

# The Effect of TCP's Adaptivity on the Throughput

Zoltán Szabó, Sándor Molnár

High Speed Networks Laboratory

Dept. of Telecommunications and Telematics

Budapest University of Technology and Economics

H-1117 Magyar tud. krt. 2, Budapest, Hungary

Email: {szabo,molnar}@ttt-atm.ttt.bme.hu

Zsolt Kenesi

Traffic Analysis and Network Performance Lab.

Ericsson Hungary Ltd.

H-1037 Laborc u. 1, Budapest, Hungary

Email: zsolt.kenesi@eth.ericsson.se

**Abstract**—This paper analyses how the throughput of TCP connections is affected by the adaptation of the protocol to fluctuations of the background traffic. Due to its congestion control TCP adapts to the changes of the background traffic on several time-scales. However, TCP flows suffer significant throughput reduction if the background traffic fluctuates around a certain time-scale that depends on network conditions. In this paper the phenomenon and the mechanisms of the throughput reduction are demonstrated and analyzed. Furthermore, two different techniques are shown to eliminate the effect. Our arguments are supported with several kind of simulations.

**Index Terms**—TCP throughput, congestion control, adaptivity, performance

## I. INTRODUCTION

Data traffic based on Transmission Control Protocol (TCP) [1] is dominant in IP networks. The detailed analysis of TCP's features and behavior is a hot topic of recent research programs [9] [10]. The most important information about TCP can be found in RFC 793 [2] in which TCP originally was defined, while RFC 1122 [3] and RFC 2001 [4] contain further additions. In [5] and [6] several additional features are described.

Since TCP connections are capable of providing and interpreting feedback they adapt to different network conditions that are relevant to the actual scenario they experience. It was published in papers [7] and [8] that if a TCP connection shares a bottleneck link with non-adaptive background traffic flow, TCP adapts to it inheriting and propagating the correlation structure and statistical properties of the background traffic flow above a characteristic time-scale. This time-scale of adaptation depends on the end-to-end path properties, i.e. round-trip time, widow size, etc.

We found that TCP's adaptation to background traffic effects the throughput of the connection, as well. Especially when the background traffic fluctuates around the characteristic time-scale mentioned above, the throughput of TCP connections suffers significant reduction. We also found that this effect is almost independent of the commonly applied versions of TCP. We also analyzed how network parameters affect the reduction of the throughput.

Finally, we show two methods for compensating the effect of throughput reduction either by applying larger buffers to avoid consecutive losses or decreasing the maximum Retransmission Time-Out (RTO) value to reach shorter recovery periods of the TCP window.

The paper is organized as follows. Results of previous works about the adaptation of TCP is discussed in Section 2. Section 3 investigates how TCP's throughput is affected by fluctuations on different time-scales. In Section 4 we analyze the phenomenon and show the effect of network parameters in Section 5. Finally, in Section 6, we present how throughput reduction can be compensated by limiting maximum RTO or applying larger buffers.

## II. ADAPTIVITY OF TCP

There is a feedback mechanism implemented in TCP which is responsible for congestion control of the flow. It adjusts the rate of the source according to changing network conditions. The rate of the source is decreased in case of packet loss.

Adaptivity can be observed clearly if a TCP flow shares the bandwidth with other traffic in a bottleneck. In case the background traffic is not adaptive, e.g. based on User Datagram Protocol (UDP), TCP utilizes the free capacity on the link according to a special characteristics described and discussed in details in [7] and [8].

It was also shown in the cited papers that TCP inherits the statistical properties of the background above a characteristic time-scale that depends on the basic period of the cwnd determined by network parameters, such as buffer size, bandwidth, delay, etc. Even below this time-scale TCP adapts to fluctuations, though the effectiveness of adaptation is quite limited. Additionally, it was also presented that TCP propagates self-similarity in the network according to the adaptation characteristics.

## III. PHENOMENON OF THROUGHPUT REDUCTION

In most of the networks adaptive and non-adaptive traffic is transmitted through the same infrastructure. In networks where diffserv architecture [11] is not implemented TCP shares the path with flows based on different protocols like UDP. Our motivation is to understand how adaptive and non-adaptive traffic interacts from the perspective of TCP's adaptation and its effect on the throughput of the protocol.

