MODELING THE IMPACT OF DELAY SPIKES ON TCP PERFORMANCE ON A LOW BANDWIDTH GPRS WIRELESS LINK

Pirkko Kuusela and Pasi Lassila
Networking Laboratory, Helsinki University of Technology,
P.O. Box 3000, FIN 02015 HUT, Finland
Email: {Pirkko.Kuusela,Pasi.Lassila}@hut.fi

ABSTRACT

We model the goodput of a single TCP source on a low bandwidth lossless GPRS link experiencing sudden increases in RTT, i.e., delay spikes. Such spikes trigger spurious timeouts that reduce the TCP goodput. Renewal reward theory is used to derive a straightforward expression for TCP goodput that depends on the bandwidth limitation, RTT and delay spike properties through average spike duration and distribution of the spike intervals. The basic model is for i.i.d. spike intervals, and correlated spike intervals are modelled by using a modulating background Markov chain. Also a simple deterministic p-formula is given. Validation by ns2 simulations shows excellent agreement and good accuracy even when modeling assumptions are mildly violated,e.g., regarding our lossless assumption.

1 Introduction

TCP is the most widely used transport protocol for reliable data transmission over the Internet. In addition to offering a reliable packet delivery service, TCP also includes functionality for controlling the packet sending rate to avoid congesting the network. However, this functionality has been originally designed based on the characteristics of the fixed network. There the basic operational principle of TCP rate control is roughly that TCP increases its rate as long it is receiving acknowledgements correctly from the receiver, and as soon as a packet loss is detected, it is interpreted to indicate that there is congestion in the network and hence the sending rate should be reduced considerably. As a result of a loss the sending rate is either halved or it can even be reduced back to an initial small value, as happens in the case when the loss is detected via a coarse grained timeout timer. To set the value of the timeout timer TCP relies on an adaptive estimation algorithm of the RTT (Round Trip Time). The underlying assumption in the estimation is that changes in the RTT due to random fluctuations of the traffic in the Internet can be tracked but sudden unexpected considerable increases are interpreted as a sign of congestion and hence the timer will rightly expire, triggering TCP's Go-Back-N retransmission scheme and the sending rate is initialized according to the configured initial window size.

When considering the operation of TCP over a wireless link, the delays in the observed RTTs of TCP can be highly variable. Moreover, the pattern of variability can be such that the measured RTTs contain very sharp spikes which can be even an order of magnitude larger than the typical measured RTT [6]. The RTT estimation algorithm of TCP can not track such sudden considerable increases in the measured RTTs. These sudden increases in RTTs are called *delay spikes* and potential reasons for their occurrence can be for example [8]: handovers typically result in delay spikes of several seconds occurring at a time scale of minutes in urban environment; link layer error recovery (reliable link layer protocols are usually used in modern cellular systems) may also cause delay spikes, especially when the radio channel conditions change abruptly due to the mobile's movement, e.g., when entering a tunnel; scheduling of radio resources between circuit switched calls and data (as in GPRS) can cause delay spikes. These delay spikes trigger so called *spurious timeouts* and result in unnecessary retransmissions and congestion control actions on the part of TCP, as the packets are not lost, they are simply delayed. To enable TCP to handle RTT spikes, some experimental algorithms have been proposed, namely the Eifel algorithm [8],[10], or F-RTO algorithm [13], but even they can not completely remove the effect of delay spikes, and hence evaluating their impact is necessary.

In this paper, we provide a simple modeling framework to study the impact of such spurious timeout events on the long run (steady state) *goodput* of TCP, i.e., on the amount of successfully sent traffic per time unit. Additionally, we consider a low bandwidth link (typical bandwidth of a GPRS link is 40 kbps [8]) such that a realization of the evolution of TCP sending rate is determined by the following: An RTT spike creates a period of silence during which there is a (spurious) timeout. Then the source starts increasing its sending rate exponentially (slow start) until it reaches a window size corresponding to the maximum sending rate of the physical link given the RTT. The source keeps sending at the maximum rate until there is an RTT spike again and the process repeats. In modeling, the first observation is that these spurious timeouts are caused by random events occurring in the mobile's environment, which are also independent of TCP's packet sending characteristics. Therefore, we model these events as an outside disturbance with various levels of generality regarding the inter-occurrence time of the RTT spikes which trigger spurious timeouts. First the case of i.i.d. times between RTT spikes is considered in the framework of renewal reward processes. Then a generalization by using Markov renewal reward theory is given, where the distributions between RTT spikes can be modulated by a background discrete time Markov chain (thus producing correlated times between the RTT spikes). Additionally, a simple explicit formula for the goodput is derived which assumes that there is a per packet probability that an RTT spike occurs, i.e., no explicit assumptions are made on the distributions between RTT spikes. The models are validated through extensive ns2 simulations. Also, the impact of different distributions on the performance is investigated.

