of the parameters of TCP congestion control so as to minimize the variations of the window without adding to the aggressiveness of the protocol. This will be useful for multimedia applications using TCP, or possibly a new version of TCP adapted to such applications.

The second advantage of TCP modeling is that it permits network designers to improve the reaction of their network to incoming packets at the moment of congestion, given the current control policy of TCP. Important work has been done in this direction (e.g., [5], [8], [18], [17], [19]). The focus was and is always on the dimensioning of buffers in network routers and on the management of their occupancy. The objectives [12] are to maximize the utilization of network resources, to protect the TCP transfers from each other and to minimize the queueing time in routers so that delay sensitive applications (e.g., Telnet, short WEB transfers) get a better service. A typical problem in this direction is the tuning of RED parameters [8], [5], [17]. RED [12] is an interesting buffer management technique presented as a solution to achieve the above objectives. For example, using the expression of TCP throughput calculated in [21], it has been shown in [8] that when the number of TCP connections exceeds a certain threshold, the RED buffer gets into an instable regime. This instability has been explained by the sudden jump in the drop probability when the average queue length reaches the maximum threshold, see [12] for more details on RED parameters. This has motivated the introduction of a new parameter to RED (the *gentle_* parameter) in order to avoid such jump. Other works have used the explicit expressions of TCP throughput to improve other parts of the network as the link layer on a wireless interface [6].

Recently, a new application of TCP modeling has emerged [10], [11]. It consists in using the explicit expressions of TCP throughput to control the rate of real time flows (e.g., an audio flow) in a TCP-friendly way. The reaction of the network (e.g., packet drop probability) is averaged over a certain time interval and the rate of the real time flow is set to the expected throughput of a TCP connection running in the same network conditions. Clearly, a tradeoff exists between choosing a long averaging time interval to get a smooth variation of the rate — hence a low jitter and a good quality — and choosing short intervals to get a fast reaction to changes in network conditions.




Fig. 1. Measurement testbed

The modeling of TCP can be seen as a two-step procedure. First, we find the modeling of the evolution of TCP window between congestion events as well as upon congestion. This includes the modeling of some particular mechanisms of TCP such as timeouts and the limitation of the rate due to the window advertised by the receiver [13]. Second comes the modeling of network reaction which consists, directly or after some transformation, in a modeling of the process of congestion events. The TCP connection lifetime can be seen to be a succession of congestion events between which TCP increases its window and upon which it reduces it. As we will see later, a model for the network also has to describe how the round-trip time of the connection varies. The variation of the round-trip time determines the variation of the window between congestion events. Note that TCP increases its window by the same amount during a round-trip time regardless of its length. An example of congestion events and window evolution is depicted in Figure 4.

If the network is well defined (e.g., a single RED buffer), the modeling of TCP can be achieved with a considerable degree of correctness. One can come with an accurate model for the network, combine it with a model for TCP, and solve the overall model for TCP performance. All the works assuming a single bottleneck router crossed by only TCP connections are of this type (e.g. [8], [18]). The difficulty exists when we want to approximate the throughput of TCP in the real Internet. Some assumptions are required in this latter case to model the reaction of this huge and heterogeneous environment. The difficulty of the analysis as well as the accuracy of the modeling change with the assumptions we make. For a given assumption, we must expect our results to hold on some Internet paths while not on others. The results may hold due to the correctness of our modeling on these particular paths. They may also hold due to another phenomenon that we observed during our work in this direction, that is the cancellation of errors introduced by the different blocks of a model for TCP. A good model for TCP must have all its blocks validated separately and these blocks must work well on a wide range of Internet paths.

We address in this paper the problem of TCP modeling in its general form. We focus mainly on the modeling of TCP in a real network rather than a particular environment as a single Drop Tail buffer or a RED buffer. Based on measurements we conducted over the Internet, we discuss the different issues to be considered for a correct modeling of TCP. We outline the different approaches in this domain and we show the need for a general approach that copes with the heterogeneity of network reaction we observed. We will see how the heterogeneity of the Internet can change the performance of a model from one path to another, and even from one hour to another on the same path. We introduce the notion of separate validation of each block of a model for TCP instead of an overall model validation. We also discuss some issues related to the use of throughput expressions in practice, particularly in TCP-friendly applications.

In the next section, we present our measurement testbed that we will refer to along the paper. In Section III, we discuss the issues related to the modeling of TCP window evolution. The issues related to network modeling are discussed in Section IV. In Section V, we explain our technique for the validation of a model for TCP. The paper is concluded in Section VI.

II. MEASUREMENT TESTBED

Our testbed consists of three long-life TCP connections run between a machine at INRIA Sophia Antipolis in the south of France and three other machines over the Internet (Figure 1). The first machine is located at ESSI next to INRIA. The second machine is located at Paris at about 800 Km from INRIA. The third machine is located at the University of South Australia. These can be considered as respectively LAN, MAN and WAN environments. In the

sequel, we use the terms LAN, MAN and WAN to distinguish between the three connections. The source is located at INRIA and runs the New-Reno version of TCP [7]. New-Reno is a recent version of TCP able to recover from multiple packet losses without timeout and with only one division of the window by two. The source of our connections is fed by a simple application that has always data to send. The connections were run for many hours in different days during the month of January 2000. We developed and we ran a tool at INRIA that looks at the trace of every connection (packets and ACKs) and calculates a certain number of statistics on the connection such as the total number of packets acknowledged, the total number of retransmissions, the moments at which the window is reduced, the variation of window and round-trip time over time, etc. We accounted for all the mechanisms of New-Reno when developing our tool, particularly the fact that a New-Reno source can recover from multiple packet losses in the same window of data. We stored the statistics at fixed intervals in separate files. We chose these intervals in a way to get enough data in each file. We shall assume that the network conditions are approximately the same during each interval. These intervals are set to 20 minutes for LAN, 40 minutes for MAN, and 60 minutes for WAN.