Analysis of the phenomenon is not trivial. It is a difficult task especially for real traffic which has complex structure including short and long connections limited by several bottlenecks in different network parts. To present the problem and understand the mechanism behind we apply a simplified model that is in line with the previously cited papers.

As a first step of our analysis we investigate TCP's throughput when a single, long and greedy TCP flow shares bottleneck with periodically changing non-adaptive background traffic. The background streams are based on UDP and are constructed in a way, such that they fluctuate on a limited, narrow time-scale. To limit the time-scale under investigation, the background traffic approximates a constant amplitude sine wave of a given frequency  $f$ . The power spectrum of this process consists of a single frequency component at  $f$ . In the simulation the background process was approximated by a packet stream with a packet size of 1000 bytes. In this way, we perform a kind of harmonic analysis and investigate the effect of only one component in the spectrum of background traffic stream.

The background UDP traffic rate changes between zero and the maximum bandwidth. As a result of this method, 50% of the link is free independently of the frequency of the fluctuations. By applying such kind of stream we are able to measure how the throughput of TCP flow depends on the fluctuations of background UDP traffic.

Figure 1 shows the test configuration. The buffers are identical with a size of 40 packets. Service rates are 1Mbps at each link, while propagation delays are 5ms. (We analyzed stationary state of the configuration, therefore all the simulations were performed for tenth of minutes and the transient interval was removed from the beginning of the log files.)

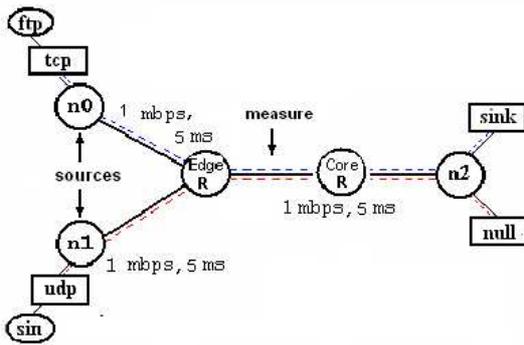


Fig. 1. Simulation model for analysis of throughput characteristics

By changing the frequency of the sine wave we are able to measure the throughput of TCP as a function of frequency components in the background streams. We refer this function as throughput characteristics of TCP in the rest of the paper. Figure 2 shows the results. We present normalized throughput where value 1 represents the average free capacity in the link which is 500kbytes in this case.

As it can be observed, throughput characteristics depends on the frequency of changes in background traffic rate. The normalized throughput varies between 90-100% in most of the cases meaning that TCP is able to utilize free link capacity quite efficiently. However, there is an interval around 0.4 1/sec where the throughput is reduced to 60-70% of free capacity. Moreover, those reduction is almost invariant to the commonly used versions of TCP, i.e. Reno and New Reno.

Since in real networks background traffic is not limited to a single time-scale, we analyze in the following scenario when more than one frequency components are present in the spec-

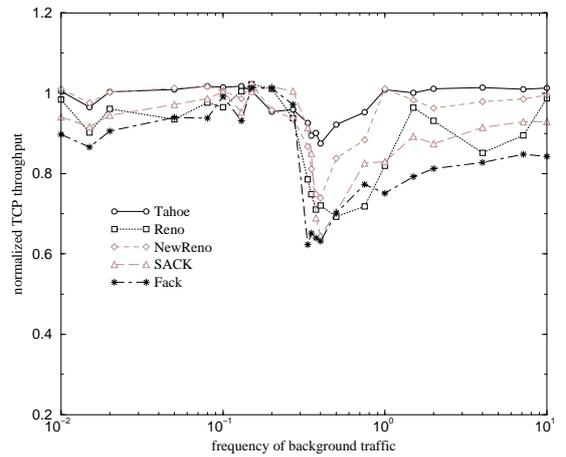


Fig. 2. Throughput characteristics of different TCP versions

trum of the background traffic stream. Our motivation for this analysis is to investigate whether TCP's throughput is determined by fluctuations of background traffic on different time-scales independently of each other or not.

To test the independence of effects a wide range of traffic mixes are generated composed of two frequencies  $f_1$  and  $f_2$ . The rate of background stream was changed based on the following equation:  $A_{background}(t) = C_1 \sin(2\pi f_1 t) + C_2 \sin(2\pi f_2 t) + C_2$  where  $C_1$  and  $C_2$  are constants.

