Traffic Engineering for Internet Applications

Ajit K. Jena and Adrian Popescu

Department of Telecommunications and Signal Processing
Blekinge Institute of Technology
371 79 Karlskrona, Sweden

ABSTRACT

The focus of the paper is on resource engineering for supporting Service Level Agreements (SLAs) in IP networks. SLA at both link level and application level are considered. Using an object-oriented simulation model a case study is presented for client-server interactions generated by mixed traffic conditions in a Frame Relay (FR) WAN. Performance issues of Short Range Dependence (SRD) and Long Range Dependence (LRD) traffic under different resource control regimes are compared. The results show that major portion of the end-to-end delay comes from the queueing delay at the WAN ingress points, which is due to the significant bandwidth differences that may exist between LAN and WAN link layers. The results also highlight the role TCP window size and FR PVC control mechanisms play in the provision of delay performance for Internet services.

Keywords: Traffic Engineering, Service Level Agreement, Long Range Dependence, Delay Performance

1. INTRODUCTION

The recent advent of IP-based telephony and multimedia services has offered providers of teleservices a great opportunity to push for the development of new services that integrate different media and also exploits the power of intelligent terminals. Communications can now encompass not only voice, but also audio, video, shared applications, and even virtual reality. At the same time, Internet, with the access to information that is offered to everyone, is acquiring an everincreasing importance for bussiness community as well as for the population at large.

Recent proliferation of heterogeneous networking technologies has also created new opportunities and challenges for network deployers. On one hand, there is the choice of extension of the new teleservices to cellular phones by developing more powerful user interfaces and also gateways to the public network. On the other hand, there is the challenge of the provision of Quality of Service (QoS) guarantees on an end-to-end basis, i.e., at the level of applications running in user space of general purpose operating systems. Furthermore, while the access to Internet has become a commodity, there is also a strong need for differentiating the service offerings of different Internet Service Providers (ISPs), under the frame of QoS as well as of value-added Internet services.

The focus of the paper is laid on resource engineering for supporting Service Level Agreements (SLAs) in IP networks. Two levels are considered here, namely SLA at the link level and SLA at the application level. A case study is considered for these purposes, which is based on an object-oriented simulation model for client-server interactions that is able to capture real message exchanges as well as traffic payloads as seen in typical Internet access scenarios. The OPNET simulation environment is used and the parameters for the simulation are derived from real Internet traffic flows. Internet applications like HTTP and SMTP are considered in the study, and statistical models have been developed that are based on real traffic measurements. The core of the network is formed by a Frame Relay (FR) WAN, with a topology that allows for diverse traffic mixtures.

The focus of the experiments is laid on the delay performance together with controls and traffic engineering to cater for a specific SLA. The impact of link and transport layer controls on Internet traffic is studied. The influence of Long Range Dependence (LRD) properties on the delay performance at packet and application level is reported as well. Starting with Short Range Dependence (SRD) traffic conditions, a bandwidth and buffer allocation scheme is worked out such as a specified SLA is obtained at both the packet and application level. Subsequently, under

Further author information: (Send correspondence to A.P.)

A.K.J.: E-mail: akj@its.bth.se

A.P.: E-mail: adrian@nada.kth.se; Telephone: +46 8 7906560; Fax: +46 455 385667

LRD traffic conditions, the utilization profiles and the resource allocation schemes are adjusted so as to achieve similar performance levels.

The results show that major portion of the end-to-end delay comes from the queueing delay at the WAN ingress points, which is due to the significant bandwidth differences between LAN and WAN link layers. The results also bring out the significant role TCP window size and Frame Relay PVC control mechanisms play in the provision of delay performance for Internet services.

The paper is organized as follows. The first part (Section 2) is devoted to Service Level Agreement (SLA) definitions and descriptions. Performance metrics for SLA are also considered. The second part (Section 3) is focused on the description of the case study considered for experiments. In particular, specific settings for various aspects in the model (e.g., mix of applications, TCP parameters, PVC settings, and diverse router parameters) are presented in a detailed way. The next part (Section 4) reports the performance obtained at the packet level, for both SRD and LRD traffic. SLA at the packet level is also reported. Finally, section 5 is devoted to the performance obtained at the application level.

