# The Impact of Delay Variations on TCP Performance

Michael Scharf

Institute of Communication Networks and Computer Engineering⋆
University of Stuttgart, Germany
scharf@ikr.uni-stuttgart.de

**Abstract.** This paper studies to which extent variable delays in mobile networks affect TCP performance. Our main contribution is a simple analytical model that determines the TCP timeout duration from given network parameters. Based on the model, we quantify the risk of spurious timeouts and their impact on TCP goodput. Our results show that TCP is quite insensitive to variable delays.

## 1 Introduction

Delays in mobile networks can be highly variable because of effects such as handovers or resource preemption. The transmission control protocol (TCP) suffers from performance degradation if such delay variations trigger spurious TCP timeouts, i. e., if timeouts occur even though no packets are lost. In this case, TCP may waste scarce bandwidth by unnecessarily retransmitting segments or underutilize the available resources due to the needless reduction of the sending window [1]. There are two important questions in this context: First, how often do severe delay fluctuations occur in real networks? Even though measurements in GPRS [2, 3] have observed delay spikes of up to several seconds, their frequency remains an open research issue. And second, do the variations indeed trigger spurious timeouts? This mainly depends on the duration of the TCP retransmission timeout (RTO). As specified in RFC 2988, TCP determines the RTO value taking into account both the low-pass filtered round-trip time (RTT) samples and the observed delay variance. However, the dynamics of this algorithm have hardly been addressed so far, since in wireline networks the RTO duration is usually dominated by a minimum value of 1 s [4].

In this paper, we quantify the sensitivity of TCP to delay variations in order to address the second issue. In Section 2, we analyze the RTT estimation for bulk data traffic over links with rather low data rates and high latencies. Section 3 outlines an analytical model that accurately predicts the RTO duration from the properties of the path. Based on upper and lower bounds for the RTO duration given by the model, we quantify the risk of spurious timeouts for a wide parameter range in Section 4. We also estimate the performance degradation due to spurious timeouts. From this we conclude that TCP is quite insensitive to variable delays, and that optimization approaches such as the Eifel-Algorithm [5] are only beneficial in case of frequent and extreme delay variations.

## 2 Modeling the round-trip time measurement

When accessing the Internet through a cellular network, the radio channel usually is the bottleneck. If a single TCP connection utilizes this link, the *congestion window* follows a regular saw-tooth pattern because of the *additive increase, multiplicative decrease* mechanism of the TCP congestion control [6], provided that it is not restricted by the *receiver advertised window*. This window evolution results in different queuing delays at the bottleneck. We argue that, in one *cycle* [6], the round-trip time samples $x(n)$ can be approximated by $x(n) = \frac{1}{2}\,x_{\text{max}} + \frac{1}{2}\,x_{\text{max}}\,\frac{n}{N-1}$ with $n \in [0; N-1]$. The maximal round-trip time $x_{\text{max}}$ therein is the maximum path capacity $C$, i. e., the bandwidth-delay product counted in segments, divided by the service rate $\mu$ of the radio link. $N$ is the number of RTT samples. By default, TCP takes one sample per RTT. Alternatively, TCP may measure the RTT with help of timestamps. From an abstract point of view, these two methods only differ in the number of samples:

$$N_{\text{default}} \approx \left\lceil \frac{b}{2}\,(C+1) + 2\,b \right\rceil \quad \text{or} \quad N_{\text{timestamps}} \approx \left\lceil \frac{3}{8}\,(C+1)^2 + \frac{C}{b} + 1 \right\rceil . \tag{1}$$

These formulae are derived in [7]. Note that, due to the mechanism of *delayed acknowledgments*, one acknowledgment may either refer to $b = 1$ or $b = 2$ segments, depending on the service rate $\mu$.

## 3 Analysis of the TCP round-trip time estimator

The RTO computation is done by a non-linear filter. Using the RTT samples $x(n)$ as input function, the RTO duration $R(n)$ can be determined analytically, e. g. by applying the Z-transform. As shown in Fig. 1, the maximum $R_{\text{max}}$, the mean value $\overline{R}$, the minimum $R_{\text{min}}$ and the minimal difference $\Delta_{\text{min}}$ between $R(n)$ and the RTT

⋆ The research work of this paper was done in cooperation with Alcatel, Research and Innovation Department, Stuttgart, Germany.

