

Performance Analysis of ATM ABR Service Under Self-Similar Traffic in the Presence of Background VBR Traffic

George Benke†‡, Jim Brandt††, Helen Chen††, Siamak Dastangoo†,
Gregory J. Miller†

†The MITRE Corporation
7527 Colshire Drive
McLean, VA 22102

††Sandia National Labs
P.O. Box 969
Livermore, CA 94551

‡Georgetown University
Department of Mathematics
Washington, DC 20057

RECEIVED

APR 18 1996

STI

{hycsw, brandt}@ca.sandia.gov, {gbenke, sdastang, gmiller}@mitre.org

Abstract — Recent empirical studies of traffic measurements of packet switched networks have demonstrated that actual network traffic is self-similar, or long range dependent, in nature. That is, the measured traffic is bursty over a wide range of time intervals. Furthermore, the emergence of high-speed network backbones demands the study of accurate models of aggregated traffic to assess network performance. This paper provides a method for generation of self-similar traffic, which can be used to drive network simulation models. We present the results of a simulation study of a two-node ATM network configuration that supports the ATM Forum's Available Bit Rate (ABR) service. In our study, we compare the state of the queue at the source router at the edge of the ATM network under both Poisson and self-similar traffic loading. Our findings indicate an order of magnitude increase in queue length for self-similar traffic loading as compared to Poisson loading. Moreover, when background VBR traffic is present, self-similar ABR traffic causes more congestion at the ATM switches than does Poisson traffic.

1. Introduction

It has been observed that Ethernet LAN traffic, WAN traffic, and coded video traffic are statistically *self-similar* [1], [2], and [3]. Measured Ethernet traffic exhibits *fractal* behavior, which is not captured by conventional traffic models (e.g., BMAP, MMPP, and Poisson). The self-similar nature of aggregated LAN traffic may have a significant impact on the design, control, and analysis of the high-speed, cell-based backbones by which LANs are expected to be inter-connected. Self-similar traffic may have an adverse effect on network performance since its burstiness and long-term correlation may stress the network resources such as memory buffers and link capacity. In fact, it has been shown that burstiness decreases much more slowly when aggregating self-similar traffic streams as opposed to aggregating conventional traffic streams [1]. This persistent burstiness can potentially cause poorer performance than is predicted by conventional traffic models. To quantify the effects of self-similarity, the use of self-similar traffic models should be explored in conducting network performance analyses.

We report here on one such network performance analysis, in which we study the impact of self-similar traffic on the ATM ABR service. Asynchronous Transfer Mode (ATM) is an emerging high-speed, cell-based transport technology envisioned to support integrated services over both WAN and LAN environments. The ATM Forum has defined four service categories for ATM: Constant Bit Rate (CBR), Variable Bit Rate (VBR), Unspecified Bit Rate (UBR), and Available Bit Rate (ABR). The CBR service category is used by sources that require a static amount of bandwidth that is continuously available during the connection lifetime. The VBR service category is intended for real-time applications that require tightly constrained delay and delay variation guarantees. The UBR

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DISCLAIMER

**Portions of this document may be illegible
in electronic image products. Images are
produced from the best available original
document.**

and ABR service categories are both intended for non-real-time applications that do not have strict delay requirements

The ABR service is designed to provide the service subscribers with rapid access to unused network bandwidth while maintaining a low Cell Loss Rate (CLR). The ATM Forum has adopted a rate-based flow control protocol [4] that uses a closed-loop feedback mechanism to allow a traffic source to adapt to dynamically changing network conditions (see Figure 1). To receive network feedback, a source must periodically probe the network by sending resource management (RM) cells, which are reflected to the source by the destination. Switches along the path indicate the current worst-case network condition by marking these RM cells with relevant information. While a simple switch would do nothing more than set the congestion-bit in the ATM header of passing cells [5], a more elaborate switch could mark the RM cell with explicit supportable rates that are calculated using per-VC accounting information. Based on the feedback information, traffic sources calculate new sending rates using a multiplicative-decrease and additive-increase algorithm. All new sources must start with an Initial Cell Rate (ICR), which is negotiated during connection setup and is usually small in comparison to the available link rate. This start-up mechanism is designed to prevent switch queue overflow when multiple sessions start sending simultaneously.

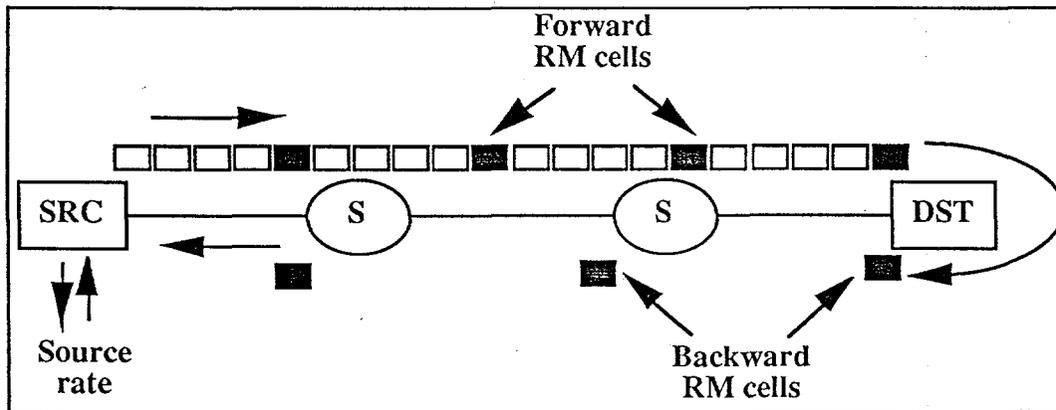


Figure 1. ABR Rate-Based Flow-Control Scheme

The purpose of our simulation study is to observe the queue behavior at the ATM ABR Service Access Point (SAP) using both a simple traffic model and an accurate model, which reflects the self-similar properties of measured traffic. The traffic source simulates as an IP router where multiple Ethernet connections are multiplexed, see Figure 2. It has been shown that large queues can develop in the simplest models of constant-rate servers when the traffic arrival process is self-similar [6]. Large queues translate directly into high delay and possibly buffer overflow if memory buffers are not sized appropriately. The ABR SAP has complex service characteristics because service times depend both on the local demand and on external factors which constrain the resources available in the ATM network. Moreover, due to the source-destination round-trip time there is an unavoidable lag or hysteresis in the network response to increased demand. Queueing characteristics at an ABR source for self-similar traffic arrival processes have not yet been sufficiently investigated.

