On The Effect of The Background Traffic on TCP's Throughput

Zsolt Kenesi
Ericsson Hungary Ltd.
Traffic Analysis and Network Perf. Lab.
H-1037 Laborc u. 1, Budapest, Hungary
zsolt.kenesi@ericsson.com

Zoltán Szabó, Zsolt Belicza, Sándor Molnár
Budapest Univ. of Technology and Economics
Dept. of Telecommunications and Media Informatics
H-1117 Magyar tudósok krt. 2, Budapest, Hungary
{szabo,belicza,molnar}@tmit.bme.hu

Abstract

The utilization of the free capacity by the TCP congestion control is not perfect due to the complex fluctuations of the background traffic resulting in reduced throughput. Different frequency components of the background traffic have different effects on the TCP performance. In this paper we present a comprehensive TCP performance evaluation study to understand the nature of the TCP adaptivity mechanism in frequency domain. Through simulations we analyze the impact of network parameters and also some alternatives for compensation of the throughput reduction. Finally, based also on the analysis of measured TCP traffic we investigate the robustness of TCP to fluctuations occur in actual network traffic.

1. Introduction

Due to the implemented congestion control algorithm TCP has an adaptive mechanism,which tries to utilize the free bandwidth on its path determined by the network parameters and the background traffic. Since congestion control was introduced in the Internet [2], it has proved its efficiency in keeping network-wide congestion under control in a wide range of traffic scenarios. However, full adaptation and therefore complete utilization of the free bandwidth is not possible, as the network does not provide prompt and explicit information about the amount of free resources.

In short, the 'adaptivity' of TCP describes how the protocol is able to adapt to the different network conditions and how it is able to utilize the changing free capacity. The term 'adaptivity' was introduced for the above-mentioned feature of TCP in [9] and [8]. It was also shown, that if a TCP connection shares a bottleneck link with a 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 the adaptation depends on the end-to-end path properties, i.e. round-trip time, widow size, etc. In other words, the presented analysis shows that TCP is sensitive to the dominating frequency components in the background traffic.

In this paper, our main purpose is to analyze the robustness of TCP to the different frequency components in the background traffic. In other words, we study the capability of TCP congestion control adaptivity to the free capacity in the path of the flow. Deep understanding of this effect is important for designing and dimensioning high performance TCP networks and making decisions on the applied mechanisms.

Considering the related work according to our knowledge there is no published work exactly on the specific topic of our paper. However, there are a number of papers which analyze the behavior of TCP from many different points of view. For example, a possible application of our analysis is the denial of service attack discussed by Kuzmanovic et al. in [4]. In this paper the authors show that TCP's congestion control algorithm is highly robust to diverse network conditions, its implicit assumption of end-system cooperation results in a well-known vulnerability to attack by high-rate non-responsive flows. Rubenstein et al. [6] present techniques based on loss or delay observations at end-hosts to infer whether or not two flows experiencing congestion are congested at the same network resources. The background traffic characterization from different purposes are also discussed in several papers: Gunawardena et al. [1] investigate several background traffic models in their admission control design; Ribeiro et al. [5] develop a multifractal parametric cross-traffic model.

2. Throughput characteristics

Individual TCP flows are transmitted through several hops in actual networks. The throughput of the TCP connections is determined by a bottleneck in the path of the flow. Although the location of the bottleneck is changing with the time, it is a good approximation if the behavior of the TCP flow is analyzed in a single bottleneck.