2. SERVICE LEVEL AGREEMENT

An SLA is a formal definition of the relationship that exists between a supplier of services and customers. SLA addresses issues like the nature of service provided, reliability, responsiveness, the process for reporting problems, time frame for response to reported problems and problem resolution, methodologies for auditing service levels, penalty and escape clauses, etc.¹ In the context of Internet, service issues have become an important tool for Internet Service Providers (ISPs) to retain their client portfolio and to attract new customers. An ISP may provide specific SLAs to customers who may use the ISP network in different ways, which can be generally described as a combination of three basic modes: 1) to access the public Internet; 2) to interconnect two or more sites; and 3) to access proprietary, industry specific networks such as in the case of enterprise networks. However, the terms of SLA that govern each of the access profiles are likely to be different. For example, when the objective is to connect two sites of a given customer, the assurances about the performance level will focus on the path between the pair of access routers and is likely to be more clearly defined and controlled. On the other hand, when the goal is to access the Internet in general, the ISP may not be in a position to control the performance obtainable. However, in such cases, the ISP may be in a position to certify that the performance obtained will be more or less decided by the external Internet cloud and that the ISP's access network will not be the bottleneck.

In the context of computer networks, service specification for SLA can be done at different protocol layers. At the application layer, the focus is on application sessions. Primitives like throughputs, object access delays, and transaction update turnaround time are relevant here. Further, the overall performance aspects obtained at the application layer is a combination of: 1) the stochastic nature of payload contents sizes; 2) the application layer protocol; and 3) the transport and lower layer protocols. Application layer protocol works as a feedback loop on the end-to-end path and thus influences the performance significantly.

Another key issue is that the performance at the application level is also dependent on the number of concurrent application sessions. A SLA specification geared towards the *busy hour traffic* may be more meaningful here. At the transport layer, connection throughput and goodput are considered for SLA. On the contrary, at the link layer, the focus is on packets. Given the likely size of user population at the client site, the major points considered here are the bandwidth allocation and port buffer sizing. Also, there may be rate and burst control issues as well (e.g., in ATM or FR). Depending on particular environments, a combination of these layer specific SLAs may be chosen.

There are two aspects related to SLA of networked systems: availability and responsiveness. These are applicable at both the application level and the network layer. Host or router uptime, link outage frequency, and error rates, they all fall in the category of system availability. The other important issue of network connectivity can be monitored and detected using diverse network management tools (e.g., SNMP). Further, some of the major points that can be grouped under the category of responsiveness are: one way end-to-end delay at the link layer; application level turnaround time; TCP connection setup latency; TCP retransmissions; packet loss rates; and available bandwidth. Some of these metrics can be collected from regular SNMP or RMON type statistics databases whereas others (e.g., TCP and application layer metrics) can only be audited via sophisticated monitors. Some SLAs take the form of guaranteed transmission profiles (such as CBR in ATM and CIR in FR PVCs and SVCs) to allow users to inject specific amount of data into the network with an assured bandwidth and loss characteristics.

Figure 1. Network architecture for SLA experiments

3. CASE STUDY: SLA ON A BEST-EFFORT NETWORK

The goal is to study a client-server model driven with Internet applications like HTTP and SMTP. This involves understanding and simulation of application sessions and the constituent transactions. Towards this end, an object-oriented view has been developed into a client-server simulation framework.² The model closely mimics the real-life events that occur in an Internet consisting of clients and servers. The client and the server nodes running the above-mentioned applications are connected by an Internet cloud (fig 1). The OPNET simulation tool³ is used as a simulation environment.

The ingress and egress points to the global Internet is via routers, which feed into a FR WAN link via Frame Relay Access Devices (FRADs). The clients and the servers are organized into two separate sites. The client nodes are 10BaseT Ethernet hosts and the server nodes are 100BaseT hosts. Further, a fractional Brownian model (fBm) is used to model the network traffic (according to the so-called "Joseph effect"),⁴ whereas the application traffic is modeled with a mixture of Uniform (or Lognormal) and Pareto distributions.