A large number of simulations were performed, covering a whole plane with the two frequencies in the range of [0.05; 10] 1/sec, see Figure 3. The figure shows that in the area where one of the frequency components is around the critical frequency the throughput is reduced. This result indicates that TCP's throughput is determined by the two frequency components independently.

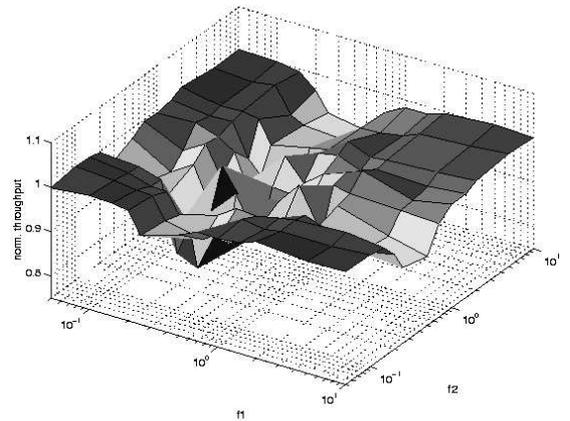


Fig. 3. Throughput characteristics of TCP when the background process is composed of two frequencies

#### IV. ANALYSIS OF THE PHENOMENON

The throughput characteristics presented in Figure 2 can be separated into three different intervals. The first is in the left side, above the critical time-scale of reduction, i.e. at low frequencies. The second interval is around the critical time-scale and the third one is below the critical time-scale, i.e. at high

frequencies. We discuss each interval separately and present results based on simulation with Reno version of TCP.

Above the critical time-scale the background traffic rate varies with relatively low frequency. It varies relatively slow, thus TCP is able to follow the changes of background and utilize the free capacity left in the bottleneck. Rates of both TCP and UDP traffic is depicted in Figure 4.

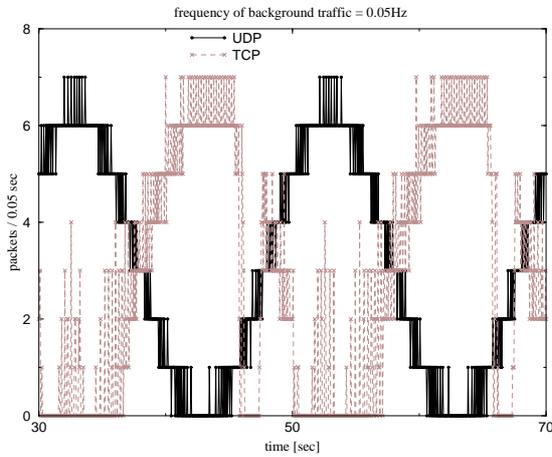


Fig. 4. TCP and background traffic at large large time-scale

Since the speed of changes of the background traffic is relatively slow, TCP is able to adapt almost perfectly to it. As the throughput of background traffic increases TCP packets are dropped. When UDP utilize the whole bandwidth TCP is unable to send packets and is forced into time-out phase. After time-out TCP sends packets again starting from a slow-start phase. The process can be followed clearly in Figure 5 that shows the congestion window of TCP.

Please note, that we refer as 'congestion window' to the number of TCP packets in the network. The internal *cwnd* variable of TCP differs from the number of packets in the network during short periods after packet loss. In those periods *cwnd* is set for artificial values and does not refer to the number of the outstanding packets.

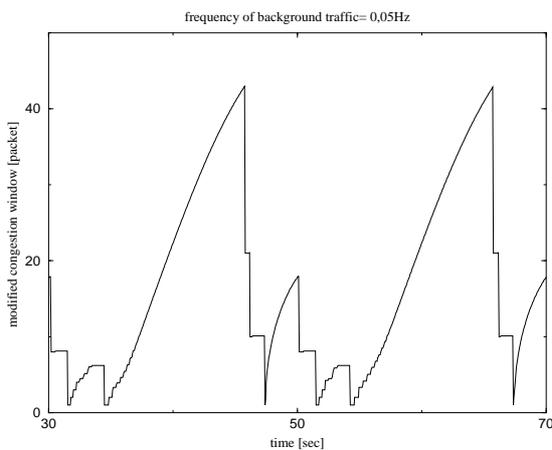


Fig. 5. TCP and background traffic at large time-scales

The second part of the characteristics is the critical interval. Similarly to the previous case, TCP losses packets and performs

timeout period when the background traffic fills the link in, see Figure 6. The figure shows the growth of sequence number of TCP packets in time.