The paper is organized as follows. The derivation of the models is in Section 2. Model validation and other numerical results are given in Section 3, and Section 4 contains our conclusions and suggestions for future research.

1.1 Related work

A lot of research has been done in modeling the steady state throughput of TCP in fixed networks with a given packet loss process. The simple 'square-root-p' formula was derived, e.g., in [11]. Using more complex assumptions on the nature of the packet loss process more refined formulas have also been derived, see e.g., [12] and [4]. Notably, in [4] the loss process can be an arbitrary point process and the throughput of the congestion control phase can be shown to depend on the correlations such that any positive correlations between loss events actually improves the throughput (deterministic loss process being the worst case).

In TCP models for wireless channels, the per packet loss process is typically presented by a two-state Markov process, representing the wireless channel's alteration between good/bad states, resulting in correlated packet losses, see, e.g., [2],[9], [15]. However, these models do not include spurious timeouts (RTT delay spikes).

Work related to modeling spurious timeouts has been done in [5], where a model for TCP Reno experiencing spurious timeouts has been given which very closely follows the operation of the actual TCP Reno protocol. The model is based on an extension of the well known model in [12] and considers a large bandwidth delay product, with no sending rate limitations (we include the sending rate limit). The model takes into account packet losses, but the effect of delay spikes are modelled only through the mean length and mean time between the spikes, whereas our model allows to distinguish between different distributions and non-i.i.d. processes. In [3] the measured RTT process has been modelled with a semi-Markov process and from the dynamics of the model the behavior of TCP has been deduced, including the occurrence of spurious timeouts and spurious duplicate ACKs (packets can also be reordered due to delay spikes caused by handoffs). The model does not consider the impact of congestion losses.

In comparison to [5] and [3], using a more abstract representation of how the TCP's goodput accumulates (and not attempting to model the protocol itself), we obtain a model that is straightforward to utilize and which allows quite general RTT spike processes.

2 Stochastic TCP model

We model the TCP goodput of one persistent source on a low bandwidth lossless wireless link. Typically, TCP models are discrete and consider the behavior per packet or per RTT round. Our abstraction is a continuous-time model describing the actual TCP sending rate that produces goodput.

In this scenario, the goodput of TCP is determined by an external stochastic process that generates the RTT spikes that lead to TCP timeouts. The intervals between the spikes are random and the duration of the spike is modelled using its mean duration, denoted by SD. The evolution of TCP's (successful) sending rate is described by a) a silence for time SD due to the RTT spike, b) an exponential increase in the sending rate (the slow start phase), and c) reaching an upper bound on the sending rate due to the low bandwidth of the link. Note that the congestion window can actually increase to a quite large value during the time the source is limited by the physical link rate. There can even be 'congestion losses' and corresponding window halvings, but these are ignored in our modeling since we assume that it is the physical link rate that limits the packet sending rate.

Remark 1: During the RTT spike the TCP protocol is likely to perform an exponential back-off. If the RTT spike duration is short (less than 10 seconds), it is convenient to model the time when no packets get through as the duration of the RTT spike. In this case the times when TCP attempts to send packets is close to the time when the link becomes capable of delivering packets again. If the RTT spike duration is long (15 sec or more) the TCP exponential back-off determines when sending becomes possible.

2.1 Renewal Model

To calculate the long term average goodput of the connection we put our model into the framework of renewal models [14]. Let Z_n denote the nth RTT spike, and take $Z_0 = 0$. These are the renewal times. The length of nth renewal period is denoted by $T_n = Z_n - Z_{n-1}$. The evolution of the time derivative of the TCP's actual sending rate during nth renewal period, X_n, is modelled according to the phases described earlier

$$\dot{X}_n(t) = \begin{cases} 0, & t < SD, \\ \frac{\ln 2}{R} X_n(t)dt, & SD \le t \le L + SD, \\ 0, & t > L + SD, \end{cases} \quad (1)$$

where

$$\begin{aligned} X_n &= \text{TCP sending rate, with } X_n(SD) = 1, \\ R &= \text{TCP's RTT}, \\ SD &= \text{the RTT spike duration (no goodput increase)}, \\ L &= \text{time to reach maximum sending rate } C_{\max}. \end{aligned}$$

For simplicity we have taken $X_n(SD) = 1$. Also note that in (1) the term $\frac{\ln 2}{R}$ implies that the rate doubles for every RTT (slow start). The notation introduced above is illustrated in Figure 1.