Most of Internet client-server applications can be modeled as ON-OFF sources. The client makes a REQUEST for a service at a specific server (client side ON state) and the server responds with a RESPONSE (server side ON state). The ON periods on both sides are interspersed with OFF periods. The stochastic aspects of ON duration derive directly from those of the application layer message sizes.² The high variability in the content sizes results in clusters of packets that enter the network. The OFF periods are typically due to user inactivities and the variability in the OFF process may be due to different types of users accessing network services concurrently. The level of activity of human users generally varies from user to user.

In order to understand the end-to-end transactions, it is important to quantify the effects of network control mechanisms on the application-level properties. Transporting packets across the network between the application end points is subject to delays, errors, and losses that occur at congested router/switch. Even though the presence of LRD has been known for quite some time now, the impact on end-to-end performance at the packet or application layer is not fully known. The aim of the simulation experiment is therefore focused on the end-to-end delay performance while maintaining packet losses within specified limits. The motivation for this is that as far as end-user performance is concerned, packet losses and delay are related. Packet losses result in retransmissions at the TCP layer and indirectly affect the object access latency at the application layer. By minimizing losses and targeting the delay performance the overall performance may be optimized and at the same time may result in simplification of the simulation experiment setup.

3.1. Setting Up the Environment

Global Internet is an example of best effort collaborative service spanning across many geographical, cultural, and ethnic limits. Individual organizations own and administer parts of this network as per their own policies. When the users access services on the Internet, the end-to-end path sometimes literally spans across continents. Along the end-to-end path, the packets cross many different: 1) organizations; 2) network technologies; and 3) routing and switching policies. Because of these diverse conditions, in general, it is very difficult to make performance evaluations to a specific degree of accuracy. Simulation oriented investigations may however be possible under some specific operating conditions. Even so, it is extremely important to identify important environmental conditions and choose appropriate parameter values for them. Some of the most important parameters, and suitable choice for values of the same are presented in this section.

3.1.1. Network architecture

The case study focuses on a network architecture that is appropriate for a corporate Intranet or an Interconnection of two sites of an organization separated by hundreds of kilometers of geographical distance (fig 1). The network consists of three components: the local networks labeled “Site 1” and “Site 2”, and the “IP Cloud” interconnecting them. The local networks are owned and administered by the corporate organization whereas the ISP handles the interconnectivity. The ISP’s infrastructure to support this connectivity is considered to be a single T1 link.

The relation between the corporate organization and the ISP takes the form of a SLA. The focus of the simulation study is to formulate the SLA and also find out the implied performance bounds when the operating conditions are varied in a certain range.

3.1.2. Mix of applications

WWW constitutes today about 70-80% of the Internet traffic load and the second application in terms of traffic volumes seems to be E-mail.² On the contrary, the use of FTP is diminishing because of: 1) the effectiveness of WWW as a means of content distribution; and 2) security holes related to the FTP service.

Given this background, a mix of WWW and E-mail traffic has been initially considered as an application environment for the experiment. The proportion of the traffic belonging to the respective categories is controlled by controlling the number of client hosts on either sites and/or by adjusting the session arrival rates. The simulation investigates the effects during a *busy hour*. During the *busy hour*, session arrivals for both classes of applications are modeled as Exponential processes with a mean inter-arrival rate of 15 seconds.

The second important aspect considered here is the choice of the parameters for content distributions at the application layer. This has been adjusted so as to create two separate environments: one with LRD traffic characteristics (Hurst estimate of roughly 0.9) and the other with SRD traffic (Hurst estimate of roughly 0.55). The WWW traffic is considered as the main traffic and E-mail traffic serves only as a background traffic. The LRD properties of the mixture traffic are determined by the traffic with highest LRD component.⁴ Accordingly, the SMTP traffic parameters were selected so as to result in a traffic mix where the behavior of the WWW traffic is dominant.