III. MODELING TCP WINDOW EVOLUTION

Most of the effort on TCP modeling [3], [14], [15], [16], [17], [18], [21] is devoted to long transfers, namely to the congestion avoidance mode of TCP [13]. The slow start phase [13], due to its fast window increase and its short duration, is often ignored. The window is assumed to increase linearly in time between congestion events and to decrease to half its size upon congestion. The moments of congestion are determined by the underlying model for the network. These are the moments at which the source detects the loss of a packet and decides to reduce its window. Ideally, and this is what the new versions of TCP try to approximate [7], these must be the moments of detection of the first loss in a window of data packets.

TCP is known to increase its window in congestion avoidance mode in a linear manner but as a function of round-trip time number rather than time. This increase is approximately equal to $1/b$ packets per round-trip time, where b is the number of packets covered by an ACK [21]. For the time linear increase to hold, the round-trip time is often supposed to be constant or to vary independently from the window size, so it is substituted in the analysis by its average during the transfer [14], [15], [21].

For the multiplicative decrease factor, it is indeed equal to one half when the source receives three duplicate ACKs for congestion detection and when the Fast Recovery phase succeeds [7]. If it is not the case, a timeout occurs, the window is set to one packet, and slow start is called to reach quickly the slow start threshold [13]. Some authors [3], [21] keep the decrease factor in the case of timeout equal to one half while others [16] set the window to one packet and assume a linear increase from this low value. We believe that there will be no difference between the two approaches in the future given the expected low probability of timeout with the new enhancements proposed to TCP congestion control [1]. Our measurements showed that on WAN, most of the losses are detected with timeout and this is due to the small size of TCP window. The same result is noticed in [21]. On LAN and MAN, the window is large and the timeout phenomenon is quite absent. In our work on TCP modeling [3], we set the multiplicative decrease factor to a constant value ν ($\nu < 1$) for both kind of congestion detection methods. This factor ν , together with the general window increase rate α we considered, permits our results to be used for the study of other congestion control policies.

Clearly, the above model is simple since it only accounts for the linear-increase multiplicative-decrease part of TCP congestion control. Other issues have to be considered for a better modeling of TCP. We will enumerate those that are the most important in the following sections and explain when possible how they can be introduced into the above model.

A. Dependency between window and round-trip time

On paths where the window of the TCP connection is small compared to the bandwidth-delay product and where packets cross multiple congested routers, one should expect that the assumption on the linearity of the window increase will hold. Indeed, on such paths, the connection does not contribute to the queuing time in network routers and the congestion is mostly caused by other aggressive connections. The round-trip time varies then independently from the window size and can be substituted in the analysis by its average [21] resulting in a time linear window growth at a rate $1/(bRTT)$, with RTT being the average round-trip time and b being the number of packets covered by an ACK. However, this independency is not expected to hold on paths where the window of the connection is large compared to the bandwidth-delay product. The increase in the window in this latter case will result in an increase in the round-trip

time which in turn will result in a sub-linear increase in the window with time [14]. One should expect that, in presence of a linear model for TCP window evolution that uses the average round-trip time for the calculation of the window slope, this sub-linearity will result in an overestimation of the real throughput.

To understand such dependency, we plot the variation of the round-trip time as a function of the window size on our WAN and LAN connections (Figures 2 and 3). The WAN connection is a typical example of the first connection whereas the LAN connection is a typical example of the second one. We measure the round-trip time for one packet per window of packets and we note the window size upon this measurement. We average then the round-trip times obtained for close window sizes. This gives the thick line in both figures. We see clearly how on WAN, the round-trip time is on average constant and independent of the window. However, it is an increasing function of the window on LAN and a sub-linear window increase should be seen on this connection. Indeed, Figure 4 shows well this sub-linear behavior. We plot in this latter figure the variation of the real window on LAN for some seconds. The straight line corresponds to the expected window evolution if the round-trip time were independent from the window size. This line is given by the time linear model for TCP window evolution.

The problem with the sub-linearity of window increase is that it complicates the analysis and makes difficult (if not impossible) to obtain simple explicit expressions for TCP throughput. The works in the literature that consider such sub-linearity (e.g., [17], [18]) give the throughput as a solution of some implicit differential equations. Only the works considering linear models (e.g., [3], [21]) succeed to find closed-form expressions for TCP throughput. Moreover, it is not clear for the moment how such sub-linearity can be modeled in a complicated network as the Internet. The importance of the sub-linearity depends on many factors as the intensity of exogenous traffic sharing the path with the TCP connection, the buffering capacity in network routers, and the available bandwidth on the path. In a real environment, one has to look at the traces of round-trip time and window and try to infer from these traces some model for the dependency between round-trip time and window size. In previous works [17], [18], this inference was not necessary since the content of the network was well defined (e.g., a single RED buffer with given bandwidth and propagation delay) and hence the variation of the round-trip time as a function of the window size was well known. The model for round-trip time, together with the model for window evolution and the model for network reaction, provide the required information for the calculation of TCP throughput. For example, one can use the models developed in [4], [17] for such calculation.