In this paper, we define a self-similar traffic model, describe a method for generating self-similar traffic, and present a simulation study of an ATM network with self-similar arrival traffic at the ABR SAP under the presence of VBR background traffic. The next sections discuss the concept of self-similarity, the self-similar traffic model, and the traffic

generation method. Section 4 describes the simulation study, and section 5 summarizes the preliminary results and analysis. Finally, section 6 summarizes our results and suggests plans for future studies.

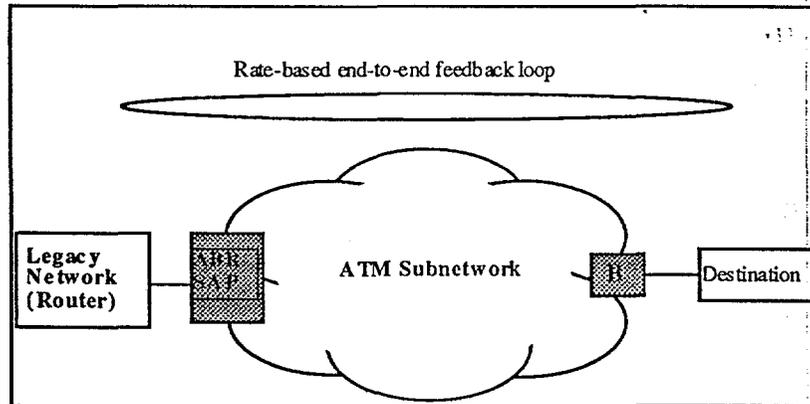


Figure 2. Closed-Loop Flow-Controlled ATM Network.

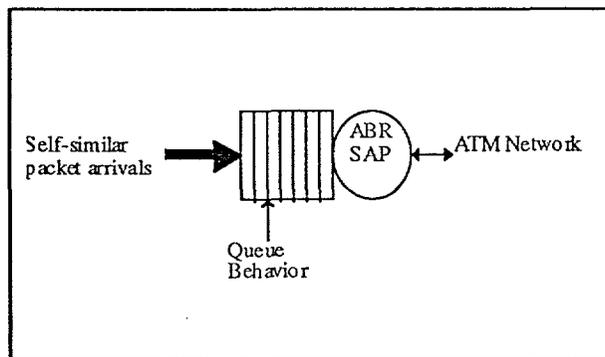


Figure 3. An ABR Queue.

2. A Self-similar Traffic Model

Recent measurements of data traffic on LANs [1] and WANs [7] have shown the existence of significant long-range dependence among arrivals, a characteristic not found in standard Poisson traffic models, see Figure 4. Consequently, self-similar processes have recently become the subject of study in an effort to model observed network traffic more accurately. In the remainder of this section, we define self-similarity and describe a self-similar traffic model.

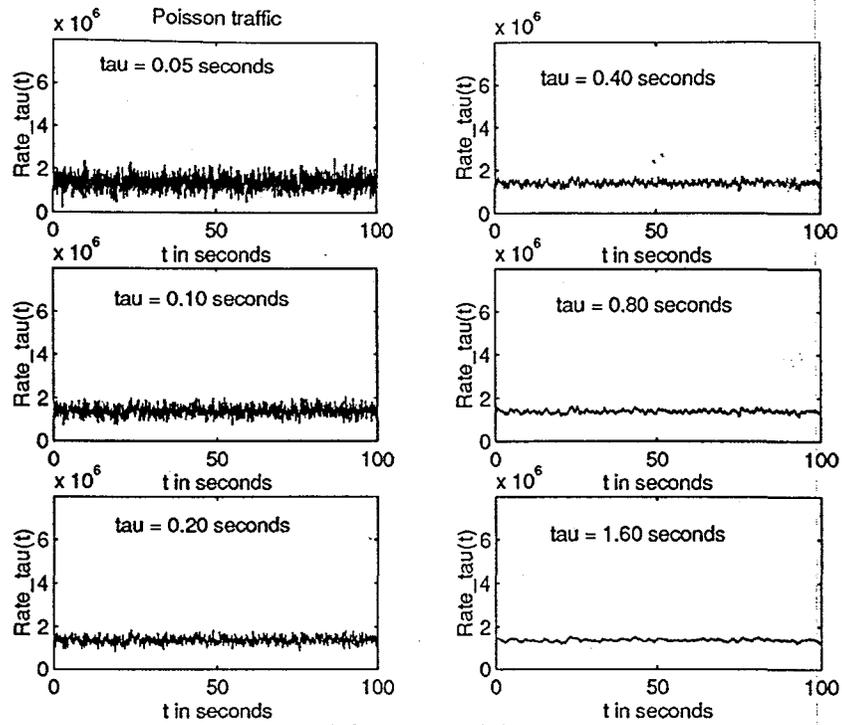


Figure 4 (a).

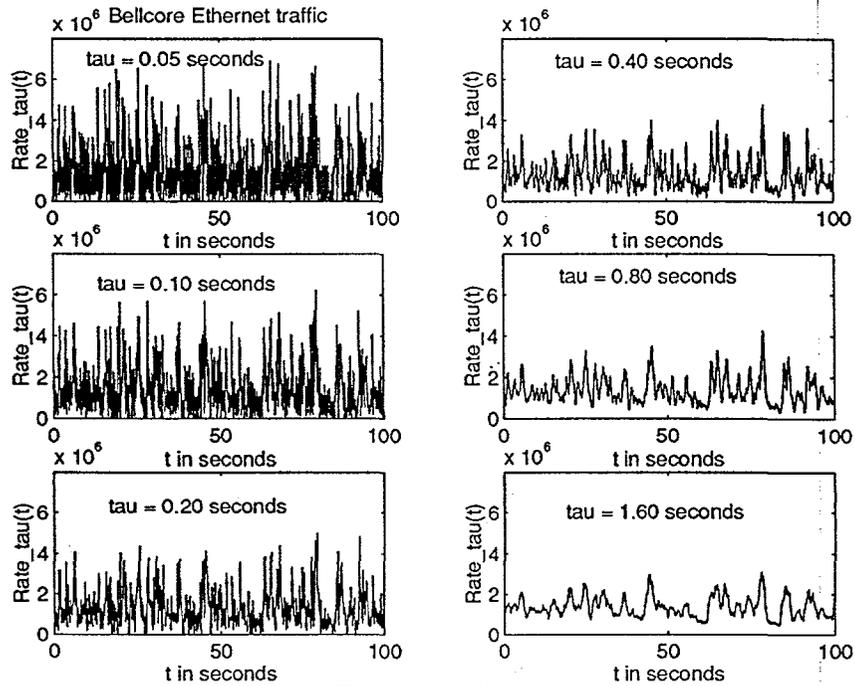


Figure 4 (b).