3.1.3. TCP parameters

The TCP layer is an important element in the end-to-end path of the application service loop. Each connection uses TCP window mechanism to dynamically determine the available capacity in the network and accordingly to decide how much data it can safely have in transit. The TCP window size is an important parameter in flow and congestion control of the Internet traffic.

However, when implementing an application service, there is no guiding principle as to how to select a particular window size. The implementation of the same application service under different operating systems tends to use very different window sizes when initiating TCP connections (e.g., Sun Solaris uses a window size of 8760 whereas RedHat Linux uses a value of 32120).⁵ It is unclear how this kind of shaping mechanism interacts with the application layer properties and what is the effect on performance. Accordingly, the simulation experiments are performed for window sizes of 8760, 16384, and 32120.

3.1.4. PVC settings

The physical connectivity between the two sites of the network under investigation is provided by a standard T1 link with a FR PVC creating a logical channel over it. The characteristics of the logical channel are described as a set of three parameters CIR , B_c , and B_e .⁶

To the best of our knowledge, there are no methodologies actually available as to how to select parameters to ensure optimal resource allocation. In order to find a meaningful value for CIR , a calibratory approach has been adopted. To begin with, a *Zero CIR* allocation strategy is chosen for the PVC. Under such conditions, the FR rate control does not take any effect. All packets that enter the PVC are marked *discard eligible*. Even though all packets are admitted into the network, their delivery is subject to resource conditions on the physical path. A set of about 100 client hosts are used to populate the two sites in network to allow for sufficient number of TCP connections on the link. In the limiting case, this behavior will approach fBm conditions. Further, to start with, the application parameters for the sources are chosen so as to create more or less Markovian conditions. The link

resource conditions are basically the port buffer size at the routing or switching stages. The buffers were optimized to keep packet loss within prescribed limits. With this setup, the simulation experiment was performed and the packet trace collected. From this, the byte count was collected over non-overlapping intervals of 10 seconds and the average bandwidth computed. A mean value of about 50 Kbytes was obtained for the bandwidth. Using this as a thumb rule, a *CIR* of 64 Kbytes/sec seems to be adequate. There are a few intervals where the *CIR* limit is indeed exceeded but they can be easily accommodated by the excess physical bandwidth available on the link. The bandwidth allocation will be essentially studied as a scaling problem. Having decided on a *CIR* allocation for the traffic with SRD properties, the number of hosts will be reduced gradually for the LRD case while maintaining similar end-user QoS levels, thus determining the bandwidth allocation applicable for such traffic.

The choice for the burst control parameters B_c and B_e has also no formal principle. The ratio $\frac{CIR}{B_c}$ is denoted as T_c , the measurement interval. Survey of relevant literature has revealed that prescribed values for T_c vary in the range 0.1 to 1.0 but the reasoning behind the specific limits are unknown.⁷ For the parameter B_e , no suitable ranges were found. Again, choosing an exploratory strategy, the values for the parameters T_c were varied over a range and B_e was varied in relation to T_c . Simulation runs were carried out for several such profiles in order to estimate the performance impacts of such variations.

3.1.5. Router/Switch parameters

The choice of parameters for router or switching elements along the end-to-end path is an important factor in performance estimations. The first parameter of interest is the time taken per packet by the routers and switches. This is a value independent of the packet size. The routers use their FRAD interface to put the packet onto the WAN and similarly the switches also have to just do a table lookup to determine the output port. As each packet arrives at the input port of a switch, it raises an interrupt. The switching element determines the output port for the specific packet and schedules the packet to join the queue at that port exactly after one switching time delay. However, the packets may have to wait at the output port because of multiplexing. The time delay per packet at the switching or routing stages has been set to 20 μsec . This is a typical value for the modern fast packet processors. The other important parameter is the size of the port buffer. Here too, a heuristic approach was applied. For different combinations of TCP window sizes and PVC burst control parameters, the loss performance was monitored via simulation. The port buffers were dimensioned not in bytes or packets but based on the speed of the link to which the port is attached. For example, for a T1 link, a port buffer dimension of 100 milliseconds indicates a port buffer size of 15360 bits. Through this experimental procedure the link buffers are gradually increased so as to keep the losses limited to well below 0.5%. For the experimental environment used in the case study, a buffer of 250 milliseconds seems to be adequate for our purposes.

