Effect of Controls on Self-Similar Traffic

Parag Pruthi† and Adrian Popescu‡

†NIKSUN, 180 Tices Lane, East Brunswick, NJ 08816, USA.

‡University of Karlskrona/Ronneby
Department of Telecommunications and Mathematics
371 79 Karlskrona, Sweden

June 18, 1997

Abstract

Tremendous advances in technology have made possible Giga- and Terabit networks today. Similar advances need to be made in the management and control of these networks if these technologies are to be successfully accepted in the market place. Although years of research have been expended at designing control mechanisms necessary for fair resource allocation as well as guaranteeing Quality of Service, the discovery of the self-similar nature of traffic flows in all packet networks and services, irrespective of topology, technology or protocols leads one to wonder whether these control mechanisms are applicable in the real world. In an attempt to answer this question we have designed network simulators consisting of realistic client/server interactions over various protocol stacks and network topologies. Using this testbed we present some preliminary results which show that simple flow control mechanisms and bounded resources cannot alter the heavy-tail nature of the offered traffic. We also discuss methods by which application level models can be designed and their impacts on network performance can be studied.

1 Introduction

Recent traffic measurement studies from a wide range of working packet networks (including Ethernet LANs, WANs, CCSN/SS7, ISDN, Frame Relay and VBR video over ATM) have convincingly shown the presence of fractal or self-similar properties in both local area and wide area traffic traces [2, 5, 6, 10, 23, 28, 31], which means that similar statistical patterns may occur over different time scales that can vary by many orders of magnitude (i.e., ranging from milliseconds to minutes and even hours). Such fractal phenomenon have been observed in many other branches of science and engineering [39]. Examples of such systems are found in sociology (Lotka's law), linguistics (Zipf's law), economics (Pareto's law), medicine (cardiac spectrum) and biology (Willis' law). Despite the efforts in these other areas of science, the application of methods developed in these fields to fractal traffic modeling and analysis is relatively new. While most of the previous work in science and engineering has almost exclusively focused on statistical and practical features of fractal models (e.g., data analysis, mathematical modeling) and numerous advancements have been made, the door is wide open for research and advancements in the area of engineering and performance analysis of such systems. This is mainly because of the lack of mathematical tools to analyze such systems leading to difficulties in obtaining practically useful results.

This discovery of these features in network traffic calls to question some of the basic assumptions made by most of the research done in control, engineering, and operations of broadband integrated systems. The fact that network traffic is inherently fractal or long-range dependent (LRD) and most broadband systems have been designed based upon the assumption that traffic flows in high-speed networks are characterized by short-range dependence (SRD) leads one to question the extent to which the results of these studies are applicable in practice. At the time being, there is mounting evidence that LRD is of fundamental importance for a number of teletraffic engineering problems, such as traffic measurements [40], queueing behavior and buffer sizing [4, 9, 12, 25], admission control [8] and congestion control [17, 35, 37]. Unlike traditional packet traffic models, which give rise to exponential tail behavior in queue size distributions and typically result in optimistic performance predictions, fractal traffic models predict hyperbolic or Weibull (stretched exponential) queue size distributions. These models therefore will lead to different control and management algorithms of present and future packet networks within an integrated framework where end-to-end Quality of Service (QoS) is important.

Under the present circumstances where current data networks are designed and operated based on models which are not representative of the traffic measured from actual working data networks, our ability to manage and control broadband networks and services in “real time” when subjected to “real traffic” may be very limited. Also, practically useful management knowledge and engineering methods have played a major role in the general acceptance, rapid deployment and successful operation of new technologies and services and are expected to impact the growth of broadband services and technologies as well. If these new systems are to succeed in the market place a wealth of management knowledge must accompany the deployment of broadband technologies and services.

In this paper we will describe our preliminary attempts to investigate the effect of controls on self-similar traffic flows, and discuss implications to controllability in the context of a realistic network environment. We will begin with a discussion on the role of controls in the network context (section 2). Following that discussion we will describe some methods of modeling realistic applications and sources (section 3). These source models will be used to drive network simulators and in particular we will use a TCP/IP simulator we have developed to discuss the impacts of controls on self-similar traffic (section 4). We will end this paper with a few discussions on our observations and will scope out further work needed to gain a thorough understanding of network controls on self-similar traffic flows.