Based on our comprehensive analysis we located the critical interval. However, the detailed discussion on the location of the interval is not the target of this paper.

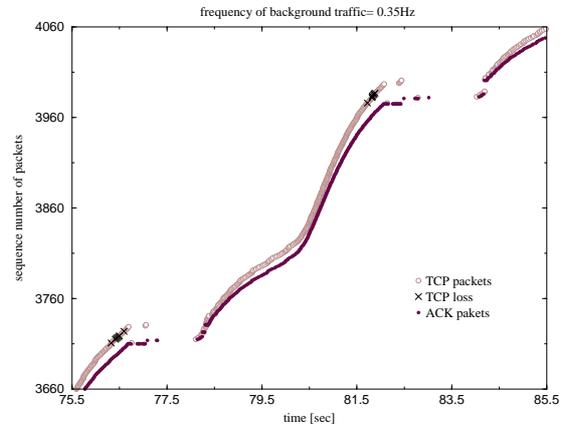


Fig. 6. TCP trace around critical time-scale

When a wave of background traffic arrives, TCP suffers consecutive packet losses and is forced into timeout phase. After a timeout period TCP starts sending packets until the next wave of background stream arrives. Since the time-scale the background fluctuates on is close to TCP's recovery period, TCP is not able to reach higher transmission rates between two waves of background traffic. This periodic process can be followed in Figure 7 that shows buffer utilization in the bottleneck.

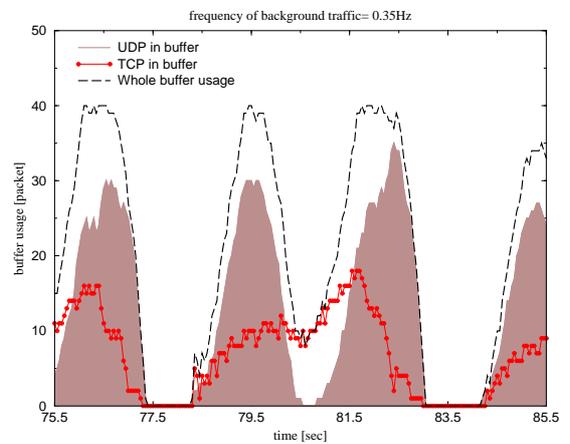


Fig. 7. Buffer utilization around critical time-scale

TCP's congestion window also shows the effect clearly. The maximum value of the window is 20 while it was 40 in the previous case, see Figure 8. It means that TCP source is unable to increase the number of packets sent by 'opening' the window even the background traffic has the same average rate.

The third part of the throughput characteristics is where background changes at high rate. Obviously, TCP is not able to adapt to more frequent background changes similarly to the case in the critical interval, but there is another effect that compensates throughput reduction below the critical time-scale.

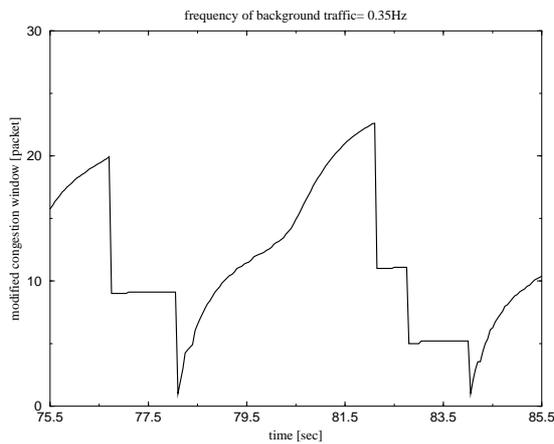


Fig. 8. TCP's congestion window around critical time-scale

The buffer utilization of background stream depends mainly on the amount of data sent during a wave. If the frequency of background traffic fluctuations is increased, the amount of data sent during one period is decreased. That leads to the effect that background traffic utilizes less buffer space compared to lower time-scales.