The solution of (1) is

$$X_n(t) = \begin{cases} 0, & t < SD, \\ 2^{\frac{t-SD}{R}}, & SD \le t \le L + SD, \\ C_{\max}, & t > L + SD, \end{cases} \quad (2)$$

where $L = \frac{RTT}{\ln 2} \ln C_{\max}$.

Consider the number of packets sent during the nth renewal cycle, denoted by A_n, as the reward. Then

$$A_n = \int_0^{T_n} X_n(t)dt$$
$$= \begin{cases} 0, & \text{if } T_n < SD, \\ \frac{R}{\ln 2}(2^{\frac{T_n - SD}{R}} - 1), & \text{if } SD \le T_n \le L + SD, \\ \frac{R}{\ln 2}(2^{\frac{L}{R}} - 1) + C_{\max}(T_n - L - SD), & \text{if } T_n > L + SD. \end{cases} \quad (3)$$

Next we briefly summarize the renewal reward theorem and state necessary assumptions on the lengths of the renewal cycles T_n, i.e., the time intervals between RTT spikes. We assume that the T_ns are i.i.d. random variables with finite mean $E(T) = \mu$ and probability density function f_T. Let the counting process

$$N_t = \sum_{n=1}^{\infty} \mathbf{1}_{\{Z_n \le t\}}$$

be the renewal process associated with sequence T_n. The total reward up to time t is given by

$$A_t = \sum_{i=1}^{N_t} A_i.$$

The renewal reward theorem states that the time-averaged total reward converges to cycle averaged rewards, i.e.,

$$\frac{1}{t} E(A_t) \to \frac{E(A)}{\mu}, \; t \to \infty,$$

where $E(A)$ is the common expected value of the cycle rewards. Hence, from now on, we omit the cycle index n from the notation.

Thus, to get the long term goodput of the TCP source it suffices to integrate (3) over the distribution of the delay spikes, f_T,

$$\begin{aligned} E(A) &= \int_0^T X(t)dt \\ &= \frac{R}{\ln 2} \left(\int_{SD}^{L+SD} (2^{\frac{y-SD}{R}} - 1) f_T(y)dy \right. \\ &\left. + (C_{\max} - 1) \int_{L+SD}^{\infty} f_T(y)dy \right) \\ &+ C_{\max} \int_{L+SD}^{\infty} (y - L - SD) f_T(y)dy. \end{aligned}$$

and then the TCP goodput is given by $G_{\text{TCP}} = E(A)/E(T)$. The parameters in the goodput formula are the RTT, the bandwidth limitation $C_{\max}$, the RTT spike duration SD and the distribution of the spikes.

The expression for the expected reward simplifies considerably if we assume that the intervals between RTT spikes follow the exponential distribution with parameter $1/E(T)$. In that case the goodput becomes

$$G_{\text{TCP}} = \frac{E(A)}{E(T)} = \frac{e^{-SD/E(T)} \left(\frac{R}{E(T)} - \ln 2 C_{\max}^{1-R/(E(T)\ln 2)} \right)}{R/E(T) - \ln 2}.$$

Observe that here the goodput actually depends on the time ratios SD$/E(T)$ and $R/E(T)$, in addition to the bandwidth limitation $C_{\max}$.

2.2 Markov renewal model

In the previous model the intervals between RTT spikes were required to be i.i.d. Next we show how we can introduce correlation between T_ns by considering the time-averages of semi-regenerative processes. In such processes the cycles are no longer i.i.d, but periods are "modulated" by an embedded Markov chain B_n, whose invariant distribution is denoted by π. We assume that the time scale of the modulating process is slower than the RTT spike process.

To illustrate the model in a simple setting we assume that the modulating discrete time Markov process has two states 0 and 1, and the possible state transition times are only the times when an RTT spike occurs. The state transition times of the embedded process are complete regeneration times of the joint process (see Figure 2).