2 Role of Network Controls

One of the most important and challenging issues in the design of broadband networks is to develop an efficient and integrated framework within which requested end-to-end QoS guarantees are fully supported. Quality of Service is one of the main reasons for “controlling” and monitoring network traffic flows and is the primary focus of our study. Standardization committies have developed many recommendations and specifications they believe will result in providing guaranteed levels of service to the end user. For instance, the ATM Forum Traffic Management (TM) working group has recently completed the TM 4.0 specification [1], according to which five service categories have been identified. These are the Constant Bit Rate (CBR), the Real-time Variable Bit Rate (rt-VBR), the Non-real-time Variable Bit Rate (nrt-VBR), the Unspecified Bit Rate (UBR) and the Available Bit Rate (ABR). Many diverse (broadband) applications are expected to be supported within these five classes of services. Each of these classes of traffic imposes various interactions at various levels of the protocol hierarchy such as interactions between the window-based flow and congestion control of TCP/IP and the preventive congestion control scheme of ATM; but such demands at various levels may result in conflicting or incompatible requirements between these levels which may make service deployment onerous. Such questions on how to map application level requirements into protocol level constraints are yet to be fully investigated. Given the state of the art one wonders whether a cost effective, integrated and distributed system can be developed which can meet strict QoS criteria for a large number of end users subscribing to a variety of services.

At the present time it is fair to say that the design and management of integrated broadband networks is complicated by the presence of disparate service requirements (some of which will be known only as future applications and services are developed). Examples of applications which lead to these disparate service requirements range from voice and interactive data to image, full-motion video and bulk data [14]. Such different classes of traffic may require deterministic or statistical bounds on various QoS parameters such as throughput, cell loss, cell delay and cell delay jitter. These parameters represent performance objectives expected from the system for the duration of a connection. The problem of traffic management when integrating communication protocols from both existing and future traffic types and when required to provide bounds on QoS on a per connection basis becomes indeed a very challenging task.

How can a network provider guarantee the potentially wide spectrum of QoS requirements? The ATM Forum has concluded that a number of traffic management mechanisms are necessary to properly provide the expected levels of service from the network. Such management mechanisms include Traffic Shaping, Resource Management, Call Admission Control, Usage Parameter Control, Network Parameter Control, Priority Control, etc. The basic philosophy is that by the use of such control mechanisms the traffic can be sufficiently regulated at the edges of the network (by shaping, feedback control, blocking, etc.) so as to make the traffic flow within the network conform to behavior such that QoS guarantees can be met using very simple strategies based on traditional models.

These topics being very important for network operations and management have received a lot of attention in the literature and a number of functional control mechanisms have been proposed for ATM traffic control. There is generally a trade-off between the levels at which QoS requirements are met and efficiency with which network resources are utilized. For instance, four general categories of services can be distinguished according to this trade-off, namely services with deterministic guarantees (so-called hard guarantees) with worst-case allocation [20, 41], services with statistical guarantees (so-called soft guarantees) with probabilistic allocation [11, 22, 45], prediction-based services with measurement-based admission control [18], and services with feedback-based control [21]. Usually, networks with deterministic guarantees provide the best QoS guarantees but at the expense of lower efficiency in utilization of network resources whereas the other three approaches offer the possibility of a better resource utilization at the cost of a potential QoS degradation.

Most studies on broadband network management have been entirely based on traditional modeling assumptions. To investigate issues of the effectiveness of network controls on fractal features we begin in the next section by modeling self-similar traffic at various levels. We will then use these traffic generators to drive network simulators and compare results to gain insights into the effectiveness of controls on "shaping" fractal traffic.

3 Application Level Modeling

Today, it has been established that traffic flows and patterns for resource usage in packet and cell based networks have fractal properties which implys that the above-mentioned control mechanisms need to be revisited and reassessed under the assumption of realistic traffic. We begin by first discussing various features of "realistic" traffic modeling which have been proposed in the literature. We will then model traffic sources around some key features and use these to drive network simulators.

Recent studies have shown that aggregations of heavy-tailed ON/OFF sources can lead to self-similar traffic flows; i.e. aggregation of many individual, albeit higly variable, on-off sources (on- and off-periods with heavy-tailed distributions with infinite variance, e.g., Pareto distributions) can lead to self-similar traffic flows [29, 38, 43]. Others have questioned this simple model under *realistic* client/server network environments where one has bounded resources; and they desire to incorporate a combination of many other details and characteristics. Although it is true that network traffic is governed by many physical factors, a "good" model should incorporate only those feature which are relevant for the problem under consideration. Below we will list some of the many factors affecting network traffic flows:

- User behaviour. One of the most important factors that affects the traffic behaviour at a session/call level and also during the session has been linked to (human) user behaviour. It was found, for instance, that the distribution of user requests (think time) and preferences for documents on Internet (WWW) show an extreme degree of fluctuations over a wide range of time scales [6, 7].