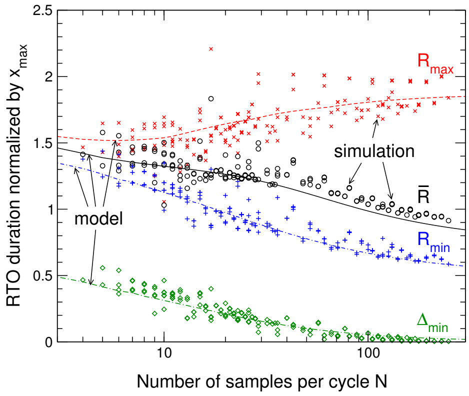


**Figure 1.** RTO duration predicted by the model

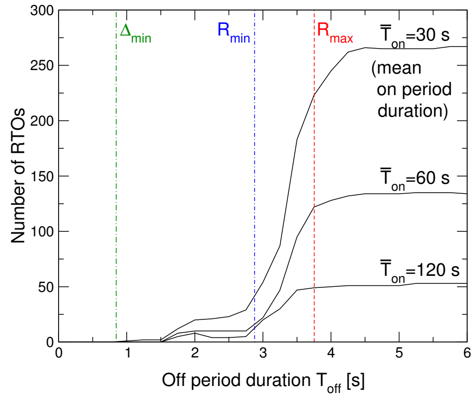


**Figure 2.** Spurious timeouts caused by delay spikes

highly depend on the number of samples $N$. The algorithm is more aggressive if the sampling rate is high, i. e., if timestamps are used. Thus, our results confirm that the parameters standardized in RFC 2988 are not well suited for high sampling rates. This effect must be carefully weighted up against the faster feedback that timestamps can provide. Fig. 1 also shows that our analytical model matches quite well to simulation results for typical GPRS and UMTS scenarios, even though there are some discrepancies. More details on the analysis can be found in [7].

## 4 Quantifying the impact of delay variations on TCP

The results of the analysis allow us to determine the probability that a timeout is triggered by an off period of duration $T_{\text{off}}$. Since the retransmission timer is usually restarted whenever a new acknowledgment arrives, $T_{\text{off}}$ must exceed the current RTO duration in order to trigger a timeout. The simulation results in Fig. 2 confirm that there virtually always is a spurious timeout if $T_{\text{off}} > R_{\text{max}}$. We therefore propose a *spurious timeout probability*:

$$P_{\text{TO}}(T_{\text{off}}) = \begin{cases} 0, & \text{if} \quad T_{\text{off}} < R_{\text{min}}, \\ (T_{\text{off}} - R_{\text{min}})/(R_{\text{max}} - R_{\text{min}}), & \text{if} \quad R_{\text{min}} \leq T_{\text{off}} < R_{\text{max}}, \\ 1, & \text{if} \quad T_{\text{off}} \geq R_{\text{max}}. \end{cases} \tag{2}$$

This formula approximates the probability that a transmission interruption of duration $T_{\text{off}}$ results in a spurious timeout, assuming that packets remain buffered in the link layer. The main input parameters are the maximum path capacity $C$ and the bottleneck service rate $\mu$. Mobile networks are typically characterized by $R_{\text{min}} \geq 1\,\text{s}$ because of the high latencies on the radio link and the lower bound recommended in RFC 2988. Thus, as a rule of thumb, delay variations are only critical if they are of the order of several seconds. Considering that $U \in [C/2; C]$ segments are unnecessarily retransmitted after a spurious timeout, the goodput of a bulk data TCP connection is

$$\lambda_{\text{G}}(T_{\text{off}}) = \lambda_{\text{max}} \cdot \frac{T_{\text{On}}}{T_{\text{On}} + T_{\text{Off}}} \cdot \left(1 - \frac{P_{\text{TO}}(T_{\text{off}}) \cdot U}{T_{\text{On}} \cdot \mu}\right), \tag{3}$$

where $T_{\text{on}}$ is the mean time between the interruptions. $\lambda_{\text{G}}(T_{\text{off}})$ is quite close to the maximum goodput $\lambda_{\text{max}}$ if $T_{\text{On}} \gg U/\mu$. Otherwise, the last factor in Eq. (3) quantifies the performance improvement of TCP enhancements like the Eifel-Algorithm. However, in terms of goodput, the performance degradation caused by sudden delays is only significant if timeouts occur quite frequently, e. g., more than once per minute in a GPRS scenario.

## References


1. H. Inamura, G. Montenegro, R. Ludwig, A. Gurtov, and F. Khafizov, “TCP over second (2.5G) and third (3G) generation wireless networks,” RFC 3481, Feb. 2003.
2. A. Gurtov, M. Passoja, O. Aalto, and M. Raitola, “Multilayer protocol tracing in a GPRS network,” in *IEEE VTC 2002 Fall*, Sept. 2002.
3. R. Chakravorty, J. Cartwright, and I. Pratt, “Practical experience with TCP over GPRS,” in *IEEE GLOBECOM*, Nov. 2002.
4. M. Allman and V. Paxson, “On estimating end-to-end network path properties,” in *ACM SIGCOMM*, Sept. 1999.
5. R. Ludwig and R. H. Katz, “The Eifel algorithm: Making TCP robust against spurious retransmissions,” *ACM SIGCOMM Computer Communications Review*, vol. 30, no. 1, pp. 30 – 36, 2000.
6. T. 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, June 1997.
7. M. Scharf, M. Necker, and B. Gloss, “The sensitivity of TCP to sudden delay variations in mobile networks,” to appear in Springer LNCS series, Proc. Networking 2004, Athens, Greece, May 2004.