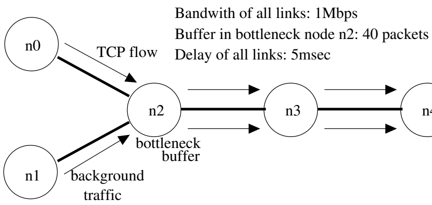

Figure 1. Topology of the applied network model

Figure 1 shows the applied network model. All the simulations were performed in Network Simulator (ns-2). In most of the networks adaptive and non-adaptive traffic is transmitted through the same infrastructure. In networks where diffserv architecture is not implemented, TCP shares the path with many concurrent flows based on TCP or on different protocols like the non-adaptive UDP. Provided, that an individual TCP flow observes the whole background traffic as a non-adaptive stream, the background traffic can be modeled by one non-adaptive traffic flow. In the applied model the background traffic is represented by a non-adaptive UDP traffic and generated in a way, such that it fluctuates on a limited, narrow time-scale. To limit the time-scale under investigation, the rate of the background traffic is approximated by a constant amplitude sine wave of a given frequency f: $BW(f,t) = Asin(2\pi ft + \alpha) + m$, where α is a uniformly distributed random variable between $[0; 2\pi]$. (The process $BW(f,t)$ is a stationary ergodic stochastic process.) The parameters of the function $BW(f,t)$ is set in the way that the background traffic rate changes between zero and the maximum bandwidth. As a result, 50% of the link is left free by the background traffic independently from the frequency of its fluctuations. By applying such kind of traffic we can measure how the throughput of TCP flow depends on the frequency of background fluctuations independently from other effects. During the simulations we investigate TCP's throughput when a single, long and greedy TCP flow shares bottleneck with the periodically changing non-adaptive background traffic that utilizes half of bandwidth in average. We define the throughput characteristic of TCP as the average throughput of the TCP flow as a function of the frequency of the fluctuations in the background traffic rate.

For determining the throughput characteristic we analyzed the stationary state of the configuration, therefore all simulations were performed at least for ten minutes and the transient interval was removed from the beginning of the log files. To be able to compare different scenarios presented later we normalized the TCP throughput, therefore the value 1 means the average free capacity in the link, which is 500 kbps in this case. Figure 2 shows the throughput characteristic for different TCP versions.

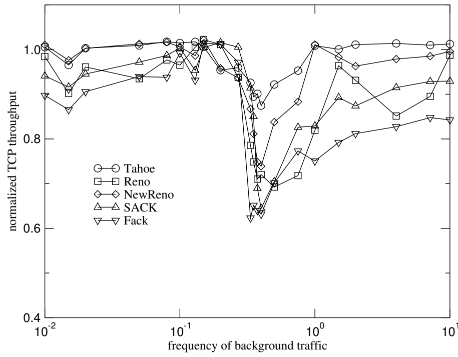

Figure 2. Throughput characteristic of different TCP versions

As it can be observed, the throughput characteristic depends on the frequency of fluctuations in background traffic rate. The throughput varies between 90-100% in most of the cases meaning that TCP is able to utilize the free link capacity quite efficiently. However, there is an interval from $0.3/sec$ to $1/sec$ where the throughput is reduced to 60-70% of the free capacity. We refer to this frequency region as critical time-scale or critical interval. Additionally, the analysis for common versions of TCP, i.e. Tahoe, Reno and NewReno shows that at low frequencies the throughput curve is almost invariant to the TCP version but there is difference in the degree of throughput reduction around $0.3/sec$ and above.

Our goal in this paper is to answer the following questions related to the throughput reduction of TCP. What causes this throughput reduction? What is the effect of different network parameters? How could it be eliminated? How significant is this effect in case of actual traffic?

3. Analysis of throughput characteristics

The throughput characteristic presented in Figure 2 can be separated into three different intervals. The first is in the left side of the figure that describes the characteristic 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. Further on we discuss each interval separately and present the results based on simulations with the Reno version of TCP.

Above the critical time-scale the background traffic rate varies with relatively low frequency. It varies relatively slowly, thus TCP is able to follow the changes of background traffic and utilize the free capacity. The rates of both TCP and UDP traffic is depicted in Figure 3.

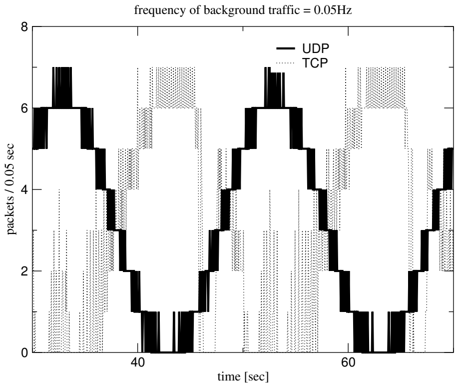

Figure 3. TCP and slowly varying background traffic

As the throughput of background traffic increases, TCP packets are dropped in the bottleneck buffer. When the background traffic utilizes the whole bandwidth, TCP is unable to send packets and is forced into time-out phase. After timeout TCP sends packets again starting from a slow-start phase and has enough time among the waves of background traffic to increase its congestion window and transmit packets on the maximum link rate. Figure 3 shows clearly this effect. Obviously, when TCP's rate follows the fluctuations of background traffic the frequency of sine wave of the background traffic appears not only in the rate but also in the process of the congestion window of TCP.