In this case the sequence $(B_n, (A_n, T_n))$ is a Markov marked chain and (B_n, A_n, T_n) is called a Markov renewal sequence. Then if $E_\pi(A) < \infty$ and $E_\pi(T) < \infty$ the Markov Renewal Reward Theorem (given in [1], see also [14]) states that

$$\lim_{t \to \infty} \frac{1}{t} A_t = \frac{E_\pi(A)}{E_\pi(T).}$$

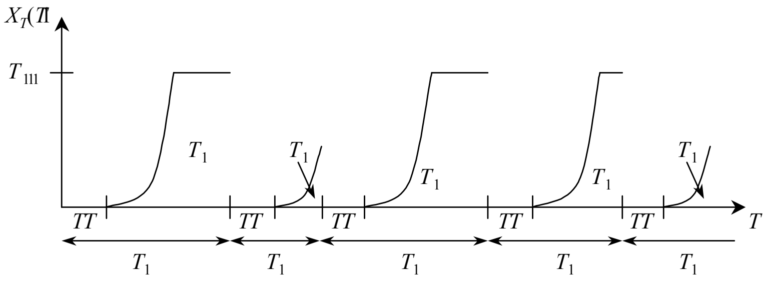

Figure 1: Notation for the renewal model.

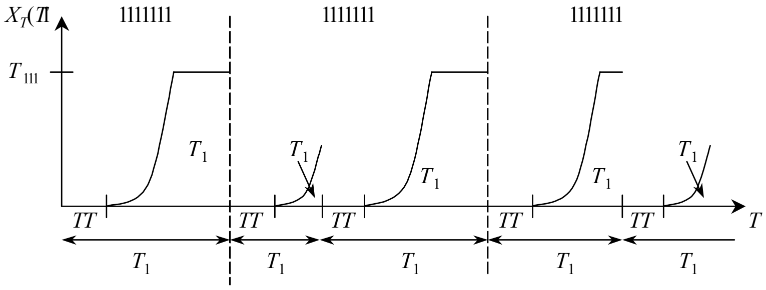

Figure 2: Notation for the Markov renewal model.

If the RTT spike intervals have a probability density function $f_0(t)$ at state 0 and $f_1(t)$ at state 1, the goodput formula becomes

$$G_{TCP} = \frac{E_\pi\left(\int_0^T X(u)du\right)}{E_\pi(T)} = \frac{\pi_0 \int_0^\infty f_0(t) \int_0^t X(u)du\ dt}{\pi_0 \int_0^\infty f_0(t)tdt + \pi_1 \int_0^\infty f_1(t)tdt} + \frac{\pi_1 \int_0^\infty f_1(t) \int_0^t X(u)du\ dt}{\pi_0 \int_0^\infty f_0(t)tdt + \pi_1 \int_0^\infty f_1(t)tdt},$$

where π_0 and π_1 are the equilibrium probabilities of states 0 and 1, respectively, and those are determined by the state transition probabilities p_{00}, p_{01}, p_{10} and p_{11}.

If we assume exponential probability densities $f_0(t) = 1/\mu_0 e^{-t/\mu_0}$ and $f_1(t) = 1/\mu_1 e^{-t/\mu_1}$, the goodput formula simplifies to

$$G_{TCP} = \frac{\pi_0 e^{-SD/\mu_0} \mu_0 \frac{R - \mu_0 \ln 2 C_{\max}^{1-\frac{R}{\mu_0 \ln 2}}}{R - \mu_0 \ln 2}}{\pi_0\mu_0 + \pi_1\mu_1} + \frac{\pi_1 e^{-SD/\mu_1} \mu_1 \frac{R - \mu_1 \ln 2 C_{\max}^{1-\frac{R}{\mu_1 \ln 2}}}{R - \mu_1 \ln 2}}{\pi_0\mu_0 + \pi_1\mu_1}.$$

2.3 Simple explicit goodput formula

In this section we derive a very simple explicit goodput formula that resembles the well known 'square-root-p' formula. In our approach we assume that each packet will experience an RTT spike with probability p. Then A_n, the number of sent packets between RTT spikes, obeys the geometric distribution with parameter p and hence $E(A_n) = 1/p$ for all n. For simplicity, we approximate the TCP sending rate at time t (after time SD has elapsed) by $g(t) = \min(2^{t/R-1/2}, C_{\max})$. We denote by H the required time to send $1/p$ packets.