A model for round-trip time variation

Using the traces of our LAN connection, we propose a technique for the inference of the dependency between window size and round-trip time. First, we have to define a model for the variation of the round-trip time. Then, we have to use some fitting technique to infer the parameters of such model from the traces of the connection. It seems that the most appropriate model for round-trip time variation is the one used in the literature for networks of one router (e.g., [14]). Denote by W the window of the TCP connection. The model considers that the round-trip time is constant whenever W is less than a certain window W_0 . For windows larger than W_0 , the round-trip time increases linearly with W at a rate β . We write the model as follows

$$RTT_e = RTT_0 + 1\{W > W_0\}\beta(W - W_0).$$

RTT_0 can be seen as the contribution of the propagation delay and the exogenous traffic to the total round-trip time. $\beta(W - W_0)$ can be seen as the contribution of the connection to the total round-trip time when its window is large. $1/\beta$ represents the bottleneck bandwidth on the path of the connection.

Given this model for the round-trip time, we use the non-linear least-square technique to infer the three parameters RTT_0 , β and W_0 of the model from the traces of the connection. Let RTT_m be a round-trip time measurement that corresponds to a window W . We can write RTT_m as follows

$$RTT_m = RTT_e + \epsilon,$$

where ϵ is a certain error. Let ϵ_n be the error introduced by the n th round-trip time measurement. We propose to find the parameters of the model that minimize the sum $\sum_{n=1}^N \epsilon_n^2$, where N is the total number of round-trip time measurements.

We solve the minimization problem for the traces of our LAN connection shown in Figure 3. We get the thick line in Figure 5. We also plot in Figure 5 the 95% confidence intervals as well as the measurements we obtained. Our results

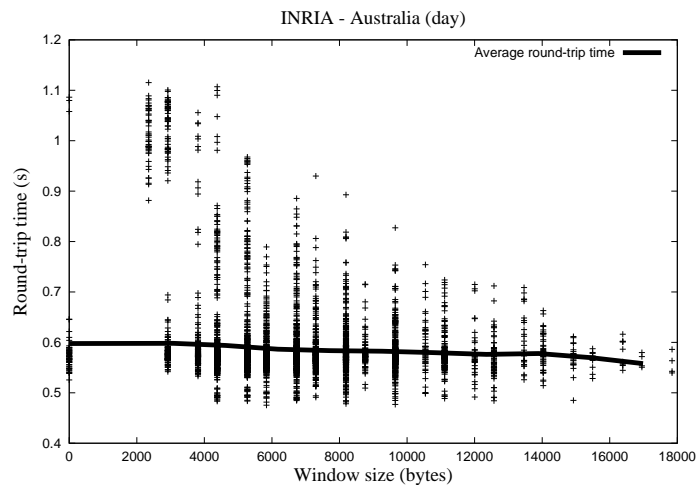


Fig. 2. WAN: Round-trip time vs. window




Fig. 3. LAN: Round-trip time vs. window

show clearly that there is a dependency between window and round-trip time on this particular path. For small windows, the expected round-trip time is constant and equal to 0.35s. Once the window of the connection exceeds 2246 bytes, the expected round-trip time starts to increase linearly by 0.0493s every 10000 bytes.

B. Modeling timeouts and Fast Recovery

The above modeling does not consider these two mechanisms of TCP [7], [13]. Due to the coarse granularity of TCP timers (500 ms in most implementations), the first mechanism introduces a certain idle time between a congestion and its detection. The second mechanism also introduces a certain time between the detection of a congestion with duplicate ACKs and the resumption of the window increase. During this latter time, the source is supposed to recover from losses and to transmit new packets in order for the ACK clock not to stop.

The modeling of these two mechanisms of TCP require a detailed description of the protocol behavior at the packet level. This behavior is quite complicated to model and varies from one version to another [7]. The Fast Recovery phase is always ignored due to the complexity of its modeling, also because it adds a negligible contribution to the throughput when it works well. It is assumed in [21] that once the source receives three duplicate ACKs, the window resumes its linear increase until the next congestion event. This seems reasonable with the new versions of TCP (e.g., SACK [7]) where the Fast Recovery phase is quite robust and fast, but it is not true with the other versions as Reno [7] where Fast Recovery may fail due to multiple losses per window. A failure of Fast Recovery results in a timeout and a slow start from a window of one packet. A model which does not account for the Fast Recovery phase will lead to throughput




Fig. 4. LAN: Window vs. time




Fig. 5. LAN: Expected round-trip time vs. window

overestimation if such failures are frequent.

The timeout mechanism is studied in [21] using the probability that a packet is dropped. Denote this probability by p . The authors focus on the calculation of the probability that the source fails to receive three duplicate ACKs to trigger Fast Recovery. This has been assumed to be the necessary and sufficient condition for the occurrence of a timeout. Timeouts are assumed not to occur during the Fast Recovery phase. We believe that this will be the case with the new versions of TCP that use the SACK option [7] and that have a robust and a quick Fast Recovery phase. We also believe that with the modifications proposed to TCP error recovery in [1], the failure to receive three duplicate ACKs will no longer be a sufficient condition to get a timeout. Sources will be able to recover from losses when more than one ACK are received per round-trip time. In the wait for these modifications to be deployed, we will show in the following a heuristic for the addition of the timeout mechanism to the above model according to the condition in [21].

To simplify the analysis, it is assumed in [21] that when a packet is dropped, the subsequent packets in the same window are also dropped. The probability to get a timeout ($Q(p)$) as well as the average duration of a timeout period ($Z(p)$) are calculated as a function of p ,

$$Q(p) = \min \left(1, \frac{(1 - (1 - p)^3)(1 + (1 - p)^3(1 - (1 - p)^{W-3}))}{1 - (1 - p)^W} \right),$$

$$Z(p) = T_0 \frac{1 + p + 2p^2 + 4p^3 + 8p^4 + 16p^5 + 32p^6}{1 - p},$$




Fig. 6. Model without timeout intervals




Fig. 7. Model with timeout intervals

with

$$W = \frac{2+b}{3b} + \sqrt{\frac{8(1-p)}{3bp} + \left(\frac{2+b}{3b}\right)^2},$$

and T_0 being the basic timeout interval which is doubled when the source backs off its timer after the loss of a retransmission [13]. In practice, $4RTT$ has been shown to be a safe value for T_0 [11].