- Data generation, organization and retrieval. Most measurement studies reported so far analyze network traffic at the link level only (Ethernet LANs, WANs, CCSN/SS7, ISDN, ATM) and show that both the packet arrival process as well as the byte arrival process exhibit long range dependence [42]. Evidence of LRD in the distribution of (session) interarrival times and/or file/packet sizes has been noted for some specific applications like VBR video MPEG-coded sources (these asymptotically self-similar processes show LRD as well as strong SRD characteristics) and data sources (exactly self-similar processes, show essentially similar correlation characteristics over all but for the shortest time scales) [12, 15]. Willinger has stated long ago that file size distributions as well as jobs on a unix workstation show heavy tails, and it has recently been shown that the distributions of document sizes in WWW, of Unix filesystem, of data bytes in FTP bursts, of packet interarrival times within Rlogin and Telnet sessions, and of sizes of reads and writes to an NFS server all possess LRD [2, 3, 6, 16, 23, 36]. These studies show that the key statistical behaviour (e.g., correlation structure) in packet networks is invariant, namely heavy-tailed, from one machine/network to another.

- Traffic aggregation. One of the most serious obstacles that prevents from keeping the traffic generated by individual sources isolated from other traffic (at least for as far into the network as possible) is due to the actual form of traffic aggregation (i.e., statistical multiplexing), used in the packet- and cell-based networks, where the basic contention for resources causes the traffic sources to interact

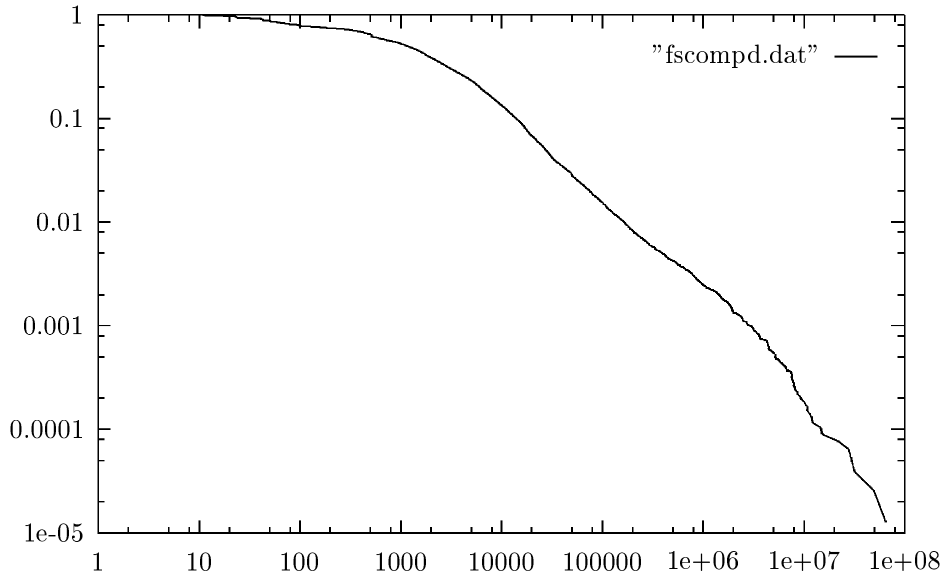

Figure 1: Typical complementary distribution of the file sizes observed in FTP and WWW servers. Note the log-log plot.

with each other. The superposition of many on-off sources exhibiting the infinite variance syndrome (Noah Effect) results in self-similar aggregate network traffic approaching fractional Brownian motion (Joseph Effect) [43]. The fractal characteristics in traffic are also shown to be highly robust with respect to a variety of network operations, such as traffic splitting, merging, queueing, policing and shaping [12, 13, 19, 26]. Self-similarity seems to be preserved under superposition of both homogeneous and heterogeneous, i.e., independent, traffic sources, and this holds over a wide range of conditions like for instance variations in the bandwidth bottleneck and buffer capacity, and mixing with cross-traffic possessing dissimilar (correlation) characteristics [32, 33]. The coupling process under realistic scenarios is a very complex one, and further work is necessary to understand whether the alteration in the short range correlations from such couplings leads to any significant effect on system operations [12].

- Network Controls. It has been speculated that by using sophisticated control mechanisms or different mechanisms for transport that the long range dependence in traffic flows can be "smoothed" out (for one example see [32]) or at the very least altered to an extent where existing methodology can be applied to the management of networks. As has been mentioned in the previous paragraph, fractal features in network traffic are robust and it would be interesting to study the extent to which controls can alter LRD traffic flows. In fact by using such mechanisms one is just trading one problem for another. This can be illustrated by the following example: suppose that in order to get around fractal traffic flows in packet networks one decides to use circuit switched networks; but in so doing one now has to solve the "long holding time problem" in circuit switching instead of the "fractal traffic problem" in packet switching.

- Network Evolution. Over time, networks are constantly being "rewired" and "upgraded" [27]. But one has to only visit Web sites which host traffic data from these networks to notice that the inherent characteristics of traffic flow on the networks before and after the "upgrade" are essentially the same.