If $C_{\max}$ is not reached,

$$\int_0^H g(t)dt = \int_0^H 2^{t/R-1/2}dt = \frac{1}{p},$$

which gives

$$H = \frac{R\ln(\frac{\sqrt{2}\ln 2}{pR} + 1)}{\ln 2}.$$

The maximum sending rate $C_{\max}$ is reached at time

$$\hat{t} = R\left(\frac{1}{2} + \frac{\ln C_{\max}}{\ln 2}\right).$$

Up to this time $\hat{s} := R(2^{\hat{t}/R} - 1)/(\sqrt{2}\ln 2) = R/\ln 4(2C_{\max} - \sqrt{2})$ packets have been sent. If $1/p > \hat{s}$, in the remaining time $C_{\max}(H - \hat{t})$ packets will be sent. Adding the packets sent at these two phases yields

$$\hat{s} + C_{\max}(H - \hat{t}) = \frac{1}{p},$$

which gives

$$H = \frac{R(\ln 2 + 2\ln C_{\max})}{\ln 4} + \frac{1}{C_{\max}}\left(\frac{1}{p} + \frac{R(\sqrt{2} - 2C_{\max})}{\ln 4}\right)$$

Finally, after taking into account the duration of the RTT spike, SD, in the total time, the simple TCP goodput formula reads

$$T_{TCP} = \frac{1/p}{H + \mathrm{SD}} = \begin{cases} C_1, & \text{if } p > \frac{\ln 4}{R(2C_{\max} - \sqrt{2})}, \\ C_2, & \text{otherwise}, \end{cases}$$

where

$$C_1 = \ln 2 \Big(p(R \ln(1 + \sqrt{2} \ln 2/(pR)) + \mathrm{SD} \ln 2 \Big)^{-1},$$

$$C_2 = (C_{\max} \ln 4) \Big(\ln 4 + pR(\sqrt{2} + C_{\max}(\ln 2 - 2 + 2 \ln C_{\max})) + \mathrm{SD} p C_{\max} \ln 4 \Big)^{-1}.$$

Again the dependence on time parameters resembles the one in previous section. Here the goodput is determined by products pSD and pR together with $C_{\max}$.

Remark 2: Above p is considered as a given parameter. However, the probability p can be related approximately to the mean of the RTT spike intervals by using

$$p = \left(\int_0^{E(T)} X(u) du \right)^{-1},$$

where X is given by (2). In our numerical examples this is used as an estimate for p.

3 Numerical results

3.1 Model validation with ns2

To validate our models we have used the ns2 simulator equipped with the RTT spike generator contributed by Andrei Gurtov[1]. The simulation scenario that is used to estimate the goodput consists of transmitting TCP data over a single link having a constant service rate, but which can be turned on/off for random periods of time (to generate the RTT spikes). The TCP source type used in the simulations is TCP SACK. To estimate the goodput, we simply send a large file ($3 \cdot 10^6$ packets) over the link and measure the time it took for TCP to send the given number of packets. Although the model is developed for a low bandwidth link, the accuracy of our model is evaluated for a small and larger bandwidth delay product scenario. Additionally, we vary the length of the RTT spikes (long/short spikes) and also study the impact of congestion losses. The fixed parameters in all the simulations are: packet size is 576 bytes and link one-way propagation delay is 100 ms. In the following, we use the name 'P-model' to refer to the simple goodput formula of Section 2.3 and 'Renewal model' refers to the renewal model of Section 2.1. In the models, we also take here R to include the packet transmission time (in addition to the propagation delay), i.e., $R = 2 \cdot \text{link propagation delay} + \text{packet transmission time}$.

First the case of long RTT spikes is considered; in the simulations the spike durations were uniformly distributed in the range $[5, 10]$ s, i.e., $SD = 7.5$ s in our models. The goodput is evaluated as a function of the mean time between RTT spikes, $E(T)$. In the simulations for each $E(T)$, the spikes are uniformly distributed in the range $[10, 2 \cdot E(T) - 10]$ s. In the P-model the value for p is obtained as in Remark 2. As seen from the results in Figure 3, for the small bandwidth delay product case ($C_{\max} = 40$ kbps, left figure) corresponding to our assumed scenario of a small bandwidth GPRS link, the agreement between the models (both P-model and Renewal model) and with ns2 simulations where no congestion losses occur is excellent. Our model does not explicitly include the effect of congestion losses, but to evaluate the impact of congestion losses on the accuracy, ns2 simulations with a loss module generating packet losses according to a given probability have been performed. The results are shown in the figures with dashed lines. In the low bandwidth case, we can observe that the models naturally overestimate the goodput and that accuracy is still quite good up to moderate loss rates (2% loss), but for high loss rates the effect becomes more clearer. For the higher bandwidth delay product case ($C_{\max} = 100$ kbps, right figure) we can still see good agreement, although our models seem to overestimate the actual goodput somewhat. We assume that the overestimation in the model is due to our assumed exponential increase up to $C_{\max}$, while, in the simulations, TCP changes the window increase into a linear one after reaching the slow start threshold. Also, in this context the impact of congestion losses becomes more pronounced (as is to be expected). Note that the P-model and the Renewal model yield results which are very close to each other (almost indistinguishable in the figures). This is to be expected as the spike interval and length distributions are uniform distributions. The impact of different distributions is evaluated later in this section.