The second part of the characteristic is the critical interval that covers the range from $0.3/sec$ to approximately $1/sec$ where the throughput is reduced. Similarly to the previous case, TCP tries to adapt to the fluctuations of the background in the critical interval when it loses consecutive packets and performs timeout period when the background traffic fills the link in. After the 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 increase its congestion window and send as many packets during a cycle as in the previous case. In other words, TCP is unable to reach high transmission rates between two waves of the background traffic. This periodic process can be followed in Figure 4 that shows the TCP trace in the critical interval. The figure shows one period of the background traffic fluctuations. Time-out phases that are repeated in each period can take about 2 seconds until TCP starts sending packets again. Since the background traffic peak arrives shortly after TCP starts with slow-start phase, consecutive TCP packet loss and time-out occurs. Figure 5 shows buffer utilization in the bottleneck. It can be observed that after most timeout phases TCP and UDP background traffic increase their transmission rate at same time. This leads to the discussed effect and, through to it, to throughput reduction.

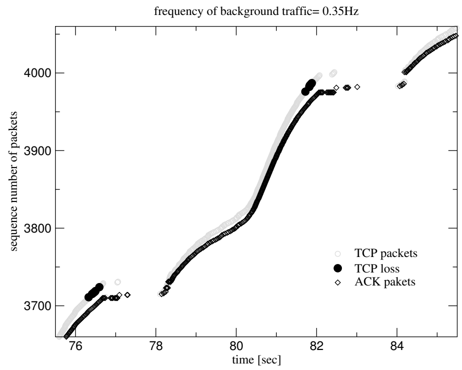

Figure 4. TCP trace in the critical interval

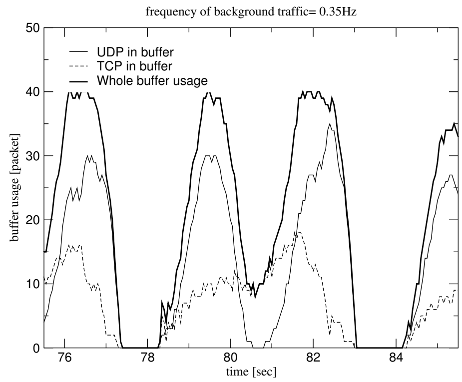

Figure 5. Utilization of the bottleneck buffer in the critical interval

If the background traffic changes at higher frequencies than the critical interval, the throughput of TCP will be growing and reduction will completely disappear. The effect that causes the increase of throughput is that the buffer utilization of background traffic is lower on higher frequencies. The buffer utilization by the background traffic depends on the amount of data sent during a peak period that is less at higher frequencies. This leads to the effect that background traffic utilizes less buffer space and TCP avoids consecutive packet losses even during the periods where the background traffic rate reaches its peak.

We also applied a simplified analytical model to describe the throughput reduction. The frequency intervals discussed above can be estimated based on modeling the adaptation of TCP and the empty space in the buffer as a function of background fluctuations. The lower frequency limit of critical interval is determined by the adaptation capability of TCP. The effect and the way of estimating this frequency that describes TCP adaptation was published in our previous works [9] and [8], as follows $f_{min} = 1/T$, where T is the period length of TCP congestion window process in case of no background traffic is present. T can be approximated as $T = (B/C + d)(B + Cd)/2$, where B is the buffer size in packets ($B = 40pkts$), C is the bottleneck link rate in packets per second ($C = 125pkts/sec$, TCP packet size is 1000 Byte), and d is the total round-trip propagation delay in seconds ($d = 30msec$). This gives us the frequency of $f_{min} = 0.13/sec$ that agrees with simulation results. The upper frequency limit of the interval where the throughput is reduced can be estimated by calculating the empty buffer space at different frequencies. In our work we found that the throughput reduction is not significant if TCP is constantly able to utilize approximately one third of buffer space. The frequency where it happens is derived in [3] and can be expressed as $f_{max} = 3C/(4\pi B)$, where C is the bottleneck link rate in packets per second and B is the buffer capacity in packets. The result of $f_{max} = 0.75/sec$ agrees with the simulation results.