All these above factors affect traffic at different time scales, but which features are most important to model? The answer to this question depends upon the problem being solved. There are essentially two levels of models which can prove useful. First, models at the application layer can be used to generate client/server traffic. These models can be useful for studying the effects of technologies and protocols on application layer performance. On the other hand if the objective is to study a switching system then a aggregate traffic model on a trunk will be useful. We will generate both kinds of traffic and show how

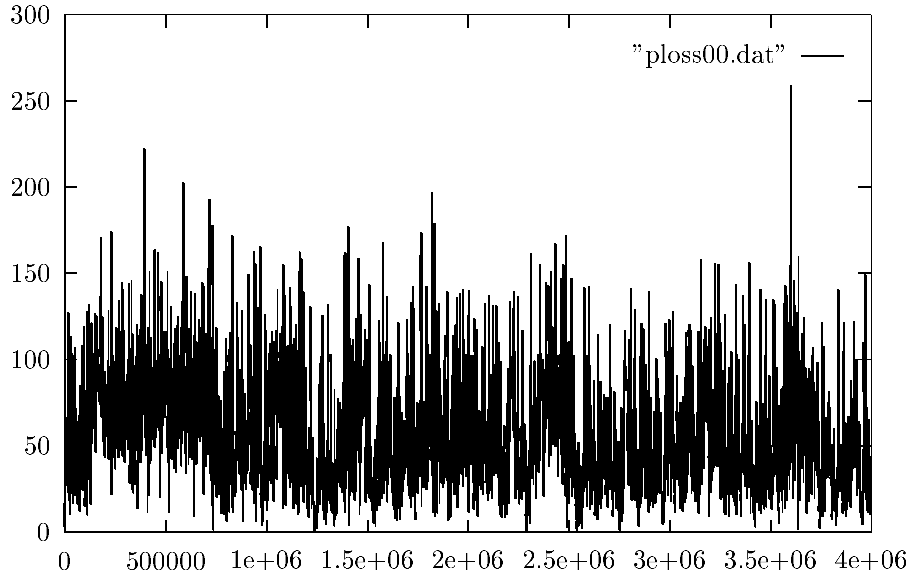

Figure 2: The aggregate server to client traffic over a lossless link in a simulation with approx. 300 clients and servers. The traffic shown is at the link layer and is being "shaped" by TCP flow control. The vertical axis shows the number of Kbytes aggregated over 1 second intervals (horizontal axis)

application traffic over realistic systems can lead to self-similar traffic on trunks. We will also show how one can directly model this aggregate trunk traffic and use these models to study switch performance.

The most natural model by which the variability over orders of magnitude (from ms. to secs.) can be easily generated is a heavy-tailed one. Heavy tail distributions where the parameter $\alpha < 2$ will naturally lead to infinite variance. Thus, we choose to model the variability in both the bytes process or the packets process (i.e. number of bytes or packets per interval (sec.)) by a Pareto distribution. Specifically we use chaotic maps to generate this process, and details on these maps can be found in [29]. We model human interactions (at time scales of secs. to mins.) by another pareto distribution but with $\alpha > 2$; using these values of α allows us to model finite as well as infinite variance processes with a single parameter.

Figure 1 shows the typical complementary distributions of file sizes observed in FTP and WWW servers [6, 28] (this data can typically be obtained in logs at the server and are typically available on many sites such as NCSA, IT.KTH.SE, etc. and may be available on your favorite site also. One can also write very simple code to download all file size data from a WWW server using a language like perl and HTTP commands). Figure 1 shows that there is an initial almost uniform distribution followed by a heavy tail. The heaviness of the tail typically implies that file sizes are highly variable and that models must be developed which capture the initial flat distribution followed by a heavy tail. We use a simple combination of light- and heavy-tails to model the qualitative features shown in Figure 1. The parameters of the model can be used to match empirical distributions. Toward that end we have studied many FTP/WWW transactions from published data and have parametrized our models to be consistent with these studies [44, 6]. For a survey of many different methods of generating heavy-tail distributions please refer to [30]. Using these methods we were able to model FTP and WWW file size distributions. We then built a simulation model of FTP and WWW client/server interactions using the FTP protocol (*RFC 959*) and HTTP 1.0 and 1.1 proposals (*RFC 1945, RFC 2068*). We also used the aggregation of ON/OFF sources to model trunk traffic [30].

Armed with application and trunk traffic generators, we performed simple simulation experiments described in the next section and studied the effect of TCP on fractal traffic flows.