From the point of view of TCP it means that TCP always finds some free buffer space and avoids consecutive losses even during the waves of background stream, see Figure 9.

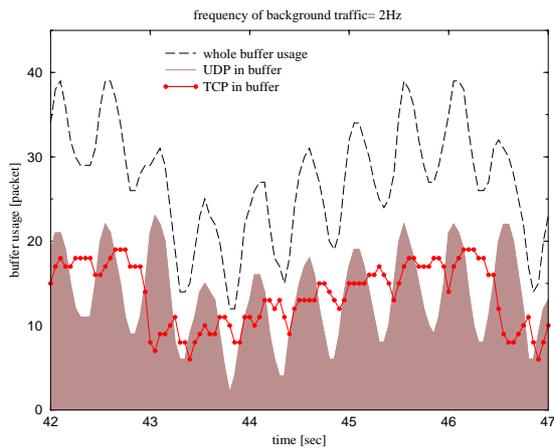


Fig. 9. Buffer utilization at small time-scale

Figure 9 clearly shows that UDP does not reach the buffer limit and TCP puts an increasing number of packets into the queue until a single packet loss. After this packet loss TCP performs fast recovery as it is shown by TCP trace in Figure 10. The figure shows the sequence number of TCP packets in time.

## V. EFFECT OF NETWORK PARAMETERS

An obvious question about the phenomenon is that how it depends on the parameters of the network? In this section we analyze how general the phenomenon of throughput reduction is and how it could be compensated by changing some of the parameters.

Firstly, we investigate the effect of link rate on the throughput characteristics. The transmission rate of the background traffic was also set by changing the link rate. This way we achieved an

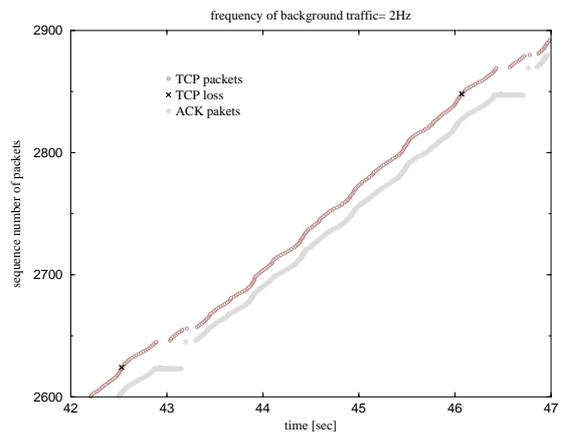


Fig. 10. TCP trace at small time-scale

average of 50% free capacity on the link similarly to previous scenarios.

As we discussed in the previous section, the throughput reduction is compensated by less buffer utilization of the background traffic stream below the characteristic time scale, i.e. at higher frequencies. As a consequence of higher link and traffic rates we can expect that the background traffic is able to fill in the buffer even at smaller time-scales, i.e. at higher frequencies. It should result that the compensation at smaller time-scales occurs only at lower values than before. Since the left part of throughput characteristics is determined by TCP's adaptation we can expect that the critical interval is extended towards smaller time-scales, i.e. higher frequencies. This effect is clearly shown in Figure 11. For comparison of throughput characteristics of different link rates the normalized throughput is depicted in the figure again.

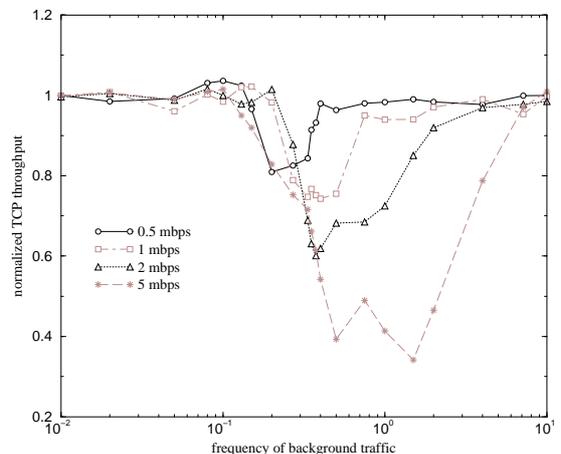


Fig. 11. Effect of link rate on throughput characteristics

Secondly, we investigate the effect of delay in the configuration. Delay has no significant effect on the background traffic and its buffer utilization but it reduces TCP's capability of adaptation to background traffic fluctuations. We can expect that TCP can reach high throughput only at larger time-scales, i.e. lower frequencies than before. That leads to the extension of critical interval towards larger time-scales. The effect of delay is shown in Figure 12.