4 Effect of network parameters

First, we investigate the effect of changes in link rate on the throughput characteristic. The peak transmission rate of the background traffic was also set to the maximum link rate in each case. By this way the average free capacity on the link was 50% which is similar to the previous scenarios. As we discussed above, the observed throughput reduction is compensated by less buffer utilization of the background traffic at high frequencies, i.e. at small time-scales. As a consequence of higher link and traffic rates we can expect that the background traffic peak contains higher amount of data packets and is able to fill in the buffer at even higher frequencies. It would result in that the critical interval is extended towards higher frequencies. This effect is clearly shown in Figure 6. For comparison of throughput characteristic of different link rates the normalized throughput is depicted. As the link rate is increased from 0.5 Mbps to 5 Mbps, the reduction of throughput can be observed at frequencies from $0.4/sec$ to $7/sec$. As the interval is more extended the effect of reduction is more significant. At 5 Mbps the TCP is able to utilize less than 50% of the free capacity on the link in a frequency interval of $0.4/sec$ to $2.2/sec$.

Second, we investigate the effect of changes in (each) link delay in the applied network model. 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 utilize the free capacity on the link only at lower frequencies than before. That leads to the extension of the critical interval towards larger time-scales. The effect of the increased link delay is shown in Figure 7 where again the normalized throughput is depicted. As the delay increases,

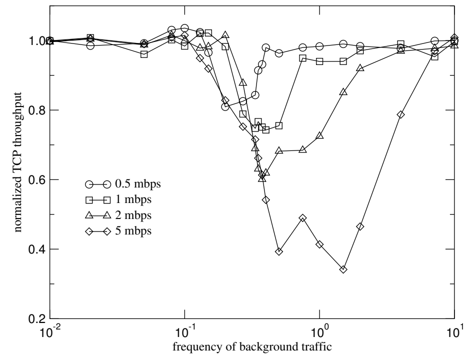

Figure 6. Effect of link rate on throughput characteristic

the reduction of throughput can be observed in wider frequency interval.

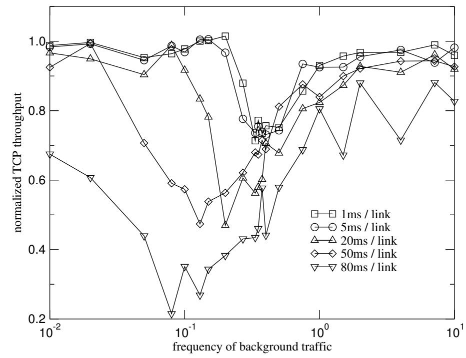

Figure 7. Effect of link delay on throughput characteristic

5 Compensation of throughput reduction

As we discussed above, consecutive packet losses and the length of recovery period of TCP play important role in the throughput reduction. We can expect that reducing the amount of consecutive packet losses or making the recovery period of TCP shorter can eliminate the throughput reduction.

Applying larger buffer space can reduce the amount of consecutive packet losses. In this case the background traffic is able to fill in the bottleneck buffer only at lower frequencies, i.e. at larger time-scales. As a result, the upper frequency limit of the critical interval is less than in the scenarios where smaller bottleneck buffer is applied. The effect of the increased bottleneck buffer size is depicted in Figure 8. We can observe, that TCP is able to utilize more than 70% of the free link capacity when a buffer of 90 packets is applied.