4 A Simple Experiment

Using the traffic models developed in the previous section we decided to simulate client/server traffic over a TCP/IP network. We implemented TCP/IP (Tahoe with Fast Retransmit) in our simulator and

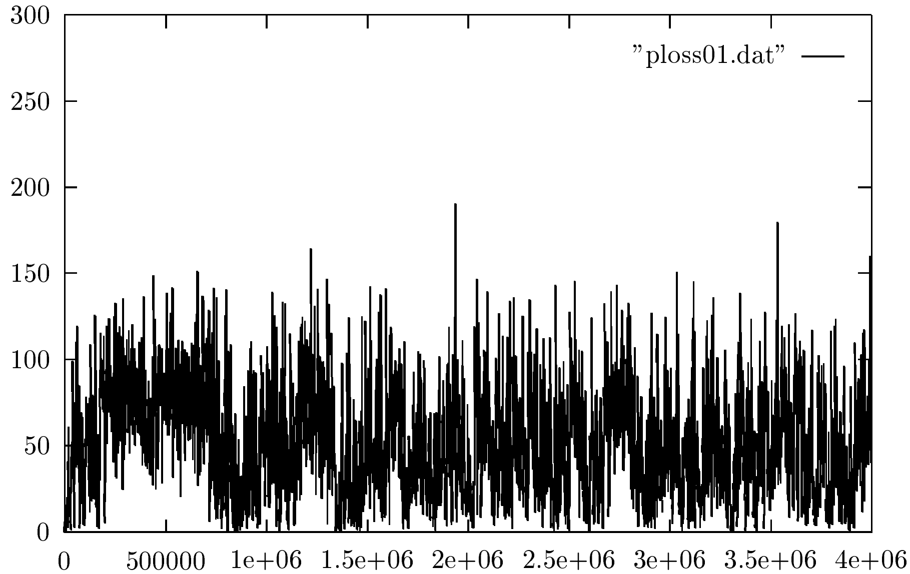

Figure 3: The aggregate server to client traffic over a lossy link (1% %losses) in a simulation with approx. 300 clients and servers. The traffic shown is at the link layer and is being "shaped" by TCP flow control. The vertical axis shows the number of Kbytes aggregated over 1 second intervals (horizontal axis)

decided to check the impacts of TCP flow control on self-similar traffic. We decided to use TCP/IP because it is one of the most common protocols being used in many different networks today. We also modeled correlated network losses using hyperbolic distributions as shown by Mandelbrot in a very early paper [24].

The FTP clients were designed to set up connections with the server and either upload or download multiple files to/from the server. The file size distributions were matched with those reported in numerous studies [28, 6]. The number of uploads/downloads per session were modeled as pareto distributions with $\alpha > 2$. The WWW clients always requested a html document from a server; this document could consist of many objects (image, audio, video, etc.) each of which were fetched as per http 1.0 specifications. User think times were modeled as pareto distributions and initially for these experiments we left the think times to have finite variance and mean (because of lack of any real data we opted not to model this in any depth). Our basic goal was to model a "busy period" and therefore we tried to keep think times short so that a significant number of clients would be active at any given time.

Having the basic framework set up we decided to conduct a few experiments. Figure 2 shows the trace monitored at the network trunk from our simulation with approximately 300 traffic sources consisting of FTP and WWW clients. For this experiment the loss rate on the trunk was kept to zero; i.e. there are no losses in the system - it represents a pure delay system yet TCP flow control and timeouts were occuring due to large delays during periods of congestion. The data shown here was monitored from the servers to the clients over a single trunk carrying this traffic. The vertical axis shows the number of Kbytes and the horizontal axis shows the number of seconds; i.e. the graph shows the number of Kbytes on the trunk per second. Note how similar this data looks to real traffic traces from actual networks. Also note the existence of low frequency components in this trace showing the presence of $1/f$-noise. Typical tests for self-similarity (variance-time plots, periodograms, and Whittle's point estimates) all confirmed the persence of LRD in these traces. We did not observe any significant change in the Hurst parameter from that of the application level sources to that observed from these trunk level traces; what we did notice was the introduction of some low-frequency components which do not appear to be present in the traffic offered from the clients and servers. Due to the preliminary nature of these expirements more work is needed to draw conclusions which would show any significant differences in the offered traffic from that carried on the trunk.