4. PACKET LEVEL PERFORMANCE

The packets generated by the application sources are aggregated inside the edge router. FR control mechanisms feed these packets into the network subject to the packets meeting the PVC control criteria. For the given case study, there is spare capacity available inside the network because the PVC *CIR* allocation is one third of the link capacity. Under such conditions, once the packets are injected into the network (as per the configured PVC settings), the performance is not much affected by link conditions. The packets also encounter switching delay that is influenced by the switch hardware. The bottleneck link does not cater to any cross traffic flows and thus behaves as a single delay line consisting of a cascade of queueing stages. At a typical switch, the queueing delay experienced by successive packets can be expressed using Lindley's equation.⁸

The simulation results reported in this section are only about the queueing delays. The transmission and propagation components of the delay are disregarded in order to focus purely on the traffic generation process.

The queueing process at the ingress router is predominantly due to: 1) bandwidth mismatch between Ethernet LAN and the FR WAN; 2) number of concurrent application sessions; and 3) burst characteristics of the generated traffic. Further, we consider the burst characteristics as being captured by the Hurst estimates of the byte count process at the link layer. The impact of Hurst estimates on the queueing process is shown in fig. 2. The empirical CCDF of the delay across the PVC for both SRD and LRD traffic scenario are plotted using a logarithmic scale on the y-axis. The almost linear behavior in the SRD case indicates an exponential behavior whereas the LRD case appears to show behavior typical of Weibull distributed random process. The delay characteristics shown in the

plot are for the case of high utilization conditions with *Zero CIR* PVC settings and when the TCP connections use a window size of 8760.

Further, for the above case, the CCDF of the end-to-end delay is compared with the delay experienced by the packets at the PVC entry point (fig 3). As the figure indicates, the queueing delay at the WAN ingress point dominates the end-to-end performance, with the conclusion that the effect of queueing and switching delays at the switches inside the WAN is marginal.

The effect of utilization levels on the PVC delay is shown in fig 4. Three different load conditions are considered. It is observed that, as the utilization levels are increased, the upper tail of the PVC delay distribution becomes more and more pronounced. This happens both with SRD and LRD traffic but the relative changes are different.

Finally, the PVC delay of the high load case of SRD traffic is plotted against the low load case of LRD case (fig 5). While the link utilization conditions may have a role in contributing to the queueing process, the profound influence of LRD contents on PVC delay is evident.

Figure 2. Effect of LRD on the end-to-end packet delay across the PVC

Figure 3. Comparison of end-to-end delay vs queueing delay at the PVC entry point

Figure 4. Effect of load level on end-to-end packet delay across the PVC for SRD (left) and LRD (right) traffic

The message streams that originate at the application layer are partitioned into TCP segments. The segments are channeled along the end-to-end path subject to flow and congestion control mechanisms of TCP. In other words, the TCP layer has a kind of non-linear modulating influence on the data streams passing through it. Each source tries to determine the available capacity in the network so that it can maximize the number of in-transit segments. In fact, TCP repeatedly increases the load it imposes on the network in an effort to locate the point where congestion occurs, and then it backs off from there. Thus, it creates losses to find the available bandwidth

Figure 5. Comparison of end-to-end packet delay across the PVC between high load SRD and low load LRD traffic

for the connection.⁷ In cases where TCP operates on top of links that have their own admission and flow control paradigms (e.g., Frame Relay, ATM), this can lead to situations where the two protocols react destructively resulting in performance deterioration. This is because TCP has no way of distinguishing between packet drops due to link layer control mechanism and those due to genuine congestion situations inside the network. In the present case study, the resources have been configured so as to eliminate congestion conditions happening in the network. On the other hand, the FR controls may create limited loss situations causing a throttling effect on TCP flows. To get a view of the packet delay performance that has not been impaired by the FR controls, the CCDFs of the packet delays under various load configurations are computed using *Zero CIR* settings. Fig 6 shows the behavior under SRD and LRD conditions, for similar load conditions. It is observed that the delay behavior tends to be more heavy-tailed in the case of LRD traffic as the window size is increased whereas in the case of SRD the behavior is same irrespective of the window size. It is also noted that, in the case of LRD traffic, the buffer at the PVC entry point must be increased to two to three times (for medium and high load conditions respectively) more than that of the corresponding SRD traffic conditions, so as to keep the packet loss ratio under 1%.