Figure 4. Burstiness Over Different Time Scales: (a) Conventional (Poisson) Traffic and (b) Measured Bellcore Ethernet Traffic.

A stochastic process $Q(t)$ for $t \geq 0$, and $Q(0) = 0$, is said to be self-similar if for any $c > 0$, and any t_1, \dots, t_n , the joint distribution of $Q(ct_1), \dots, Q(ct_n)$ equals the joint distribution of $c^H Q(t_1), \dots, c^H Q(t_n)$ where H is some constant, called the Hurst parameter of the process. The process $Q(t)$ has stationary increments if the distribution of $Q(t + \tau) - Q(t)$ depends only on τ . The following facts about self-similar process with stationary increments can be established:

- (1) If $Q(t)$ is self-similar, with $H \neq 1$, and $E(Q(1)) < \infty$, then $E(Q(t)) = 0$, for $t \geq 0$.
- (2) $Q(t)$ has the same distribution as $t^H Q(1)$.
- (3) $Q(t + \tau) - Q(t)$ has the same distribution as $\tau^H Q(1)$.
- (4) If $V(Q(1)) < \infty$, then $V(Q(t)) = t^{2H} V(Q(1))$.
- (5) If $V(Q(1)) < \infty$, then $V(Q(t + \tau) - Q(t)) = \tau^{2H} V(Q(1))$.
- (6) If $V(Q(1)) < \infty$, then the autocovariance of $Q(t + \tau) - Q(t)$ is

$$\begin{aligned} \text{Cov}(Q(t_1 + \tau) - Q(t_1), Q(t_2 + \tau) - Q(t_2)) = \\ \frac{V(Q(1))}{2} [(t_2 - t_1 + \tau)^{2H} - 2(t_2 - t_1)^{2H} + (t_2 - t_1 - \tau)^{2H}] \end{aligned}$$

In order to relate self-similar processes to traffic models, we establish the following notation: $N(t)$ = total number of bits put on the network in the time interval $[0, t]$. If $N(t)$ has stationary increments, and $E(N(1)) < \infty$, then $E(N(t)) = E(N(1))t$. Denoting $E(N(1))$ by λ we have

$$E(N(t)) = \lambda t.$$

The constant λ is the traffic rate measured in bits per second if t is measured in seconds. Note that strictly speaking, $N(t)$ assumes only non-negative integer values and is non-decreasing. In a self-similar traffic model we assume that the process $Y(t) = N(t) - \lambda t$ is self-similar according to the earlier definition. Therefore by (2)

$$N(t) = Y(t) + \lambda t$$

has the same distribution as $t^H Y(1) + \lambda t$ which has, except for a dilation and shift, the same distribution as $Y(1)$. In self-similar traffic modeling, a popular choice for the distribution of $Y(t)$ is Gaussian, which implies that the increment process, $N(t + \tau) - N(t)$, also follows a Gaussian distribution. We note that this choice violates the non-decreasing requirements for $N(t)$.

This can be seen as follows,

$$N(t + \tau) - N(t) = Y(t + \tau) - Y(t) + \lambda \tau$$

which by (3) has the same distribution as the Gaussian random variable $\tau^H Y(1) + \lambda \tau$.

Since any Gaussian random variable assumes negative values with positive probability, $N(t+\tau) - N(t)$ assumes negative values which means that $N(t)$ is not strictly non-decreasing. Histogram studies of the Bellcore Ethernet traffic [1], an extensively studied collection of traffic measurements, shows that the distribution of $N(t+\tau) - N(t)$ for large τ is close to Gaussian, but for small τ departs greatly from a Gaussian shape, and, in fact, is closer in shape to an exponential curve. See Figure 5, which shows the distribution of $N(t+\tau) - N(t)$ for a large τ (4 seconds) and a small τ (0.05 seconds). On the other hand, even at small scales, the second order statistics of measured Ethernet traffic are in remarkably close agreement with the second order statistics of self-similar finite-variance traffic models. For example, according to (5), a self-similar traffic model gives

$$\begin{aligned} V(N(t+\tau) - N(t)) &= V(Y(t+\tau) - Y(t)) \\ &= \tau^{2H} V(Y(1)) \end{aligned}$$

This relationship means that a log-log plot of $V(N(t+\tau) - N(t))/\tau$ versus τ should be a straight line of slope $2H - 2$, see Figure 7. This property may be used to estimate H directly from the plot.

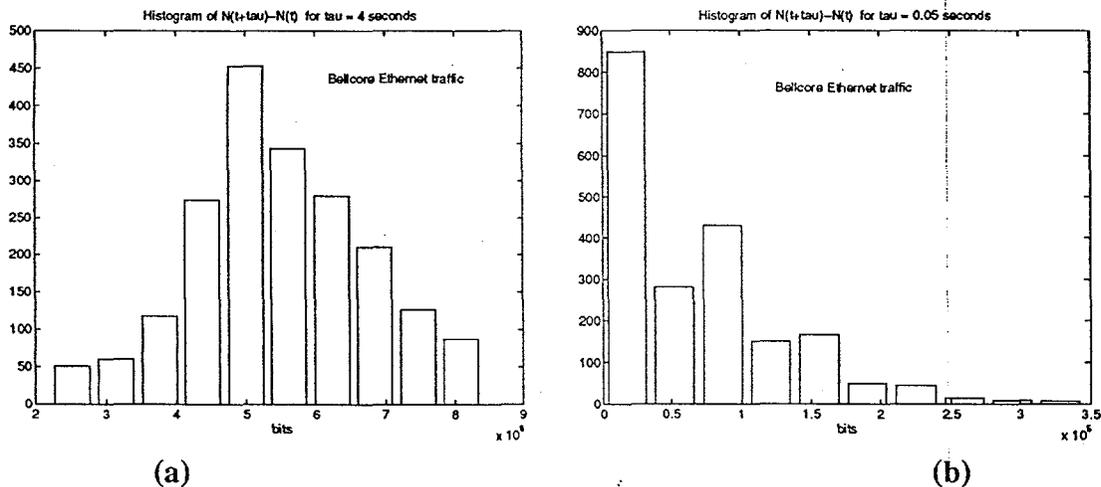


Figure 5. Histogram of the Increments of the Bellcore Ethernet Traffic for Different Values of τ .

Since a self-similar process is not fundamentally a counting process, it has nothing to say about discrete arrivals and consequently about interarrival times. Nevertheless with the addition of some "small scale structural assumptions", self-similar traffic consisting of a sequence $(t_1, n_1), (t_2, n_2), \dots$, where the t_i are interarrival times and the n_i are "packet sizes", can be generated. This may be done as follows. Suppose the following parameters are chosen:

- (1) λ = rate in bits/sec.
- (2) H = Hurst parameter (dimensionless)
- (3) τ = "characteristic time" in seconds.
- (4) n_p = packet size in bits.
- (5) distribution for the random variable $Y(\tau) = N(\tau) - \lambda\tau$.