Figure 3 shows a similar plot when we introduced 1% losses on the trunk. This simulation as well as the previous one was initialized with the same random seed, hence leading to almost identical client/server uploads and downloads. But due to the losses there were a number of TCP retransmissions along with

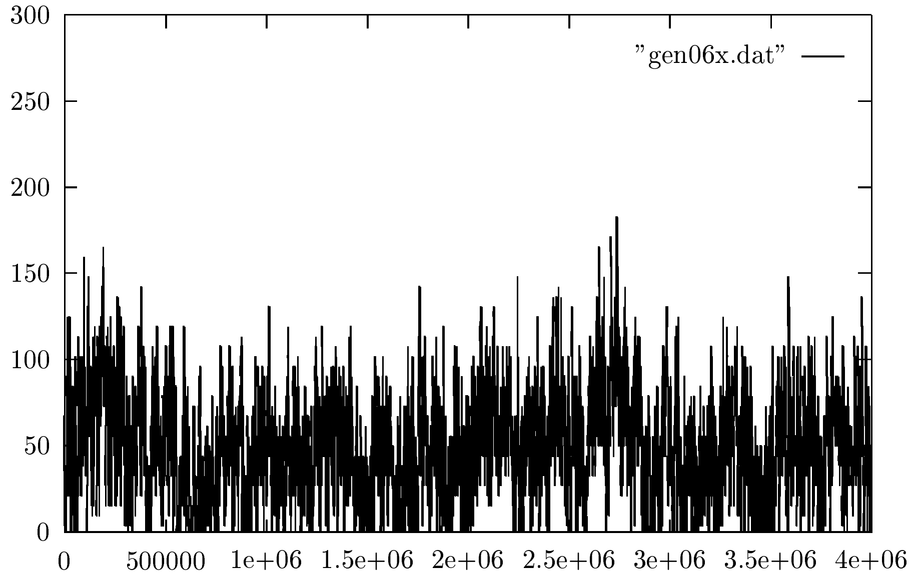

Figure 4: Aggregates of 150 ON/OFF sources whose parameters were matched from the trace shown in Figure 3. Note the good qualitative match in the mean rate, the fluctuations about the mean and the low-frequency component.

timeouts during periods of congestion. Note the differences in the two traces which show that the mean is lower in Figure 3 and the peak fluctuations of the process in Figure 3 is slightly lower than in Figure 2. A careful analysis of the trunk traffic shows that the mean rate over the link has dropped from about 55 Kbytes/sec to around 50 Kbytes/sec; but we were unable to determine any significant changes to the Hurst parameter. We also intend to investigate the effect of losses on the peakedness parameter a. Thus, it appears that a combination of network controls and losses do produce generally expected results but further work is necessary to determine the effect on the correlation structure. Note that it has been shown theoretically [12] that such controls cannot change the distribution in the tail; our emphasis is more on how the bulk of the distribution is affected and what impacts it has on network performance.

Having monitored the traffic on the trunks we decided to model the aggregate traffic shown in Figure 3 with aggregates of ON/OFF models. We matched the parameters of the trace shown in Figure 3 to the parameters of the ON/OFF models (mean, peakedness and Hurst parameter) and used aggregates of 150 ON/OFF sources [29] to see how well simiple ON/OFF models (without regard to client/server interactions) matches a trace which has been shaped by client/server interactions, TCP/IP and network losses. The result is shown in Figure 4. What is remarkable is that a simple ON/OFF model captures most of the features of the trace shown in Figure 3. The only slight differences can be found upon close examination where there is higher high-frequency component in Figure 3 than in Figure 4. Both these facts should come as no surprise because we start with applications characterized by heavy-tail distributions and the multiplexed and "controlled" traffic is known to be self-similar from link-level measurements over real working packet networks. What we have shown here is that depending upon the problem at hand there may be no need to further complicate the traffic generation by modeling endless details; simple aggregations of ON/OFF sources with appropriate parameters suffice to capture many of the features found in actual traffic traces. This is not to say that application level models are not important, because having application level models allows us to study application level performance in the real network context.

5 Conclusions

Using the framework developed here we have shown that real world protocols and impairements do modify the structure of traffic offered to the network entities, but these controls cannot alter the tail of the distributions of the traffic flow processes. At best one can only trade problems, i.e. by incorporating shapers and policers at the edges of the networks one needs huge buffers which results in large delays which may be unacceptable in practice. The effects of controls on network traffic need to be investigated in more detail in order to quantify these tradeoffs. The claims by various people in various forums that fractal traffic will be "smoothed" at the edges of the network and that flows within the network will not exhibit fractal features needs careful reconsideration. Our results indicate that the problem may not be so easily dismissed and that indeed many of the mechanisms currently being proposed for guaranteeing QoS will need to be revisited. We plan to continue to investigate the effectiveness of various schemes in light of these results. The work here should not be regarded as suggesting that controls are not useful, but rather that controls should be used for enforcing fairness and their limitations in real operations should be well understood.

References

[1] ATM Forum Traffic Management Specification, Version 4.0, 1996.

[2] Addie, R., Zukerman, M. and Neame, T.,"Fractal Traffic: Measurements, Modelling and Performance Evaluation," in Proceedings of *IEEE INFOCOM'95,* pp. 977 - 984, 1995.