The effects of FR rate and burst control mechanisms on SRD and LRD traffic are presented in fig 7. In each category, five different control structures are studied in approximately similar utilization profiles. Again, the LRD traffic seems to be sensitive to controls. Even though the packet loss has been limited to below 1%, it acts indirectly in regulating delays because each loss puts the TCP source into the congestion control phase (*multiplicative decrease*) wherein the sources try to readjust to the prevailing conditions. For FR controls, the ratio $\frac{B_e}{B_c}$ seems to be an important parameter. Additionally, under high LRD contents, the absolute values of B_c and B_e are also significant in their impact on the packet delays across the two ends of the PVC.

The median, mean and 90% quantile of the end-to-end delay and loss percentages are presented in table 1. As shown in fig 7, in the case of LRD, under different FR control mechanisms and TCP window sizes, the delay figures fluctuate over a specific range. By cutting off big bursts via FR controls, the delay levels may be controlled. But, in doing so, the utilization rates also come down. The figures for delays in the table are for the FR control mechanisms where the utilization levels are not affected too much. As mentioned above, most of the delay occurs due to bandwidth mismatch between the LAN and WAN. Thus these figures could be useful as rough packet level SLA limits (irrespective of the link layer technology) where similar bandwidth ratios may exist.

5. END-USER PERFORMANCE

The perception of the user when accessing Internet services is mainly concerned with reliability and responsiveness. The main concerns in reliability aspects of Internet is with irregularities in domain name resolution. Further, the main criteria in responsiveness are: 1) page access turnaround time in WWW; and 2) throughput rates in case of SMTP and FTP. In the previous section, the dependence of packet level performance on various factors such as:

Figure 6. Effect of TCP window size on SRD (left column) and LRD (right column) traffic at different utilization levels: low (top), medium (middle), high (bottom)

1) LRD properties of the traffic; 2) utilization profiles; 3) TCP and FR layer controls; and 4) sizing of buffers was presented. The same factors also influence the end-user performance.

The dependence of access turnaround of WWW on the utilization levels is presented in fig 8. The network resources and the load profiles were dimensioned such that at least 90% of the WWW page requests complete within a deadline of 1 second. In other words, the deadline of 1 sec was selected as something equivalent to the application layer SLA. Given the simulation conditions, and under LRD conditions, this is achieved by reducing the number of clients to approximately 50% of the SRD levels. The CCDF of the access delays is shown in fig 9. The response delay seems to be rather insensitive to the utilization levels. As explained in the previous section, the main component of delay comes from the delay experienced by the packets at the WAN ingress point. This delay

Figure 7. Effect of FR PVC controls on SRD (left column) and LRD (right column) traffic at medium utilization level and different TCP window sizes: 8760 (top), 16384 (middle), 32120 (bottom)

is mainly determined by the Hurst parameter and is not so sensitive to utilization levels.

It is very important to consider the fact that big payloads will entail longer response times. A big postscript document may take a few minutes to download. Thus, it is meaningful to correlate the bytes downloaded per page to the respective response times. Accordingly, in fig 8 the relationship is plotted on a log-log scale. A typical usage pattern in WWW may involve a user periodically accessing a particular page (e.g., reading a newspaper or checking stock prices). In such a case, the user would associate a particular response profile with the specific page and thus will be rather sensitive to big variations in that. In such a case, repeatability of the response time is an important performance issue. Given a narrow range of page size in bytes, the spread of response times is observed. The spread widens as the utilization level increases, and this indicates the wide variability in service levels.