In the above simulations the spike process produces rather long spikes (with a mean of 7.5 s), where as the typical RTT equals roughly 300 ms. Then it can be expected that the spikes are indeed dramatic enough such that TCP's RTT estimator mechanism can not learn the properties of the spike process. Hence, our assumption that every spike results in a timeout and slow start is mostly valid. However, if the duration of the spikes is shortened, this may not necessarily be the case anymore. To check this, same simulations as earlier (with low/high bandwidth) were performed where the spike durations were generated from a uniform distribution in the range $[1, 5]$, i.e., with a mean $SD = 3$ s. The results are shown in Figure 4 and it can be seen that, indeed, if the spike durations are shorter compared to the typical RTT,

[1]Patches are available from `http://www.cs.helsinki.fi/u/gurtov/ns/`

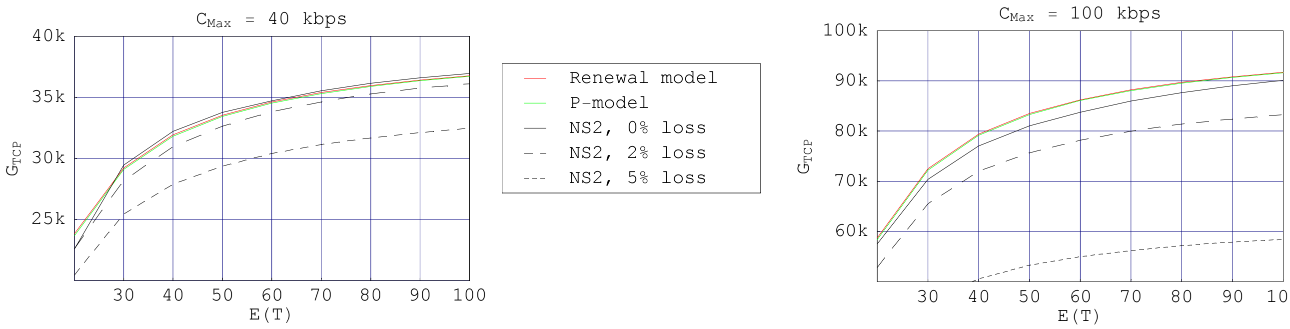

Figure 3: Goodput of TCP in low (left) and higher (right) bandwidth delay product scenarios when delay spikes have relatively long durations.

TCP's retransmission timer adapts to the spike process. Hence, there are less timeouts and less slow starts (slow starts cause loss of goodput). Our model assumes a slow start after each spike duration and the end result is that in circumstances, where TCP's retransmission timer learns the properties of the spike process, our model underestimates the goodput. However, as seen from the numerical results, the influence of this is not substantial. The pictures also show how congestion losses affect the accuracy, and the impact can be seen to be (more or less) as earlier.

3.2 Distribution Sensitivity

We have verified the accuracy of the renewal model and the simple P-model against ns2 simulation when the RTT spike distribution was uniform. Now we illustrate in Figure 5 the effect of the RTT spike distribution on the TCP goodput. The x-axis is the mean of the given distribution and the y-axis is the corresponding estimate for the TCP goodput in bps. In the Pareto distribution the shape parameter equals 1.5; the uniform distribution is located at the interval $[E(T) - \min(E(T), 5), E(T) + \min(E(T), 5)]$. Markov modulating process is illustrated by using exponential distributions with $\pi_0 = \pi_1 = 1/2$ and $\mu_0 = 0.25E(T)$ and $\mu_1 = 1.75E(T)$, i.e., in state 0 RTT spikes are more frequent than in state 1. In the P-model the value for p is obtained as in Remark 2. The RTT spike duration is SD $= 5$ and the networking parameters are: maximum rate 40 kbps, packet size 576 bytes, propagation delay 0.6 sec, which give rise to $R = 720$ ms and $C_{\max} = 8.68$ pkts/s.