Using these two functions of p , the impact of timeout intervals can be introduced into the model in the following way. We calculate first the throughput of the connection when excluding these intervals, in other words when assuming that the window resumes its linear increase directly after a congestion event whatever is the method of detection (Figure 6). Denote this throughput by \bar{X}_d . Denote by \bar{X}_f the throughput of the connection in the presence of timeout intervals (Figure 7). The throughput in both cases is equal to the ratio of the average number of packets that cross the network between two congestion events ($\simeq 1/p$) and the average time between congestion events, see [3], [21] for details. Let $\{S_n\}$ denote the process of times between congestion events when timeout intervals are excluded (Figure 6). Let $\{S'_n\}$ denote the process of times between congestion events when timeout intervals are considered (Figure 7). Given that the average number of packets that cross the network between two congestion events is the same in both cases ($\simeq 1/p$), it follows that

$$\bar{X}_f = \bar{X}_d \frac{E[S_n]}{E[S'_n]},$$

and

$$\bar{X}_d = \frac{1/p}{E[S_n]}.$$

Using these two equations and the fact that

$$E[S'_n] = E[S_n] + Q(p)Z(p),$$

we can write

$$\bar{X}_f = \frac{\bar{X}_d}{1 + p\bar{X}_d Q(p)Z(p)}.$$

Note that our expressions for the throughput are in terms of packets/s. Note also that our throughput corresponds to the rate at which TCP packets leave the network, also called the receiving rate.

C. Modeling window limitation

Another problem with the previous model is that it does not consider the fact that a TCP source cannot inject into the network more packets than the window advertised by the receiver [21]. This phenomenon can be seen in Figure 4 where we plot the number of bytes in the network as a function of time. The model becomes sub-linear and the calculation of a simple expression for the throughput seems to be impossible except for some particular models for the network [4], [16], [21]. For example, the calculation is straightforward when the time between congestion events is constant. Using

techniques from queueing theory, the authors in [4], [17] succeed to calculate the throughput when congestion events appear according to a homogeneous Poisson process.

For more general models for the network (e.g., a generally distributed time between congestion events), one can always think about finding some approximations of the throughput. One possible approximation is to assume that the receiver window is always reached between congestion events. The problem is then automatically transformed into a simpler one with constant times between congestion events equal to the average of these times. This is the kind of approximation used in [21]. In [3], we find bounds for the throughput which are also a good approximation. The advantage of our bounds is that they are valid for any receiver window. The first approximation is only valid for small receiver windows, so one needs to condition on the average window upon congestion in order to use either the approximation or the expression of the throughput obtained when there is no window limitation.

Note that the problem of window limitation exists because of the small windows advertised by current Internet receivers (less than 64 Kbytes). But, this is supposed to disappear in the future with the trend to increase the window field in the TCP header and with the trend to change dynamically the buffer size allocated to the TCP connection at the receiver [23]. It is even recommended in [2] that future studies on TCP congestion control should consider an infinite receiver window. Models that account for the limitation of the rate will not be of big importance.

D. Fluid models versus discrete models

Some of the models for TCP assume that the window increases continuously between congestion events [15], [16], [17], [22]. The use of a continuous model for the window facilitates the analysis since it permits the use of tools from the theory of continuous functions as integration and differentiation. The continuous increase may hold for the congestion window at the source but it does not hold for the volume of data in the network. Indeed, the former quantity increases in small values (in bytes) upon ACK arrivals, whereas the latter one increases in steps of one packet every b round-trip times [21]. This is due to the Nagle algorithm [20] which prohibits the TCP source from injecting small packets into the network. At the moment of congestion, it is the window at the source which is divided by two rather than the number of packets in the network [13]. So models assuming continuous increase in the window, also called *fluid* models, are appropriate for the prediction of the window variation at the TCP source. The throughput obtained with these models is the throughput that the TCP connection would realize if it is not limited by the Nagle algorithm.

The correct throughput can be obtained by one of two ways: either by using a discrete model for the window evolution as the one in [21], or by using a fluid model and introducing some corrections into the calculated throughput so as to account for the difference between the congestion window at the source and the volume of data in the network. In our work on TCP modeling [3], [4], we chose the second way for the calculation of the throughput given the simplicity of the analysis the use of fluid models permits. We present here the technique we used in order to account for the packet nature of TCP. To illustrate this technique, we plot in Figure 8 the number of packets that cross the network between two congestion events for both a fluid model and a discrete (or a packet) model. The line for the discrete model is taken from [21]. On average, half of the window is assumed to be dropped in the network upon congestion (the last round-trip time in the figure). Suppose that all the transmission rates are expressed in terms of packets/s. Denote by $E[W_n]$ the average window size upon congestion and by $E[S_n]$ the average time between congestion events. Let \bar{X} denote the throughput obtained with the fluid model and \bar{X}_d the real throughput of TCP. A good approximation of \bar{X}_d can be obtained by

- Shifting down the fluid window by 0.5 packet which results in a decrease in the throughput \bar{X} by $0.5/RTT$.
- Subtracting the rate of dropped packets which is approximately equal to

$$\text{Rate of dropped packets} = \frac{1}{2} \frac{E[W_n]}{E[S_n]}.$$

It follows that

$$\bar{X}_d = \bar{X} - \frac{0.5}{RTT} - \frac{0.5E[W_n^*]}{E[S_n]}.$$

Using one of our results [3, Equation (3)] which says that

$$E[W_n] = \frac{E[S_n]}{bRTT(1-\nu)},$$




Fig. 8. Fluid model vs. discrete model

we can write

$$\bar{X}_d = \bar{X} - \frac{0.5}{RTT} \left(1 + \frac{1}{b(1-\nu)} \right).$$

With this correction and as we will see in Figures 15 and 16 when validating our model, a fluid model for TCP is able to give the same throughput as a detailed discrete model.