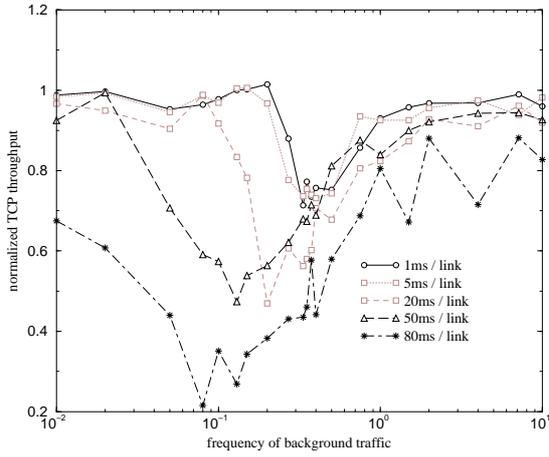


Fig. 12. Effect of delay on throughput characteristics

Thirdly, we analyze the effect of packet size on the throughput characteristics. Since packet size has no significant effect neither on transmission mechanism of TCP or on the buffer utilization of background stream, we can not expect great changes in the characteristics as the packet size varies. That is presented in Figure 13.

In this figure we also can see a strange effect where the throughput increases above 1. At that time-scale TCP is quite aggressive and gains even higher throughput than what would be left unused by the non-adaptive background flow. This is called as 'resonance effect' and discussed in details in [8]. However, it is out of the scope of this paper.

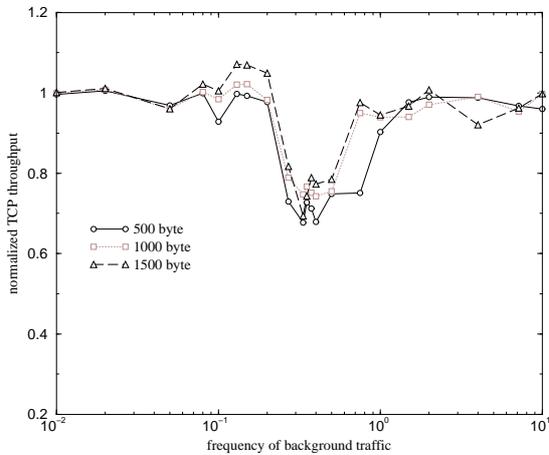


Fig. 13. Effect of packet size on throughput characteristics

## VI. COMPENSATION OF THROUGHPUT REDUCTION

As we discussed above, consecutive packet losses and the length of recovery period of TCP play important role in the shape of throughput characteristics. In this section we analyze the effect of those parameters.

The amount of consecutive packet losses can be reduced by applying larger buffer space. In this case the background stream is able to fill in the buffer only at larger time-scales, i.e. at lower frequencies, causing the compensation of throughput reduction below the time-scale where adaptation of TCP is not

proper enough. Effect of increased buffer size is depicted in Figure 14.

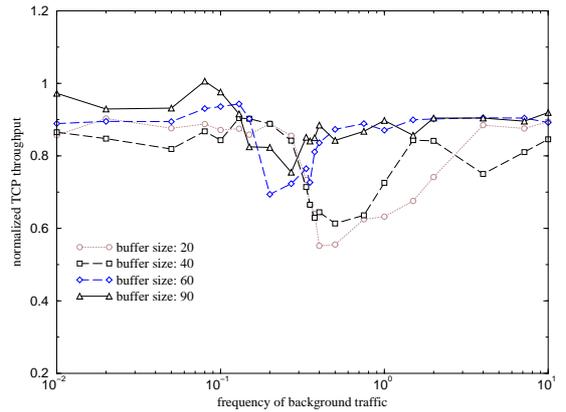


Fig. 14. Effect of buffer size on throughput characteristics

Another parameter that is effective in compensation of throughput reduction is the maximum RTO value of TCP source. RTO is calculated continuously based on the Round Trip Time (RTT) values of successfully transmitted packets. In case the calculated RTO value is higher than the maximum RTO value, the maximum RTO is an upper limit for time-out periods.