For the TCP goodput the worst case of RTT spikes is the uniform distribution, which resembles closely the simple P-model (in which the delay spikes occur at constant intervals of length $H + SD$ as discussed in Section 2.3). Although the P-model is attractive in its simplicity, we notice that it may underestimate the goodput. This is because n bursty spikes yield a better goodput than n uniformly distributed spikes on a time interval of the same length. The benefit of bursts can be seen by comparing the Pareto and exponential distributions. Moreover, the exponential distribution and the Markov modulated exponential distribution have the same mean, but in the Markov modulated version spikes arrive frequently in state 0 and very seldom in state 1. Thus the modulated RTT process yields throughout a higher goodput and extending the renewal model to include some correlation in the spike intervals is important. The benefit of correlation in RTT spikes on the goodput is in agreement with the result on the TCP AIMD-component analyzed in [4].

4 Conclusions

On a wireless link, the observed RTTs of TCP can be highly variable and the pattern of variability may contain sharp spikes, called delay spikes. These result in spurious timeouts that lower the TCP performance. We have provided a facile modeling framework to study the impact of such RTT spikes on the goodput of a TCP source. In the modeling, we have considered a lossless low bandwidth link on which the rate of successfully sent TCP packets is described by a) a silence for the mean duration of the delay spike, b) exponential increase (the slow start phase) and c) reaching the maximum rate due to bandwidth limitation.

Spurious timeouts are triggered by random events occurring in the mobile's environment. Hence we modelled the delay spikes as an outside disturbance. First, the case of i.i.d. spike intervals was considered in the framework of renewal rewards and we derived an expression for the TCP goodput. Correlation between RTT spike intervals was incorporated using embedded discrete Markov chain modulation and Markov renewal reward results. Above models require the RTT spike interval distribution. However, a simple explicit formula for the goodput was derived based on per packet probability for an RTT spike, thus requiring no explicit knowledge on the spike distribution.

Validation with ns2 showed that our models closely approximate the TCP goodput in the presence of RTT spikes. In the exact modeling scenario considered (low bandwidth link, no packet losses) the agreement with simulation is excellent, both in the renewal model and in the simple P-model. Moreover, the modeling assumptions can be violated mildly. With moderate packet losses the agreement is still good. For a higher band-

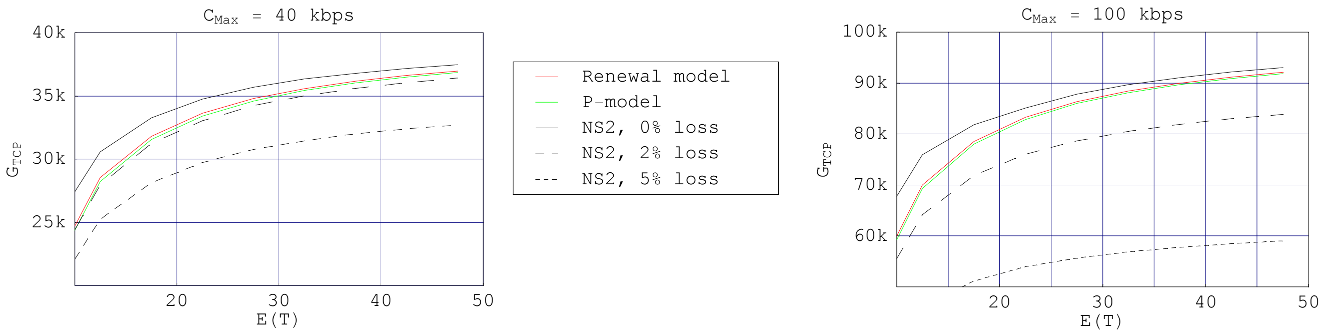

Figure 4: Goodput of TCP in low (left) and higher (right) bandwidth delay product scenarios when delay spikes have shorter durations.