Suppose moreover that the self-similar random numbers $Y(\tau), Y(2\tau), \dots$, with the characteristic distribution (5), can be generated. Then we have the increments

$$N((k+1)\tau) - N(k\tau) = Y((k+1)\tau) - Y(k\tau) + \lambda\tau$$

and $N((k+1)\tau) - N(k\tau)/n_p$ is the number of packets arriving in the time interval $(k\tau, (k+1)\tau)$.

Assuming these packets to arrive uniformly in this interval gives an interarrival time of

$$\Delta t_k = \frac{n_p \tau}{N((k+1)\tau) - N(k\tau)}$$

for packets arriving in this interval. In this "small scale structure" model we have constant packet sizes and uniform arrivals within each successive "characteristic time" interval of length τ .

More elaborate small scale models may also be introduced. For example, instead of uniform arrivals in the τ -intervals, Poisson arrivals with mean rates $N((k+1)\tau) - N(k\tau)/n_p \tau$ but still constant packet sizes could be used. Or a probability distribution for packet sizes could be incorporated with or without uniform arrivals in the τ -intervals.

3. Generation of Self-similar Traffic

In this section, we describe a method for generating self-similar traffic traces that conform to the model described in the previous section. The first prerequisite in the generation of self-similar traffic is the ability to generate random numbers from a self-similar process. We achieve this as follows:

- (1) Generate m (a power of 2) standard Gaussian independent random numbers, $w(-m/2+1), \dots, w(m/2)$.
- (2) For a given value H , consider the autocorrelation function

$$\rho(k) = \frac{1}{2} \left[(k+1)^{2H} - 2k^{2H} + (k-1)^{2H} \right] \quad \text{for } -\frac{m}{2} + 1 \leq k \leq \frac{m}{2},$$

and compute its fast Fourier transform (FFT)

$$\hat{\rho}(j) = \sum_{k=-\frac{m}{2}+1}^{\frac{m}{2}} \rho(k) e^{-\frac{2\pi i j k}{m}} \quad \text{for } -\frac{m}{2} + 1 \leq j \leq \frac{m}{2}.$$

Also compute the FFT, $\hat{w}(j)$, of $w(k)$.

- (3) Define

$$\hat{v}(j) = \hat{w}(j) \sqrt{\text{Re}(\hat{\rho}(j))}.$$

- (4) Compute the inverse fast Fourier transform (IFFT) of \hat{v} by

$$v(k) = \frac{1}{m} \sum_{j=-\frac{m}{2}+1}^{\frac{m}{2}} \hat{v}(j) e^{\frac{2\pi i j k}{m}} \quad \text{for } -\frac{m}{2}+1 \leq k \leq \frac{m}{2}.$$

The m numbers $\text{Re}(v(k))$ are then a simulation of the increments $X(k)$ of a Gaussian self-similar process with Hurst parameter H . $X(k)$ has mean 0 and variance 1.

The generation of the traffic increments now proceeds as follows.

- (1) Generate, as described above, Gaussian self-similar increments $X(k)$ of mean 0 and variance 1.
- (2) Map the random variables $X(k)$ to exponentially distributed random variables $X'(k)$ with mean $\lambda\tau$. This can be done by letting

$$X' = F_E^{-1}(F_G(X))$$

where F_G is the cumulative distribution of the standard Gaussian and F_E^{-1} is the inverse of the exponential distribution with mean $\mu = \lambda\tau$. Thus,

$$F_G(x) = \frac{1}{2} \text{erf}\left(\frac{x}{\sqrt{2}}\right) + \frac{1}{2}$$

and $F_E^{-1}(p) = -\mu \log(1-p)$. Note that we choose an exponential marginal distribution because this matches closely the marginal distribution observed for the Bellcore data at small time scales, as shown in Figure 5b.

- (3) Compute interarrival times for the time interval $k\tau \leq t \leq (k+1)\tau$ as follows. Let the number of packets for this interval be

$$M_k = \left\lceil \frac{X'(k)}{n_p} - \frac{1}{2} \right\rceil$$

where $\lceil \cdot \rceil$ denotes the ceiling function.

Then the constant interarrival time for this interval is

$$\Delta t_k = \frac{\tau}{M_k}.$$

A block diagram of the self-similar traffic generation method explained above is depicted in Figure 6. Figure 7 shows a log-log plot of the variance of $N((k+1)\tau) - N(k\tau)$ versus τ for measured Ethernet traffic at the rate of 1.38 Mbps with a Hurst parameter of 0.72 and for synthetic traffic generated by the above algorithm with $\lambda = 1.38$ Mbps, $\tau = 0.05$ sec., $n_p = 3280$ bits, and $H = 0.72$. The close agreement between the two sets of traffic on the

plot shows that the synthetic traffic indeed possesses the same degree of self-similarity (Hurst value) as the measured traffic.

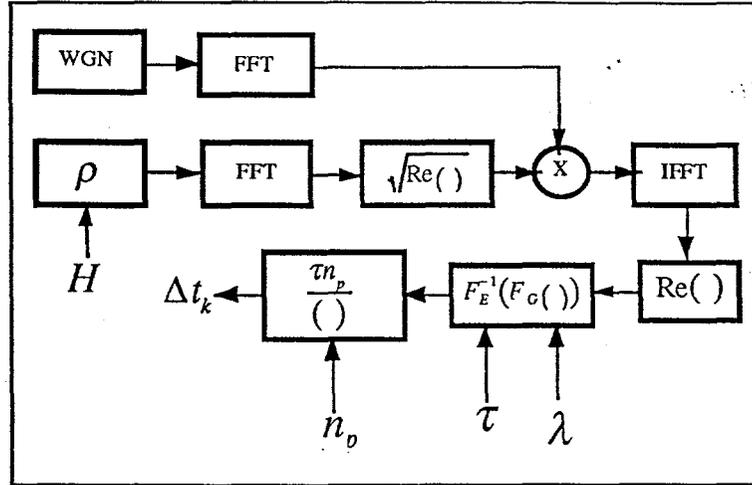


Figure 6. Block Diagram of the Self-Similar Traffic Generator.