As we discussed in Section 2, RTO plays important role in adaptation to background fluctuations. Less RTO value leads to shorter period until TCP starts sending packets again. It is obvious to expect that TCP is able to utilize the free capacity in the link by better adaptation even at smaller time-scales if the RTO value is reduced. As a consequence of that the throughput characteristics should be close to 1 even in smaller time-scales where the compensation with buffer space left free by the background happens. The effect of changing the maximum RTO value is presented in Figure 14.

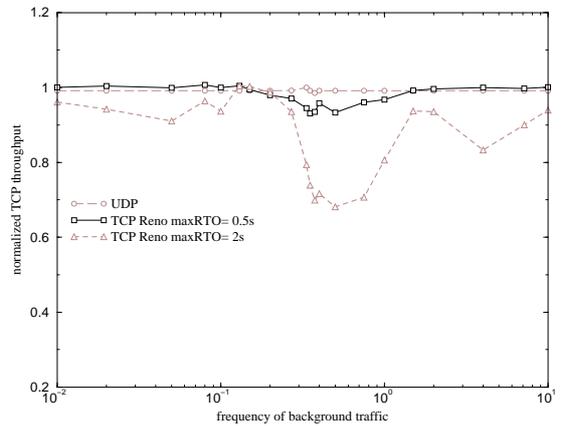


Fig. 15. Effect of maximum RTO on throughput characteristics

## VII. CONCLUSIONS

Our purpose was to analyze the effect of TCP's adaptivity on the throughput of the connection. We analyzed the throughput characteristics of TCP flows in case the background traffic fluctuates over several time-scales. We found that the throughput is reduced around a critical interval of time-scales. We also

found that the throughput is determined by frequency components present in the power spectrum of the background traffic independently.

In our research we focused on the origin of the throughput reduction, the effect of network parameters and the compensation of the effect.

We found that the origin of the effect is that around a characteristic time-scale of background fluctuations the TCP suffers consecutive losses and cannot utilize the available bandwidth. The effect does not take place at smaller time-scales as free buffer space is available for the TCP flow due to short waves of background traffic fluctuations. We presented that the effect can be compensated either by increasing the buffer size or decreasing the RTO of the source.

The detailed analysis of the effect in case of real network and traffic scenarios is subject of further studies.

#### REFERENCES

- [1] Gary R. Wright, W. Richard Stevens, TCP/IP Illustrated, Volume 2 The Implementation, Addison-Wesley Publishing Company 1995
- [2] J. Postel, Editor, RFC 793 - Transmission Control Protocol, 1981
- [3] R. Braden, Editor, RFC 1122 - Requirements for Internet Hosts Communications Layers, 1989
- [4] M. Allman, Editor, RFC 2581 - TCP Congestion Control, 1999
- [5] V. Jacobson, M. J. Karels, Congestion Avoidance and Control, Proc. ACM SIGCOMM '88, Stanford CS, August 1988
- [6] P. Karn, C. Partridge, Improving Round-Trip Time Estimates in Reliable Transport Protocols, Proc. ACM SIGCOMM '87, August 1987
- [7] A. Veres, Zs. Kenesi, S. Molnár, G. Vattay, TCP's Role in the Propagation of Self-Similarity in the Internet, Computer Communications, Special issue on Performance Evaluation of IP Networks and Services, Vol. 26, Issue 8, May 2003, pp. 899-913.
- [8] A. Veres, Zs. Kenesi, S. Molnár, G. Vattay, On the Propagation of Long-Range Dependence in the Internet, ACM Computer Communication Review, Vol. 30, No. 4, pp. 243-254, October, 2000.
- [9] T. Chahed, A-F. Canton, S-E. Elayoubi, End-to-end TCP Performance in W-CDMA/UMTS, IEEE ICC 2003, Anchorage, Alaska, USA, 11-15 May 2003
- [10] S. Floyd, V. Paxson, Difficulties in Simulating the Internet(2001), IEEE/ACM Transactions on Networking, Vol.9, No.4, pp. 392-403, August 2001
- [11] S. Blake, D. Black, M. Carlson, An Architecture for Differentiated Services, Request for Comments 2475, December, 1998