IV. MODELING THE NETWORK

This is the part of the model where the heterogeneity of the Internet has the greatest impact. The objective of this part is to find a good characterization of congestion moments or loss moments. A loss moment in our terminology is the moment at which the TCP source decides that the network is congested and that it must reduce its window. These moments can be directly characterized by making some assumptions on the way they appear. For example, one can assume that they appear according to a deterministic process or a Poisson process. This direct characterization is the approach widely used in the literature [3], [15], [16], [21]. Another possible but indirect characterization consists in finding a model for the reaction of the network to a particular TCP packet or during a small time interval. As an example of this indirect approach we find the works that suppose that TCP packets are dropped within the network with a constant or variable probability [18], or those assuming that loss moments form a Poisson process with a variable intensity function of the window size of the connection [22]. The advantage of the indirect approach is that it decouples the model for the network from the control policy at the TCP source. It happens that on some Internet paths, the process of losses seen by a TCP connection is a function of the way it increases and decreases its window, and that this process changes if another control policy is used, for example if packets are transmitted at a constant rate. The indirect approach is very useful on such paths since it permits to deduce the loss process that a TCP connection will see from the loss process seen by another connection with a different congestion control policy. We cite different applications of the indirect approach on such paths. The first application is that one can probe the network with a certain flow of packets (e.g., a constant rate flow) and calculate the parameters of an appropriate model for the network (e.g., packet drop probability, variation of loss intensity as a function of transmission rate). With these parameters, it will be possible to predict the performance of a TCP connection on the same path. Another application of the indirect approach is that from the trace of a TCP connection, one can build a model for the network and predict the performance of another TCP connection with another congestion control policy. This will be useful for a study of the impact of a change of TCP congestion control parameters. A third application of the indirect approach is that a TCP-friendly application (e.g., TFRC [11]) can deduce the loss process that a TCP flow would see from the loss process it sees, and hence can get a better estimate of the rate to use. All these applications are not possible with the direct approach since, without a model for network reaction, we cannot deduce the performance of a TCP connection from the loss process seen by another connection with another control policy. Recall that we are talking about paths where the loss process changes with the congestion control policy.

The difficulty with the indirect approach is in the definition of a correct model for the Internet and the calculation of

its parameters. This may be possible for a simple network of one router (e.g., a RED buffer is known to drop packets with a probability that increases linearly with the average queue length [12]) but it seems to be difficult for a wide network as the Internet. It is not clear if there exists a model for the Internet that works on all paths. Certainly, the Internet reaction to TCP packets changes from one path to another and along time on the same path. For example, on some paths the network may drop packets with a constant probability, on other paths with a probability that increases linearly with the congestion window, on other paths with a probability that increases logarithmically with the congestion window, etc. One can imagine different models for the network. The question is on how many paths such models are useful.

Given the difficulty to come with a model for network reaction that works on all Internet paths, we shall only focus on the direct approach. Recall that the direct approach consists in using the parameters of the loss process seen by the connection for throughput calculation. Our TCP connections are using the same control policy and thus we can decide on the heterogeneity of the Internet from the loss processes they see. In the next section, we present some measurement results to show how much the Internet is heterogenous. In the following section, we use the expression of TCP throughput we found in [3] to evaluate the impact of the assumption on the network we make. Note that in [3], we found a simple expression of TCP throughput for a general process of congestion events.

A. Diversity of loss processes in the Internet

We present in this section some of the loss processes we found on our three connections. First, we plot the distribution of times between congestion events. Figures 9, 10, 11, 12, 13 and 14 show some samples of the distributions we got. The figures contain for comparison some theoretical distributions (Exponential for LAN and WAN, Normal for MAN). On LAN, the process is highly bursty which results in this pulse close to the origin. This burstiness can be also seen in Figure 4 where the window is divided multiple times by two in short time intervals. We noticed that the congestion on our LAN connection stays for multiple consecutive round-trip times during which the network keeps dropping packets and the source keeps reducing its window. On MAN, the times between losses follow well a Normal distribution. On WAN and as one must expect, the loss process is close to Poisson. Indeed, on WAN, the source has a small window and does not contribute to the congestion of the network. The loss process it sees is the superposition of a large number of processes in all the routers it crosses. We also found some correlation of losses on the three connections. This correlation varies during the day between negative and positive values with an absolute value of the covariance coefficient reaching sometimes 0.2. Tables I, II and III show some of the covariance coefficients we saw on the three connections at different hours during the day. Recall that the covariance coefficient (of order 1) of a process $\{S_n\}$ is given by,

$$Cov = \frac{E[S_n S_{n+1}] - E[S_n]^2}{E[S_n^2] - E[S_n]^2}.$$

This coefficient varies between -1 when inter-loss times are highly negatively correlated and +1 when they are highly positively correlated.

Our measurements show some examples of losses that a TCP connection may see over the Internet. We expect to see other processes on other paths. Other distributions of inter-loss times could be found. Also, some paths as those including satellite and wireless links, may exhibit more memory than our paths which will result in a more important correlation of losses. Using one of our results in [3], we discuss in the next section the impact of the different parameters of a loss process on TCP throughput. This will give us an idea on how much a certain assumption on the occurrence of losses or congestion events impacts the result of the modeling.