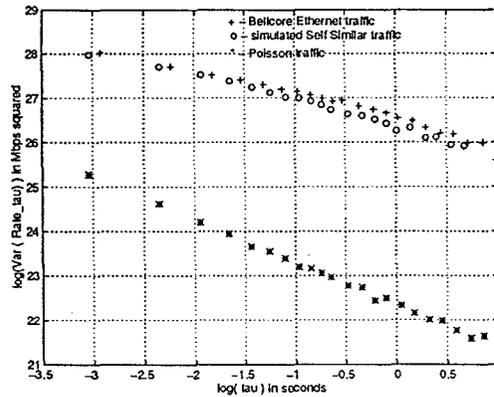


Figure 7. Log-Log Plot of the Poisson, Bellcore Ethernet and Simulated Traffic.

4. Simulation Study

We conducted a simulation study of an ATM network using both self-similar traffic, generated by the method described in section 3, and Poisson traffic. In our simulation model, the ABR service uses a feedback mechanism in which switches mark the congestion-bit (EFCI bit) of incoming cells whenever their queues grow beyond a predefined threshold of 500 cells. Further, we chose an ICR value of 10 Mbps and adopted an "optimal" ABR parameter set based on a previous study [8]. We compared the performance achieved under both the Poisson and self-similar traffic models.

As depicted in Figure 8, we used a two node configuration to collect statistics for a single session with and without background video traffic. All link capacities in the network are 155 Mbps (OC-3) and all link distances are 4 Km.

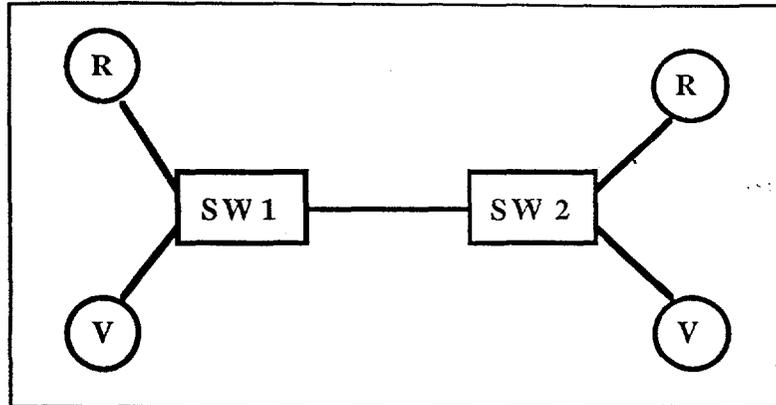


Figure 8. Two Node Network Topology With Two ABR and One VBR Connections.

Routers, denoted by *R*'s, multiplex and demultiplex traffic from attached Ethernet LANs and transmit on a single VC across the ATM link at the allowed cell rates calculated using the ABR protocol. Both the self-similar (Figure 10a) and Poisson traffic traces (Figure 11a) have an average arrival rate of 70 Mbps. Further, the self-similar trace has a Hurst parameter of 0.75. In the second study, we introduce an aggregate of ten MPEG II video streams as a higher priority background traffic. The presence of higher priority traffic induces periods of congestion in the ATM network which allows us to conduct performance evaluations of the ATM ABR flow control scheme. Figure 9 plots the aggregate video traffic of the ten MPEG II streams. As shown, the aggregated video traffic has an average arrival rate of 45 Mbps and a peak arrival rate of ~100 Mbps. In our simulation experiments, we were able to service our video traffic at its peak rate of 100 Mbps. We constructed our MPEG II streams by increasing, by an order of magnitude, the measured bandwidth of their corresponding MPEG I streams [9]. We believe this to be a reasonable method for generating MPEG II streams since the encoding schemes of MPEG I and MPEG II are identical and the two formats differ only in the image resolution, which translates into higher bandwidth.

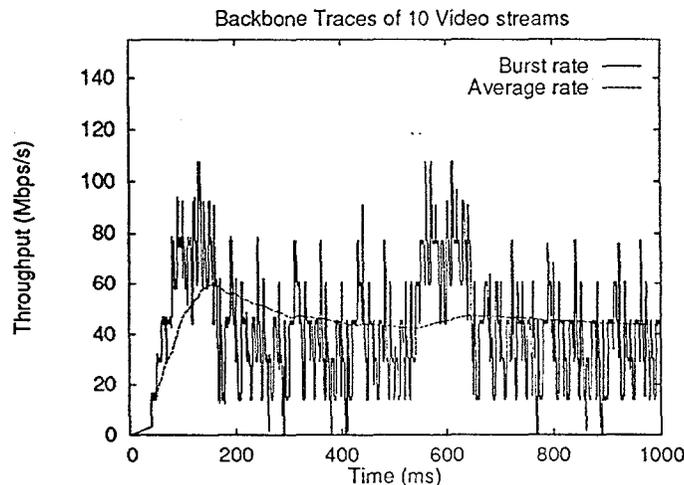


Figure 9. Aggregate Traffic Trace of 10 MPEG II Video Streams.

5. Results and Analysis

This section presents results and analysis of our simulation experiments. We focus our performance evaluation on buffer requirements as well as the ABR algorithm's response-time. First, we present a set of baseline results that do not include background traffic, followed by those that include background video traffic.

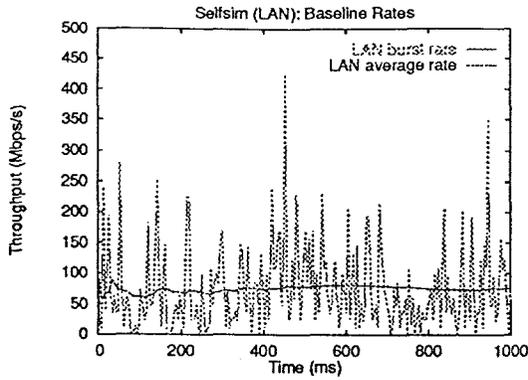
5.1 ABR without Background Video Traffic

Figures 10 and 11 plot our simulation results for the self-similar study and the Poisson study respectively. All rate measurements were made at intervals of 100-cell times, or 280 microseconds. In both figures, all four plots cover the same interval of simulation time, so that the sequence of events that affected the ABR performance can easily be observed.