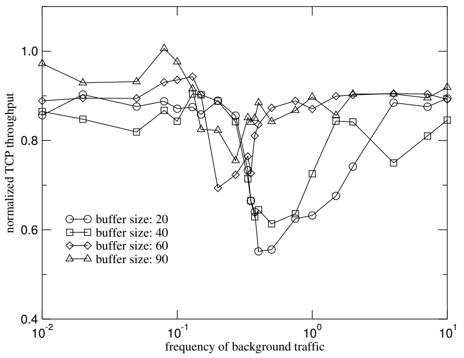

Figure 8. Effect of bottleneck buffer size on throughput characteristic

Another parameter that is effective in the compensation of throughput reduction is the maximum RTO value of TCP source. This parameter can be set independently from the network parameters. The value of RTO is calculated continuously based on the measured 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. Less RTO value leads to shorter period until TCP starts sending packets again. It is obvious to expect that the throughput of TCP would increase if the time-out phase were shorter. The effect of changing the maximum RTO value is presented in Figure 9 with buffer size 40. For comparison with the normalized TCP throughput the same measure for UDP traffic is also shown in the figure.

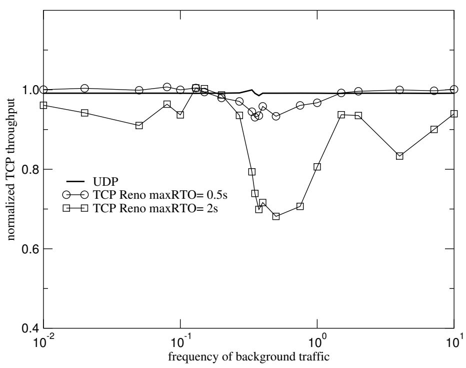

Figure 9. Effect of maximum RTO on throughput characteristic

6 Actual traffic scenarios

In this section we investigate how the effect of throughput reduction is relevant to actual network scenarios and under which circumstances it should be taken into consideration.

First, we analyze how the throughput depends on the complexity of the frequency spectrum of the background traffic. Since in actual networks the background traffic is not limited to a single time-scale, we analyze the scenario when more than one frequency components are present in the spectrum of the background traffic. Our motivation for this analysis is to investigate whether TCP's throughput is determined by fluctuations of background traffic on different time-scales independently from each other or not. We performed a large number of simulations with a wide range of background traffic mixes composed of two frequencies f_1 and f_2 in the original network model presented in Section 2. The rate of background traffic was changed based on the following equation: $BW(f_1, f_2, t) = C_1 sin(2\pi f_1 t) + C_1 sin(2\pi f_2 t) + C_2$, where C_1 and C_2 are constants. Our analysis covers a whole plane with the two frequencies in the range of $[0.05; 10]/sec$. Figure 10 shows that if at least one of the frequency components is in the critical frequency interval identified in Section 2, the throughput is reduced. If both frequency components are in the critical interval, the reduction of the throughput is more significant. This result indicates that TCP's throughput is determined by the two frequency components independently.

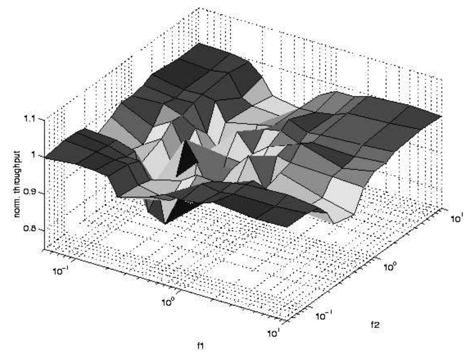

Figure 10. Throughput characteristic of TCP when the background process is composed of two frequencies

For the analysis of the effects of throughput characteristic in case of actual background traffic we applied two different white noise models to easily construct background traffic with many frequency components. We performed simulations with a special noise, the Fractional Gaussian noise (FGN). FGN process frequently appears as the limit process of traffic aggregations, which means that in Internet backbones we can often see traffic well modeled by FGN [7]. We performed several simulations with different FGN parameters, like Hurst parameter, standard deviation, mean, duration time. In all cases we found that the throughput reduction was only about 2-5%.

To get a deeper understanding of the effect of spectrum in actual traffic cases we also applied bandwidth limited white noise model. From the simulations with more frequency components and FGN we concluded that the degree of throughput reduction depends on the ratio of the spectrum within and outside the critical frequency interval. In case of white noise only a fraction of its spectrum can be found within the critical interval, therefore it may result in a TCP throughput close to 100%. To present this effect we performed simulations with different background traffic mixes whose spectrum was constructed as a bandwidth limited white noise. Since we identified the critical interval in the range of $0.3/sec$ to $1/sec$, the lower frequency limit of white noise was $0.1/sec$ and the upper limit changed from $1/sec$ to $30/sec$ in different simulations. We applied the original network model presented in Section 2. We can expect higher throughput of TCP as the fraction of the background traffic spectrum within the critical interval decreases, i.e. the upper frequency limit of white noise spectrum increases. Figure 11 shows the effect clearly. Please note that the normalized TCP throughput does not grow above 97% even in the case where the upper limit of white noise spectrum is $30/sec$ and the fraction of spectrum in the critical interval is less than 3%.