B. Impact of the loss process choice on TCP modeling

Most of existing models for TCP make simplistic assumptions on inter-loss times (e.g., deterministic [15], [21], Poisson [16]). In [3], we propose a general model for the network where we only make the simple assumption that the loss process is stationary and ergodic. We allow the inter-loss times to follow any distribution and we allow the losses to be correlated. We derive a general expression for the throughput which uses, in addition to the packet loss probability (p) and the average round-trip time (RTT), the variance of inter-loss times as well as all the covariance function of the




Fig. 9. LAN: Inter-loss time distribution

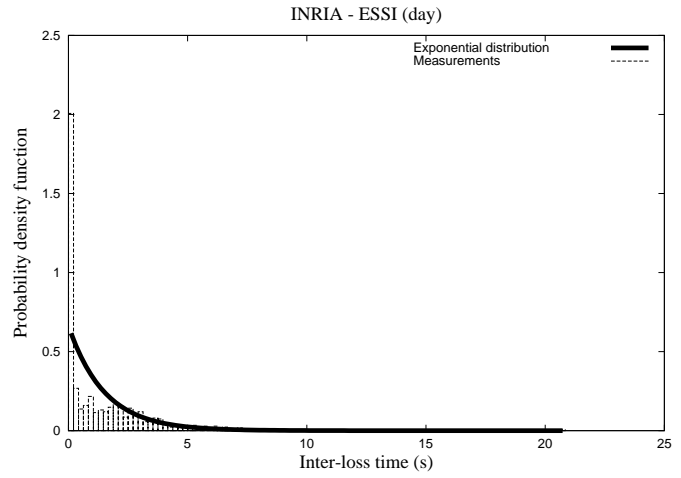


Fig. 10. LAN: Inter-loss time distribution

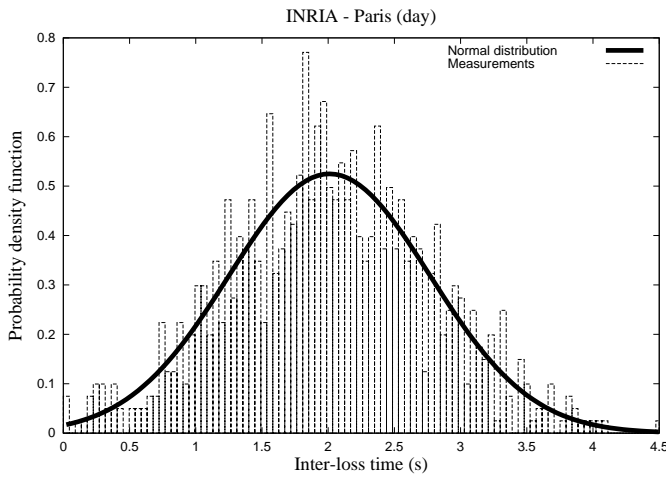


Fig. 11. MAN: Inter-loss time distribution

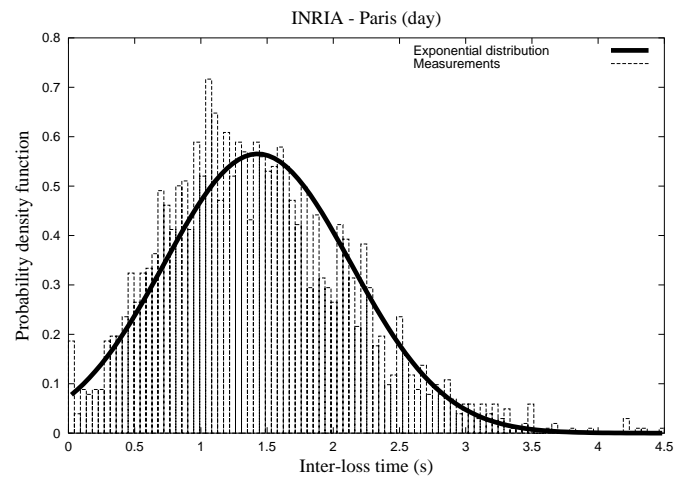


Fig. 12. MAN: Inter-loss time distribution




Fig. 13. WAN: Inter-loss time distribution

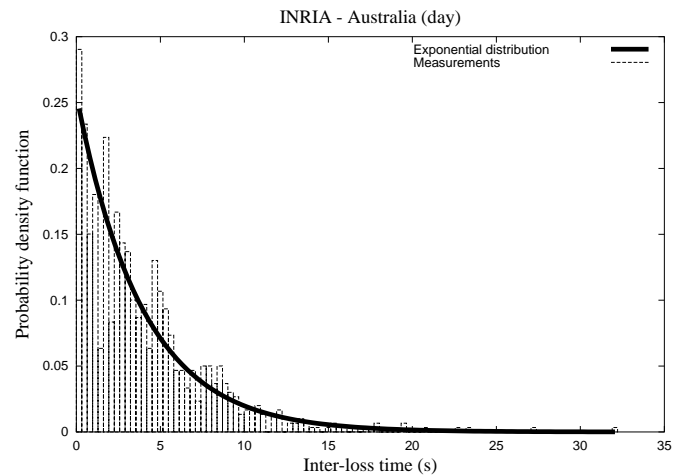


Fig. 14. WAN: Inter-loss time distribution

Hour (Traces of 20 min)	Covariance coefficient $Cov(S_n, S_{n+1})/Var(S_n)$
11:00	+ 0.034
12:00	+ 0.041
12:30	+ 0.113
13:00	+ 0.001
13:30	- 0.191
14:00	- 0.078

TABLE I

LAN: COVARIANCE COEFFICIENT

Hour (Traces of 40 min)	Covariance coefficient $Cov(S_n, S_{n-1})/Var(S_n)$
15:00	+ 0.106
19:00	+ 0.101
20:00	+ 0.015
21:00	- 0.01
22:00	- 0.048
23:00	- 0.005