[3] Bodnarchuk, R.R. and Bund, R.B.,"A Synthetic Workload Model for a Distributed System File Server," in Proceedings of *ACM SIGMETRICS'91,* pp. 50 - 59, 1991.

[4] Brichet, F., Roberts, J., Simonian, A. and Veitch, D.,"Heavy Traffic Analysis of a Storage Model with Long Range Dependent On/Off Sources," preprint, CNET, 1996.

[5] Beran, J., Sherman, R., Taqqu, M.S. and Willinger, W.,"Long-Range Dependence in Variable-Bit Rate Video Traffic," *IEEE Transactions on Communications,* Vol. 43, No. 2/3/4, pp. 1566 - 1579, February/March/April 1995.

[6] Crovella, M.E. and Bestavros, A.,"Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes," in Proceedings of *ACM SIGMETRICS'96,* pp. 160 - 169, 1996.

[7] Crovella, M.E. and Bestavros, A.,"Explaining World Wide Web Traffic Self-Similarity," preprint, Boston University, Computer Science Dept., 1995.

[8] Causey, J.W. and Kim, H.S.,"Comparison of Admission Control Schemes in ATM Networks," *International Journal of Communication Systems,* Vol. 8, pp. 165 - 184, 1995.

[9] Duffield, N.G. and O'Connell, N.,"Large Deviation and Overflow Probabilities for the General Single-Server Queue, with Applications," preprint, Dublin Institute for Advanced Studies, 1993.

[10] Duffy, D.E., McIntosh, A.A., Rosenstein, M. and Willinger, W.,"Statistical Analysis of CCSN/SS7 Traffic Data from Working Subnetworks," *IEEE Journal on Selected Areas in Communications,* Vol. 12, No. 3, pp. 544 - 551, 1994.

[11] Elwalid, A. and Mitra, D.,"Analysis, Approximations and Admission Control of a Multi-Service Multiplexing System with Priorities," in Proceedings of *IEEE INFOCOM'95,* Boston, 1995.

[12] 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, pp. 209 - 223, April 1996.

[13] Fan, Y. and Georganas, N.D.,"On Merging and Splitting of Self-Similar Traffic in High-Speed Networks," preprint, University of Ottawa, 1995.

[14] Garrett, M.W.,"A Service Architecture for ATM: From Applications to Scheduling," *IEEE Network,* Vol. 10, No. 3, pp. 8 - 14, May/June 1996.

[15] Garrett, M.W. and Willinger, W.,"Analysis, Modeling and Generation of Self-Similar VBR Video Traffic," in Proceedings of *ACM SIGCOMM'94,* London, pp. 269 - 280, 1994.

[16] Irlam, G.,"Unix File Size Survey - 1993," available at *http://www.base.com/gordoni/ufs93.html,* September 1994.

[17] Jain, R.,"Congestion Control and Traffic Management in ATM Networks: Recent Advances and A Survey," *Computer Networks and ISDN Systems,* November 1996.

[18] Jamin, S., Danzig, P., Shenker, S. and Zhang, L.,"A Measurement-Based Admission Control Algorithm for Integrated Services Packet Networks," in Proceedings of *ACM SIGCOMM'95,* Boston, 1995.

[19] Jena, A.K., Popescu, A., Pruthi, P. and Erramilli, A.,"Traffic Control in ATM Networks: Engineering Impacts of Realistic Traffic Processes," in Proceedings of *Nordisk Teletrafik Seminarium NTS-13,* Trondheim, Norway, August 1996.

[20] Knightly, E.W.,"H-BIND: A New Approach to Providing Performance Guarantees to VBR Traffic," in Proceedings of *IEEE INFOCOM'96,* San Francisco, 1996.

[21] Kanakia, H., Mishra, P. and Reibman, A.,"An Adaptive Congestion Control Scheme for Real-Time Packet Video Transport," in Proceedings of *ACM SIGCOMM'94,* San Francisco, 1994.

[22] Kesidis, G., Walrand, J. and Chang, C-S.,"Effective Bandwidths for Multiclass Markov Fluids and Other ATM Sources," *IEEE/ACM Transactions on Networking,* August 1993.

[23] Leland, W.E., Taqqu, M.S., Willinger, W. and Wilson, D.V.,"On the Self-Similar Nature of Ethernet Traffic (Extended Version)," *IEEE/ACM Transactions on Networking,* Vol. 2, No. 1, pp. 1 - 15, February 1994.

[24] Mandelbrot, B., "Self-Similar Error Clusters in Communication Systems and the Concept of Conditional Stationarity," IEEE Transactions on Communication Technology, 13, 71-90.

[25] Norros, I.,"A Storage Model with Self-Similar Input," *Queueing Systems Theory and Applications,* Vol. 16, No. 3 - 4, pp. 387 - 396, 1994.