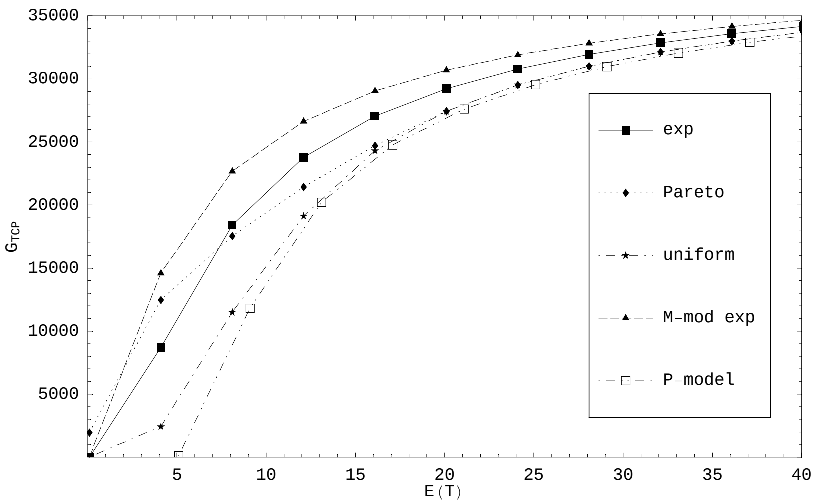

Figure 5: The impact of the RTT spike distribution on the TCP goodput: exponential, pareto, uniform, and Markov modulated exponential distributions together with the simple P-model.

width link our model slightly overestimates the lossless goodput (as we allow exponential increase up to the bandwidth limitation). By studying the sensitivity of the goodput on the RTT spike interval distribution, we observed that bursty spikes give higher goodput. This is similar to the result in [4] that studied the congestion control component of the TCP. The deterministic simple P-model is the worst case, thus giving lower bounds for the goodput.

Topics for future research include the incorporation of the distribution of the RTT spike lengths into the model. Also, the impact of congestion losses on the goodput should be included in the model.

Acknowledgement

We are grateful to Samuli Aalto (Ph.D.) for helpful discussions on semi-regenerative processes.

References

[1] S. Aalto, “Time-average properties of regenerative and semi-regenerative processes”, unpublished manuscript, 1995.

[2] A. A. Abouzeid, S. Roy, and M. Azizoglu, “Stochastic modeling of TCP over lossy links”, in Proceedings of INFOCOM 2000, Tel Aviv, Israel, March 2000.

[3] A. A. Abouzeid, and S. Roy, “Stochastic modeling of TCP in networks with abrupt delay variations”, ACM/Kluwer Wireless Networks, vol. 9, no. 5, September 2003

[4] E. Altman, K. Avrachenkov, and C. Barakat, “A stochastic model of TCP/IP with stationary random losses”, in proceedings of ACM SIGCOMM 2000, Stockholm, Sweden, August 2000.

[5] S. Fu, and M. Atiquzzaman, “Modeling TCP Reno with spurious timeouts in wireless mobile environment”, in proceedings of 12th International Confer-

ence on Computer Communications and Networks, Dallas, Texas, USA, October 2003.

[6] A. Gurtov, “Effect of delays on TCP performance”, in Proceedings of IFIP Personal Wireless Communications ’01, Lappeenranta, Finland, August 2001.

[7] A. Gurtov, and R. Ludwig, “Responding to spurious timeouts in TCP”, in Proceedings of INFOCOM 2003, San Fransisco, California, USA, April 2003.

[8] A. Gurtov, and S. Floyd, “Modeling wireless links for transport protocols”, to appear in ACM Computer Communications Review, November 2003, available at http://www.cs.helsinki.fi/u/gurtov/papers/mtp.html

[9] A. Kumar, “Comparative performance analysis of versions of TCP in a local network with a lossy link”, IEEE/ACM Transactions on Networking, vol. 6, no. 4, August 1998.

[10] R. Ludwig, and R. H. Katz, “The Eifel algorithm: making TCP robust against spurious retransmissions”, ACM Computer Communications Review, vol. 30, no. 1, January 2000.

[11] M. Mathis, J. Semke, J. Mahdavi, and T. Ott, “The macroscopic behavior of the TCP congestion avoidance algorithm”, ACM Computer Communication Review, vol. 27, no. 3, July 1997.

[12] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose, “Modeling TCP throughput: a simple model and its empirical validation” in Proceedings of ACM SIGCOMM’98, Vancouver, Canada, September 1998.

[13] P. Sarolahti, M. Kojo, and K. Raatikainen, “F-RTO: an enhanced recovery algorithm for TCP retransmission timeouts”, ACM Computer Communications Review, vol. 33, no. 2, April 2003.

[14] R. Wolff, “Stochastic Modeling and the Theory of Queues”, Prentice-Hall, 1989.

[15] M. Zorzi, A. Chockalingam, and R. R. Rao, “Throughput analysis of TCP on channels with memory”, IEEE Journal on Selected Areas in Communications, vol. 18, no. 7, July 2000.