TABLE II

MAN: COVARIANCE COEFFICIENT

Hour (Traces of 60 min)	Covariance coefficient $Cov(S_n, S_{n-1})/Var(S_n)$
11:00	- 0.197
12:00	- 0.001
14:00	- 0.102
16:00	- 0.107
20:00	+ 0.023
22:00	- 0.09

TABLE III

WAN: COVARIANCE COEFFICIENT

loss process. For a window reduction factor equal to one half we find,

$$\text{Throughput} = \frac{1}{RTT\sqrt{bp}} \sqrt{\frac{3}{2} + \frac{1}{2}\hat{V} + \sum_{k=1}^{\infty} \left(\frac{1}{2}\right)^k \hat{C}(k)}.$$

\hat{V} and $\hat{C}(k)$ are respectively the variance and the covariance function of the loss process normalized to the square of the average time between losses.

Three terms appear under the square root in our formula. The first term corresponds to the intensity of losses or the average inter-loss time. The second term represents the variation of inter-loss times. The third term represents the correlation of losses. This third term is equal to zero when the loss process exhibits no correlation. If we take only the first term, we get the famous square root formula [15] which was indeed established for a deterministic loss process.

This formula tells us exactly the error we introduce when making a certain assumption on the loss process. For example, it tells us that the throughput is an increasing function of the variance of inter-loss times. Assuming that losses are deterministic when they are actually Poisson should lead to an underestimation of the real throughput. In contrast, assuming that losses are Poisson when they are deterministic should lead to an overestimation of the real throughput.

The next section contains some measurement results that show how the use of a wrong variance results in a wrong throughput estimate. The formula also tells us that due to the geometrical decrease in the weights of the $\hat{C}(k)$, a small number of covariance functions is sufficient to calculate the throughput even if the loss process is highly correlated. The window evolution becomes independent of the past after a certain number of loss events.

V. SEPARATE MODEL VALIDATION

We introduce in this section the notion of separate validation of each part of a model for TCP. Researchers compare directly the real throughput achieved by a TCP connection to the final result of their modeling. But as we saw, a model for TCP is composed of two parts: a model for the window evolution and a model for losses. Proceeding for the validation in one step hides the errors introduced by these two parts. It gives us the sum of the two errors instead of each of them. First, this precludes us from distinguishing from which part of the model the error is mostly due. Second, and this is the most important, the errors introduced by the two parts of the model may be of opposite signs which may make the total error small and acceptable. The result will be a wrong estimation of the capacity of the model since, as we will see later, this phenomenon of error cancellation does not always exist.

To avoid the problem of error superposition and possibly error cancellation, we propose to validate separately the two parts of the model. We start first by the model for losses. Consider that a fluid linear model is used for window evolution. To get only the error introduced by the assumption on the distribution of losses, the window of TCP should increase continuously and linearly between losses and decrease multiplicatively by a factor 0.5 upon losses. But, TCP window does not have this ideal behavior in reality. What we do here is to construct this ideal behavior of the window using the moments of losses seen by the TCP connection. The average of round-trip time measurements is used for the calculation of the linear window increase rate. The ideal window is shown by the straight line in Figure 4. We call the version of TCP that has the window behaving as the ideal window “ideal TCP” or “exact fluid model”. Then, we calculate numerically the throughput obtained by ideal TCP and we compare it to the result of our modeling under a certain assumption on the loss process. The comparison gives us the error introduced by the model for losses.

The throughput of ideal TCP is calculated as follows. First, we sum over all the areas between the ideal window and time, then we divide this sum by the total time of the measurement. This gives the time average of the ideal window. The ideal throughput is obtained by dividing the average of the ideal window by the measured average round-trip time. Now, to get the error introduced by the model for the window evolution, all what we need to do is to compare the ideal throughput to the real throughput. Before this comparison, the ideal throughput has to be corrected for timeouts and for the packet nature of TCP using the heuristics we presented in Sections III-B and III-D.

We present some results to confirm the utility of such method for validation. We take first the LAN connection. We plot in Figure 15 the ideal throughput we obtained during the different hours of the day. We also plot in the same figure: the real throughput, the result of a linear fluid model assuming deterministic inter-loss times, and the result of the packet model in [21] which also considers that losses are deterministic. A model assuming deterministic losses underestimates the ideal throughput given the high variance of inter-loss times we observed on LAN (Figures 9 and 10). The ideal throughput in turn overestimates the real throughput due to the sub-linear growth of the window we discussed in Section III-A. We notice that a direct comparison of the real throughput to the result of the modeling in case of deterministic losses hides the two errors and gives us an impression that the model works correctly.

We take now the WAN connection (Figure 16). The ideal throughput approximates well the real throughput given that the window increase on WAN is quite linear (Section III-A). However, a model with deterministic losses does not give good performance in this case as on the LAN connection since the loss process is close to Poisson (Figures 13 and 14). Assuming that losses form a Poisson process should give better performance. This difference in the performance of a model with deterministic losses between LAN and WAN cannot be explained without the method for validation we introduced. Due to our method, we conclude that it is better to take the loss process as Poisson on the WAN connection.

VI. CONCLUSIONS

We presented in this paper an overview of the different issues to be considered when modeling TCP. Our main results can be summarized as follows:

- A linear window increase model does not hold on paths where the round-trip time is dependent of the window size. A sub-linear window increase model needs to be considered in this case. The modeling of the sub-linearity requires a model for the round-trip time which can be inferred from end-to-end measurements.