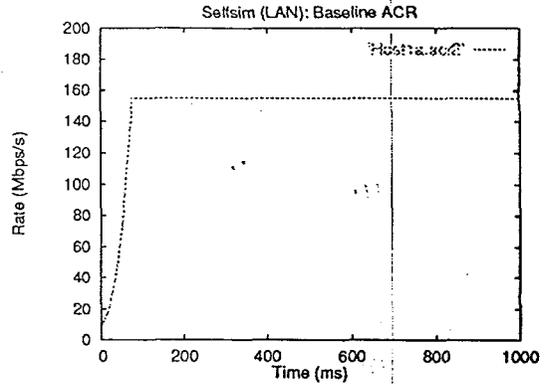
Figures 10a and 11a depict the burst as well as the average arrival rates of the self-similar and the Poisson sources at the source router. We observe that the self-similar source is burstier, but has an average rate identical to that of the Poisson source. In both cases, the burst rates occasionally exceed the available link rate of 155 Mbps. The Allowed Cell Rate (ACR) for each case, as measured at the source router, is plotted in Figures 10b and 11b. The ACR curves demonstrate that, in the absence of congestion in the ATM network (one source transmitting with an average rate less than 155 Mbps), the allowed sending rates for both traffic models achieved link speed. However, due to relatively long flow control feedback times, the increase from ICR to link rate was slow; the ramp-up time was ~75 ms for both traffic models. This behavior could lead to slow response and poor perceived performance for interactive network applications.

We plotted the source router's queue occupancy for the self-similar and Poisson cases in Figures 10c and 11c. As expected, because the ICR ramp up was slow, there was a large initial queue build up (~6000 cells) in both cases. Once the initial queue build-up cleared, the steady-state queue occupancy of the self-similar case was more than an order of magnitude larger than that of the Poisson case. This result indicates that the long-range correlation characteristics of the traffic drastically affect the nature of congestion experienced in the router.

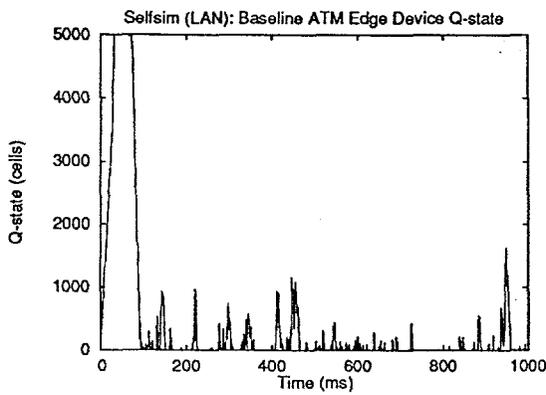
Figures 10d and 11d, which show the traffic arrival rate at the destination router, demonstrate that while the traffic for both the self-similar and Poisson models seemed to have maintained their original characteristics while traversing the network, their peak rates were limited to the available link rate. It would be interesting to statistically characterize the traffic admitted to the network by the ABR algorithm.



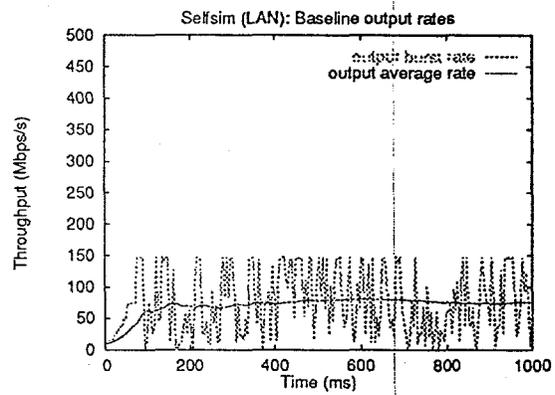
(a)



(b)



(c)



(d)

Figure 10. One Self-Similar Traffic Source, No Background Traffic. Histograms of: (a) the self-similar traffic arrival rate at the source router, (b) source router's ABR ACR, (c) source router's ATM NIC queue status, and (d) traffic arrival rate at the destination router.

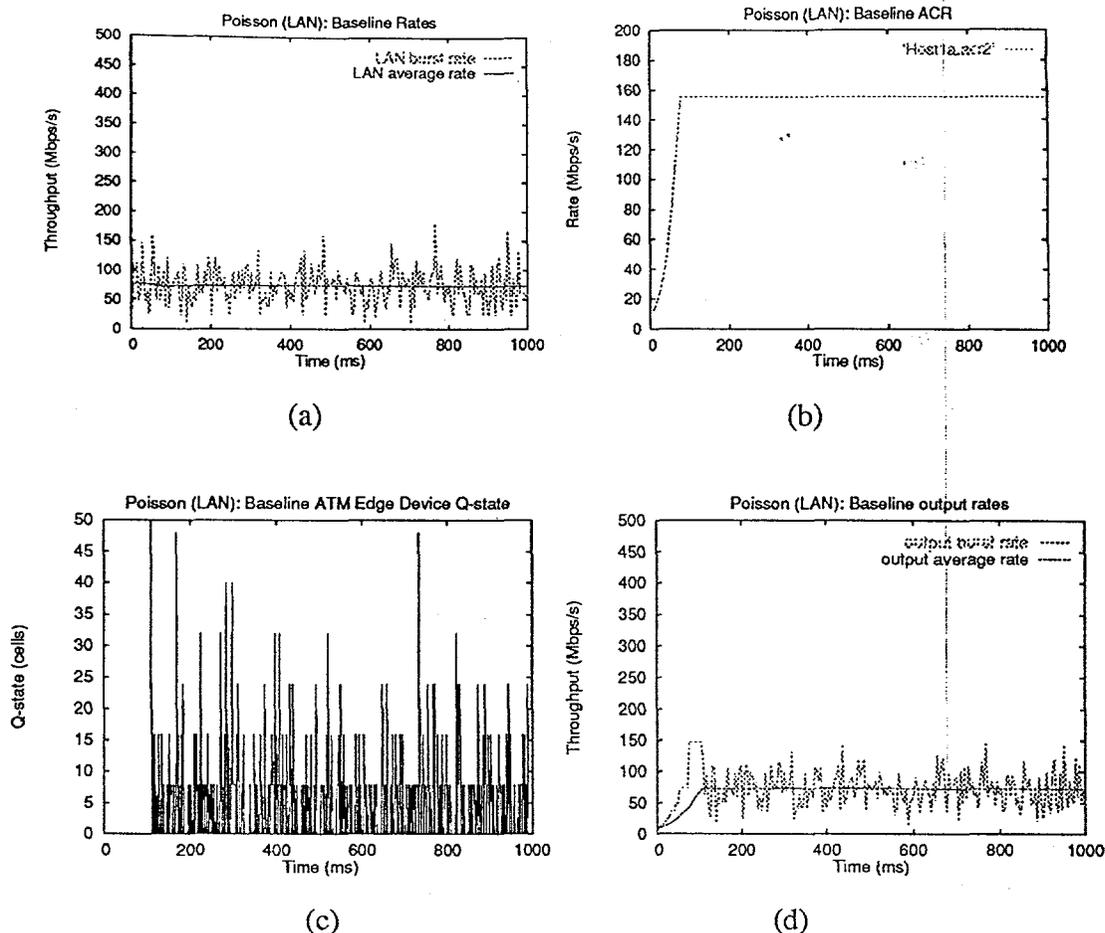


Figure 11. One Poisson Traffic Source, No Background Traffic. Histograms of: (a) the Poisson traffic arrival rate at the source router, (b) source router's ABR ACR, (c) source router's ATM NIC queue status, and (d) traffic arrival rate at the destination router.