[26] Neidhardt, A.L. and Erramilli, A.,"Shaping and Policing of Fractal Traffic," in Proceedings of *ITC Specialist Seminar '96,* Lund, Sweden, 1996.

[27] Paxson, V., "Growth Trends in Wide-Area TCP Connections," IEEE Network, 8(4), pp. 8-17, July/August 1994.

[28] Paxson, V. and Floyd, S.,"Wide Area Traffic: The Failure of Poisson Modeling," *IEEE/ACM Transactions on Networking,* Vol. 3, No. 3, pp. 226 - 244, June 1995.

[29] Pruthi, P. "An Application of Chaotic Maps to Packet Traffic Modeling," Ph.D. Thesis, Royal Institute of Technology, Stockholm, Sweden, October 1995. Source Behaviour and Self-Similar Traffic," in Proceedings of *ICC'95,* Seattle, pp. 445 - 450, 1995.

[30] A. Erramilli, P. Pruthi, and W. Willinger, "Recent Developments in Fractal Traffic Modeling," Proc. of ITC Regional Seminar, St. Petersburg, 26 June-2 July, 1995.

[31] Pruthi, P. and Devadhar, S.,"A Description of Traffic Flows from Public Frame Relay Networks," presented at the ITC Hot Topic Session on Traffic Characterization, *ITC'97,* Washington, DC., 1997.

[32] Park, K., Kim, G.T. and Crovella, M.E.,"On the Relationship Between File Sizes, Transport Protocols, and Self-Similar Network Traffic," preprint, Boston University, Computer Science Dept., 1996.

[33] Park, K., Kim, G.T. and Crovella, M.E.,"The Effects of Traffic Self-Similarity on TCP Performance," preprint, Boston University, Computer Science Dept., 1996.

[34] Parulekar, M. and Makowski, A.,"A Buffer Overflow Probabilities for a Multiplexer with Self-Similar Traffic," in Proceedings of *IEEE INFOCOM'96,* San Francisco, 1996.

[35] Roberts, J.W.,"What Traffic Handling Capabilities for the B-ISDN?," in Proceedings of *9th ITC Specialists Seminar,* the Netherlands, November 1995.

[36] Ryu, B.K.,"Implications of Self-Similarity for Providing QOS Guarantees End-to-End in High Speed Networks: A Framework of Application Level Traffic Modeling," in Proceedings of *International Zurich Symposium'96,* 1996.

[37] Stallings, W.,"Viewpoint: Self-Similarity Upsets Data Traffic Assumptions," *IEEE Spectrum,* January 1997.

[38] Taqqu, M.S., Willinger, W. and Sherman, R.,"Proof of a Fundamental Result in Self-Similar Traffic Modeling," to appear in *Computer Communications Review,* 1997.

[39] West, B.J. and Deering, B.,"The Lure of Modern Science Fractal Thinking," Studies of Nonlinear Phenomena in Life Sciences, No. 3, World Scientific Publishing Co., 1995.

[40] Wang, J.L. and Erramilli, A.,"Monitoring Packet Traffic Levels," in Proceedings of *GLOBECOM'94,* pp. 274 - 280, 1994.

[41] Wrege, D.E., Knightly, E.W., Zhang, H. and Liebeherr, J.,"Deterministic Delay Bounds for VBR Video in Packet Switching Networks: Fundamental Limits and Practical Tradeoffs," *IEEE/ACM Transactions on Networking,* Vol. 4, No. 3, pp. 352 - 362, June 1996.

[42] Willinger, W., Taqqu, M.S. and Erramilli, A.,"A Bibliographical Guide to Self-Similar Traffic and Performance Modeling for Modern High-Speed Networks," *Stochastic Networks: Theory and Applications,* Oxford University Press, pp. 339 - 366, 1996.

[43] Willinger, W., Taqqu, M.S., Sherman, R. and Wilson, D.V.,"Self-Similarity Through High-Variability: Statistical Analysis of Ethernet LAN Traffic at the Source Level," *IEEE/ACM Transactions on Networking,* Vol. 5, No. 1, pp. 71 - 86, February 1997.

[44] W. Willinger, V. Paxson and M.S. Taqqu, "Self-Similarity and Heavy Tails: Structural Modeling of Network Traffic," to appear in A Practical Guide to Heavy Tails: Statistical Techniques for Analyzing Heavy Tail Distributions, R. Adler, R.Feldman and M.S., Taqqu (Editors), Birkhauser Verlag, Boston, 1997.

[45] Zhang, H. and Knightly, E.,"Providing End-to-End Statistical Performance Guarantees with Bounding Interval Dependent Stochastic Models," in Proceedings of *ACM SIGMETRICS'94,* May 1994.