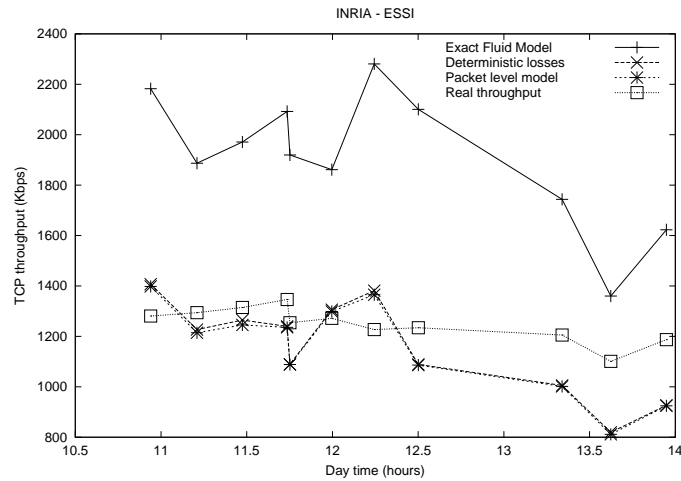


Fig. 15. LAN: Separate model validation




Fig. 16. WAN: Separate model validation

- The process of congestion events needs to be well characterized. A simplistic assumption may lead to a considerable error on some Internet paths. In particular, underestimating the variance of times between congestion events leads to an underestimation of the throughput.
- The validation of a model for TCP needs to be done in two steps. The model for window evolution and the model for the network need to be validated separately. A one-step validation hides the error introduced by each part of the model which may make the results unexplainable in some situations.

VII. ACKNOWLEDGMENTS

I would like to thank Eitan Altman and Kostya Avrachenkov for their encouragement and their valuable remarks on earlier version of this paper. I would also like to thank my colleagues at ESSI, ENST, and the University of South Australia who provided me with the required material to conduct the experimentations.

REFERENCES

- [1] M. Allman, H. Balakrishnan, and S. Floyd, "Enhancing TCP's Loss Recovery Using Early Duplicate Acknowledgment Response", *Internet Draft*, work in progress, Jun 2000.
- [2] M. Allman and A. Falk, "On the Effective Evaluation of TCP", *ACM Computer Communication Review*, vol. 29, no. 5, Oct. 1999.
- [3] E. Altman, K. Avrachenkov, and C. Barakat, "A stochastic model for TCP/IP with stationary random losses", *ACM SIGCOMM*, Sep 2000.
- [4] E. Altman, K. Avrachenkov, C. Barakat, and R. N. Queija, "State-dependent M/G/1 Type Queueing Analysis for Congestion Control in Data Networks", to appear in *IEEE INFOCOM*, Apr. 2001.

- [5] T. Bu and D. Towsley, "Fixed Point Approximation for TCP behavior in an AQM Network", *UMass CMPSCI Technical Report*, no. 00-43, Jul. 2000.
- [6] H. Chaskar, T. V. Lakshman, and U. Madhow, "On the design of interfaces for TCP/IP over wireless", *IEEE MILCOM*, 1996.
- [7] K. Fall and S. Floyd, "Simulation-based Comparisons of Tahoe, Reno, and SACK TCP", *ACM Computer Communication Review*, Jul 1996.
- [8] V. Firoiu and M. Borden, "Queue Management for Congestion Control", *IEEE INFOCOM*, Mar 2000.
- [9] S. Floyd, "Connections with Multiple Congested Gateways in Packet-Switched Networks Part I: One-way Traffic", *ACM Computer Communication Review*, Oct 1991.
- [10] S. Floyd and K. Fall, "Promoting the Use of End-To-End Congestion Control in the Internet", *IEEE/ACM Transactions in Networking*, vol. 7, no. 4, pp. 458-472, Aug. 1999.
- [11] S. Floyd, M. Handley and J. Padhye, "Equation-based congestion control for unicast applications", *ACM SIGCOMM*, Aug 2000.
- [12] S. Floyd and V. Jacobson, "Random Early Detection gateways for Congestion Avoidance", *IEEE/ACM Transactions on Networking*, Aug 1993.
- [13] V. Jacobson, "Congestion avoidance and control", *ACM SIGCOMM*, Aug 1988.
- [14] T.V. Lakshman and U. Madhow, "The performance of TCP/IP for networks with high bandwidth-delay products and random loss", *IEEE/ACM Transactions on Networking*, vol. 5, no. 3, pp. 336-350, Jun. 1997.
- [15] M. Mathis, J. Semke, J. Mahdavi, and T. Ott, "The Macroscopic Behavior of the TCP Congestion Avoidance Algorithm", *Computer Communication Review*, Jul 1997.
- [16] V. Misra, W.-B. Gong, and D. Towsley, "Stochastic differential equation modeling and analysis of TCP-window size behaviour", *Performance*, Oct 1999.
- [17] V. Misra, W.-B. Gong, and D. Towsley, "Fluid-based Analysis of a Network of AQM Routers Supporting TCP Flows with Application to RED", *ACM SIGCOMM*, Aug. 2000.
- [18] A. Misra and T. Ott, "The window distribution of idealized TCP congestion avoidance with variable packet loss", *IEEE INFOCOM*, Mar 1999.
- [19] R. Morris, "Scalable TCP Congestion Control", *IEEE INFOCOM*, Mar 2000.
- [20] J. Nagle, "Congestion control in IP/TCP internetworks", *RFC 896*, Jan 1984.
- [21] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose, "Modeling TCP Throughput: a Simple Model and its Empirical Validation", *ACM SIGCOMM*, Sep 1998.
- [22] S. Savari and E. Telatar, "The Behavior of Certain Stochastic Processes Arising in Window Protocols", *IEEE GLOBECOM*, Dec 1999.
- [23] J. Semke, J. Mahdavi, and M. Mathis, "Automatic TCP Buffer Tuning", *ACM SIGCOMM*, Sep. 1998.