5.2 ABR with Background Video Traffic

Now, we present a set of experiments that are similar to those in the previous section, but with the addition of background VBR video traffic. Here, in order to observe the congestion behavior in steady-state, we start a video stream 100 ms into the simulation run. Although the sum of the two average source rates does not exceed the backbone link rate, the sum of their burst rates often does. This condition provides us with the opportunity to evaluate ATM's ABR flow control for the self-similar and Poisson traffic models in the presence of variable-rate, higher priority background traffic. We used the background video plot for each case (Figures 12a and 13a) to identify periods of congestion, which are then used to study the corresponding ABR flow control behavior, the resulting ACR values, and the router and switch memory requirements.

The ACR curves in Figures 12b and 13b show that the sending rate of the self-similar traffic model was throttled by the ABR flow control mechanism, while the Poisson traffic was not. It is interesting that only the self-similar source's rate was affected, despite the fact that congestion was present for both types of traffic, as is clearly demonstrated by

Figures 12d and 13d. We deduce that the long-range dependence of the self-similar traffic caused prolonged periods of congestion which repeatedly resulted in switch queue buildup beyond the 500-cell threshold. The simulation results show that in the presence of the same background traffic, the Poisson traffic source was not able to cause sufficient queue growth to activate the ABR flow control's throttling mechanism.

Figures 12c and 13c present the corresponding queue status in the source routers for the self-similar and Poisson studies. Again, at steady state, the maximum queue size in the self-similar study is orders of magnitude larger than that of the Poisson study. We observed a significant increase in router buffer requirements for the self-similar source when compared to the case without VBR background traffic (compare Figures 10c and 12c). Those for the Poisson source, however, remained unchanged (Figures 11c and 13c) because the ATM switches were able to buffer the congestion caused by the video traffic without exceeding the ABR congestion threshold.

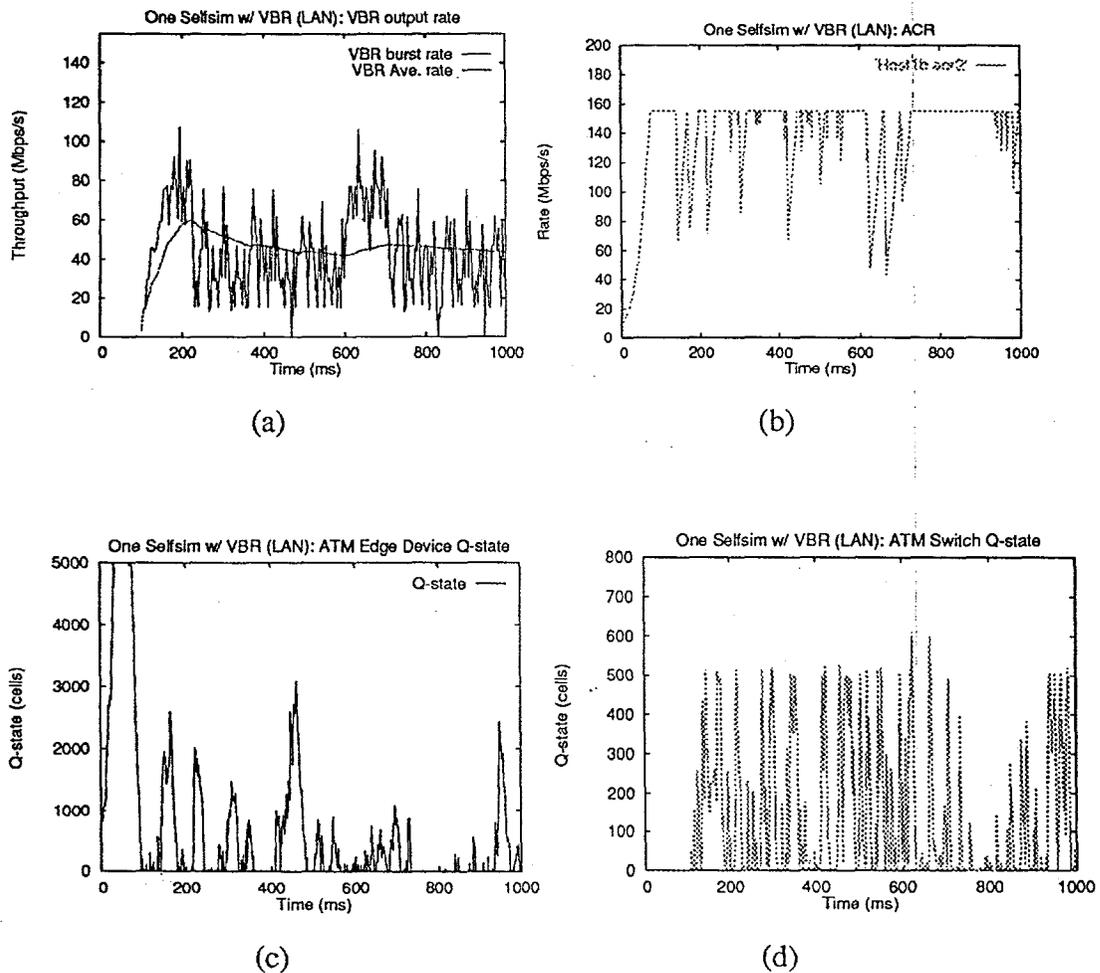


Figure 12. One Self-Similar Traffic Source With Background Video Traffic. Histograms of: (a) the MPEG II video traffic, (b) source router's ABR ACR, (c) source router's ATM NIC queue status, and (d) the queue occupancy at the congested ATM switch.