Table 1. End-to-end delay (across the two ends of the PVC) under SRD and LRD traffic environment

Traffic Class	Utilization Level	End-to-end delay across the PVC in milliseconds Median, Mean, 90% Quantile	Loss Percentage
SRD	Low	0.04, 3.5, 15.0	Highest loss of 0.24 occurring in case of TCP window size 32120 and FR control control $B_c = 204800$, $B_e = B_c$
	Medium	0.089, 5.8, 21.0	
	High	0.35, 9.5, 30.0	
LRD	Low	0.04, 9.0, 40.0	Highest loss of 0.75 occurring in case of TCP window size 32120 and FR control control $B_c = 204800$, $B_e = B_c$
	Medium	0.12, 10.0, 50.0	
	High	0.15, 15.0, 65.0

Figure 8. End-user performance for WWW service at different utilization levels: low (top left), medium (top right), and high (bottom)

The effect of changing the TCP window size on end-user performance is shown in fig 10. This figure depicts the scenario under high utilization and with *Zero CIR* FR control settings. The difference between window size 8760 and 16384 does not appear to be prominent. However, in the case of the window size of 32120, the change in variability in performance levels is apparent.

The effect of FR rate and burst controls on end user performance is shown in fig 11. The left column is for $T_c = 0.8$ bits and the right column has that set to 0.4. In both cases, the top row is for the case when $\frac{B_e}{B_c} = 2$ and that for the bottom row is 1. The effect of the ratio $\frac{B_e}{B_c}$ in controlling the variability in response times within the same page class is evident. The effect of the parameter T_c can be seen by comparing the figures from left to right within the same row. This quantity seems to have a marginal impact on end-user performance. It is also worth

Figure 9. CCDF of download times for web pages

Figure 10. WWW performance for TCP window sizes: 8760 (top left), 16384 (top right), and 32120 (bottom)

noting the increase in the population of outliers under stricter burst control regimes of FR.

6. SUMMARY

Aspects related to the impact of link and transport layer controls on Internet traffic has been reported. The influence of LRD properties at packet and application level delay performance has been studied under simulated conditions.

Figure 11. WWW performance under different FR control regimes

Starting with SRD traffic conditions, a bandwidth and buffer allocation scheme has been worked out such that a specified service level is obtained at both the packet and the application layer. Subsequently, the utilization conditions and the resource allocation schemes are adjusted for LRD traffic conditions with performance levels being maintained at those obtained under SRD scenario. The packet delay along the end-to-end path across a PVC interconnecting client and server groups has been studied when the port buffers at the routers and switches are dimensioned so as to optimize packet loss rates. The results show that major portion of the end-to-end PVC delay comes from the queueing delay at the WAN ingress points. This delay occurs due to the significant bandwidth differences between LAN and WAN link layers. The results also bring out the significant role TCP window size and Frame Relay PVC control mechanisms play in delay performance of Internet services.

In the future work we will expand the area of applicability by considering more applications (e.g., FTP) and more sophisticated topologies. An important issue is also related to validation of the simulation results. We expect finally to get to valuable results to be applied in the general framework of SLA in IP networks.

REFERENCES

1. Verma D., *Supporting Service Level Agreements on IP Networks*, Macmillan Technical Publishing, 1999.
2. Jena A.K., *Modeling and Evaluation of Internet Applications*, Licentiate Thesis, ISSN: 1101-3931, Dept. of Communications Systems, Lund Institute of Technology, Sweden, 2000.
3. Mil3 Inc., *OPNET Modeler version 6.0*, <http://www.mil3.com>, Mil3 Inc. Washington DC, USA, 1999.
4. Erramilli A., Narayan O. and Willinger W., *Experimental Queueing Analysis with Long-Range Dependent Packet Traffic*, IEEE/ACM Transactions on Networking, Vol. 4, No. 2, April 1996.
5. *Linux Man Pages*, RedHat Linux version 6.1, <http://www.redhat.com>.
6. Smith P., *Frame Relay Principles and Applications*, Addison-Wesley, 1994.
7. Stallings W., *High-Speed Networks: TCP/IP and ATM Design Principles*, Prentice Hall, 1998.
8. Kleinrock L., *Queueing Systems, Volume 2: Computer Applications*, Wiley-Interscience, New York, 1976.