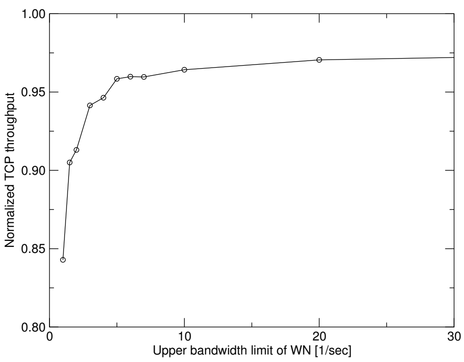

Figure 11. Throughput characteristic of TCP with white noise background process

We also investigated the spectral characteristic of different transport layer protocols, such as TCP and UDP traffic. Among the traffic traces we analyzed were the widely used log files of Internet Traffic Archive (http://ita.ee.lbl.gov/). Our general finding is that the traffic in actual cases has almost constant spectrum close to white noise. It means, that the throughput reduction of TCP flows in case of complex background traffic with wide spectrum is not significant.

7 Conclusions

Our results reveal TCP throughput reduction as a function of frequency components of the background traffic. We presented and analyzed the throughput characteristic of TCP flows in case the background traffic fluctuates as a function of sine wave on different time-scales. We found that the throughput is reduced within a critical interval of time-scales and we explained the origin of the reduction. We analyzed the effect of network parameters on the throughput characteristic and found that increasing link rate and link delay extends the critical interval. Besides the analysis of the effect we also showed that it could be compensated either by increasing the buffer size or decreasing the maximum RTO value of the TCP source. We also analyzed the significance of throughput reduction if actual background traffic is present. We found that the degree of throughput reduction depends on the ratio of spectrum components within and outside the critical frequency interval. In case of actual traffic scenarios we found that the spectrum is close to white noise and the effect does not cause significant throughput reduction. However, in case of the major frequency components are concentrated in the critical interval the effect has to be considered.

References

[1] D. Gunawardena, P. Key, and L. Massoulie. Network characteristics: modelling, measurements and admission control. *Workshop on Quality of Service IWQoS*, 2003.

[2] V. Jacobson. Congestion Avoidance and Control. In *ACM SIGCOMM 1988*, Aug. 1988.

[3] Z. Kenesi, Z. Szabó, Z. Belicza, and S. Molnár. Performance Study of TCP Adaptation. *Technical Report, Budapest University of Technology and Economics*, 2004.

[4] A. Kuzmanovic and E. Knightly. Low-rate TCP-targeted denial of service attacks. In *ACM SIGCOMM 2003*, 2003.

[5] V. Ribeiro, M. Coates, R. Riedi, S. Sarvotham, B. Hendricks, and R. Baraniuk. Multifractal cross-traffic estimation. In *Proc. of ITC Specialist Seminar on IP Traffic Measurement*, Sept. 2000.

[6] D. Rubenstein, J. F. Kurose, and D. F. Towsley. Detecting shared congestion of flows via end-to-end measurement. In *Measurement and Modeling of Computer Systems*, 2000.

[7] M. S. Taqqu, W. Willinger, and R. Sherman. Proof of a fundamental result in self-similar traffic modeling. *ACM Computer Communication Review*, 27, 1997.

[8] A. Veres, Z. Kenesi, S. Molnár, and G. Vattay. On the Propagation of Long-Range Dependence in the Internet. In *Proceedings of ACM SIGCOMM 2000*, Sept. 2000.

[9] A. Veres, Z. Kenesi, S. Molnár, and G. Vattay. TCP's Role in the Propagation of Self-Similarity in the Internet. In *Computer Communications, Special issue on Performance Evaluation of IP Networks and Services, Vol. 26, Issue 8, May 2003*, 2003.