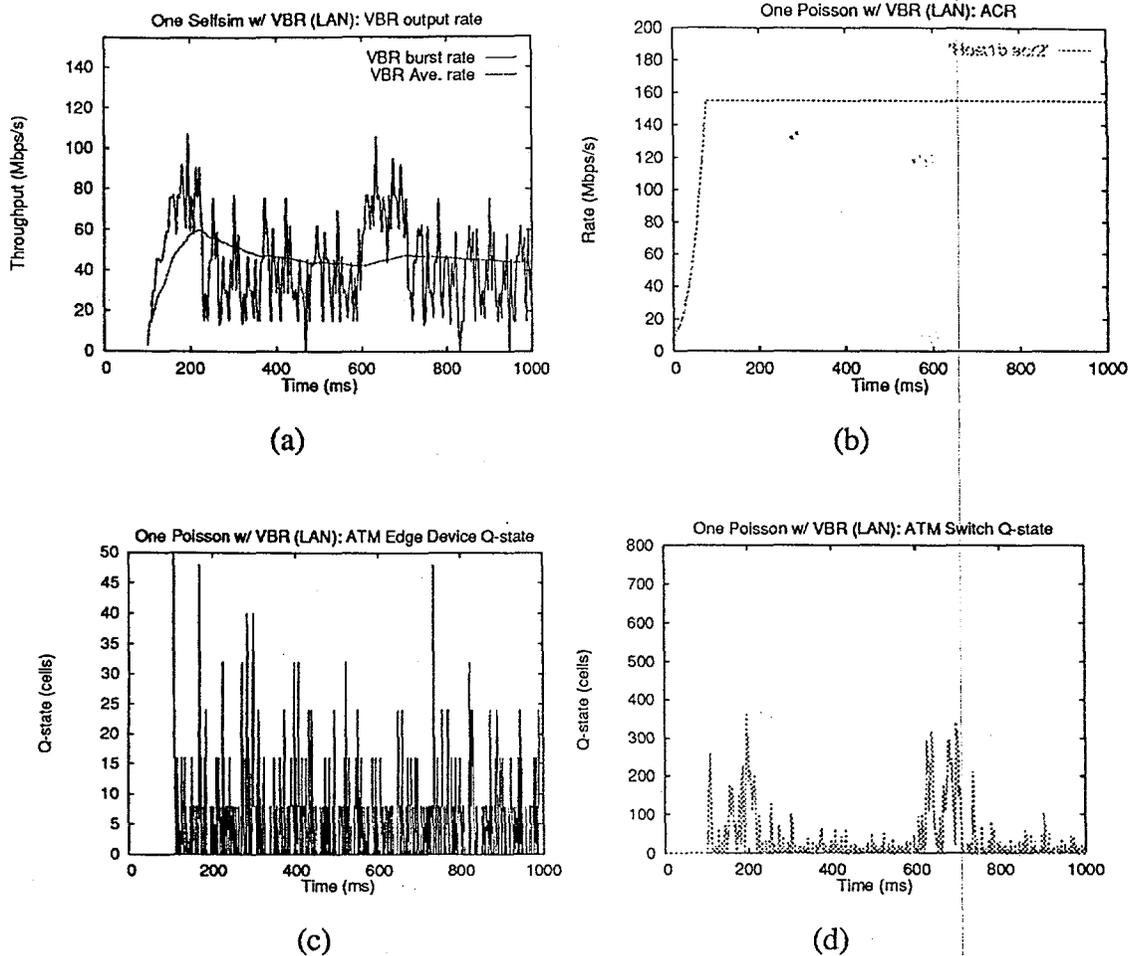


Figure 13. One Poisson Traffic Source With Background Video Traffic. Histograms of: (a) the MPEG II video traffic, (b) source router's ABR ACR, (c) source router's ATM NIC queue status, and (d) the queue occupancy at the congested ATM switch.

6. Conclusions and Future Work

Recent studies of existing Ethernet traffic suggest that conventional models are inadequate because they fail to capture correlation among arrivals. Furthermore, the emergence of ATM as a high-speed switching technology to carry both legacy LAN traffic as well as delay-sensitive application traffic motivates the study of more sophisticated traffic models to aid in predicting network performance. Self-similar models hold promise for accurately representing the correlation structure observed in measured traffic. In this paper, we presented a self-similar traffic model whose characteristics are consistent with those of actual aggregated Ethernet traffic. We described a method for generating self-similar traffic traces, and conducted a simulation study of self-similar traffic representing aggregated Ethernet traffic from a router, over ATM ABR service.

In our simulation study, we compared the performance achieved using Poisson and self-similar traffic over ATM's ABR service, both with and without the presence of VBR background traffic. We found that in the absence of background traffic, the ACR of both the Poisson and self-similar sources reached the full link capacity. However, the buffer

occupancy of the router at the edge of the ATM network was more than an order of magnitude larger for the self-similar case than for Poisson traffic. In the presence of background VBR video traffic, the self-similar source was throttled by the ABR flow control mechanism. Under the same conditions, the Poisson source was able to transmit unrestricted. The self-similar source, in conjunction with the VBR background traffic, created sufficient congestion at the ATM switches to invoke the EFCI congestion control mechanism. The Poisson traffic did not cause such persistent congestion. Also, as was true for the case without background traffic, the buffer occupancy at the edge router was orders of magnitude larger for the self-similar case.

From these results, we can conclude that self-similarity, or long-range dependence, in traffic has a profound impact on the network performance under the ATM ABR algorithm. The implication of our simulation results is that realistic traffic models are of vital importance for predicting the performance of high-speed networks. There are many areas for future work, including the refinement of our self-similar traffic generation method, and the study of the statistical properties of the traffic admitted to the network by the ABR algorithm when the source is self-similar.

7. References

- [1] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson, "On the Self-Similar Nature of Ethernet Traffic (Extended Version)," *IEEE/ACM Transactions On Networking*, Vol. 2, No. 1, Feb. 1994.
- [2] I. Norros, "On the Use of Fractional Brownian Motion in the Theory of Connectionless Networks," *IEEE Journal on Selected Areas in Communications*, Vol. 13, No. 6, August 1995.
- [3] M. W. Garrett and W. Willinger, "Analysis, Modeling and Generation of Self-Similar VBR Video Traffic," SIGCOMM '94, London, U.K., August 1994.
- [4] ATM Forum Contribution AF95-0013R10 Draft document of "UNI 4.1 ATM Forum Traffic Management Specification," February 1996.
- [5] M. de Prycker, "Asynchronous Transfer Mode - Solution for Broadband ISDN," Ellis Horwood Limited, pp. 97 - 124, 1992.
- [6] N. Likhanov, B. Tsybakov, and N. Georganas, "Analysis of an ATM Buffer with Self-Similar ("Fractal") Input Traffic," IEEE INFOCOM '95, Boston, MA, April 1995.
- [7] V. Paxson and S. Floyd, "Wide Area Traffic: The Failure of Poisson Modeling," SIGCOMM '94, London, U.K., August 1994.
- [8] C. Fang, A. Lin, "On TCP Performance of UBR with EPD and UBR-EPD with a Fair Buffer Allocation Scheme," AF/95-1328R, December 1995.
- [9] O. Rose, "Statistical Properties of MPEG Video Traffic and Their Impact on Traffic Modeling in ATM Systems," University of Wurzburg, Institute of Computer Science Report No. 101, February 1995.