Traffic Measurements from a Working ATM Network

Laura J. Bottomley and Arne Nilsson

Asynchronous transfer mode (ATM) networks present an opportunity for implementation of high speed networking for heterogeneous traffic transmission over a unifying network protocol. Mostly due to the expected diverse bandwidth requirements of ATM traffic, ATM networks need to place controls on input traffic to ensure performance of the bandwidth-on-demand network. In addition, diverse quality-of-service requirements require careful network design. Every study of network performance requires assumptions of some type to be made about network traffic. Any models which are proposed to describe ATM traffic must be verified with real traffic measurements.

As ATM networks are assembled and evolve, they will likely receive input from a variety of sources. One class of sources will almost certainly be legacy LANs. Router companies have already demonstrated their recognition of this fact as they begin to offer products designed to route from Ethernet, token ring, etc. networks to ATM. For this reason, characterization of legacy LAN traffic plays an important role in depictation of ATM input traffic.

The shape and behavior of ATM network traffic, both at the input to the network before adaptation and inside the network after packets have been broken into 53 byte cells, has significant impact on issues ranging from proper connection admission control to switch dimensioning. Peak allocation methods which rely on standardized traffic parameters may allow for guaranteed quality of service but may be grossly wasteful of network resources when the carried traffic is highly bursty. Indeed, these methods may prevent the types of statistical multiplexing gains that ATM was initially conceived to allow. Significant need exists for traffic measurements which can give a clearer picture of real traffic behavior.

This paper describes traffic measurements taken from several working networks, including the VISTANET ATM network, various Ethernet networks and a token ring network. VISTANET is a gigabit testbed sponsored by the National Science Foundation for research into high speed networking technologies and was designed to implement a medical imaging application over large distances in the state of North Carolina. The measurements were collected from both the customer-side of the network, before terminal adaptation, and also from the network side of the interface. For this reason, one of the measured ATM data sets consists of packet interarrival times and packet lengths, and the other ATM set comprises cell interarrival times.

The Ethernet traffic was measured at a university campus and at the gateway to the Internet from a supercomputing center at MCNC in North Carolina, and the token ring data was taken from a heavily loaded ring at IBM in Research Triangle Park, North Carolina. Each of these data sets consists of packet interarrival times and packet lengths.

The measurements were fit to various theoretical traffic models, as was deemed appropriate. The aim was to find mathematically tractable models which would adequately describe both the temporal and amplitude aspects of the measured stochastic process. All of the models examined were chosen for their ability to account for correlation between interarrival times. Careful attention was taken to make sure that any conclusions gleaned from the experiment are statistically justifiable.

1.0 Introduction

ATM networks present a great opportunity for networking of heterogeneous applications with a single unifying protocol. But this same opportunity brings a significant number of challenges. Connection admission, policing, monitoring, switch design, routing algorithms and decision points: all are areas where development and study is needed. Each of these diverse research and development topics requires assumptions about traffic behavior.

Few would argue that the best assumptions are those which are verified by measurement. Indeed, to that end, much emphasis has been placed on traffic measurement and modeling in the literature. The earliest studies of data traffic recognized its bursty behavior [FUCH70]. Since then work has proceeded along different directions to model traffic with the specific emphasis relying strongly on the area of interest of individual researchers (i.e. whether they are queuing theorists, switch or protocol designers, etc.). Measurements have been taken from lower speed networks such as Ethernet [e.g. SHOC80, BOTT91, LELA93, etc.] and from higher speed ATM networks, as they have become available [CROS95, JOU95, HOLT93, TALM94].

Of the ATM measurements referenced above, the last three deal with data from the VISTAnet network. The ATM data presented in this paper is from the same network. Analysis techniques are the same across all types of data. VISTAnet is a working ATM network whose driving application is Dynamic Radiation Therapy Planning (DRTP) (see figure 1). The network essentially operates as a metacomputer allowing a physician to have access to computing capabilities from a medical workstation which far exceed any previously available. The network allows treatment plans to be completed in very short periods of time using multiple sources of data.

The measurements presented here are from the traffic stream sent from the CRAY-YMP supercomputer across the network to the Pixel-Planes 5 machine. This is the network location where the greatest volume of traffic flow takes place. Section 3 presents data from both the user and network sides of the user/network interface (UNI). The authors gratefully acknowledge the assistance of MCNC in the acquisition of the traffic data.

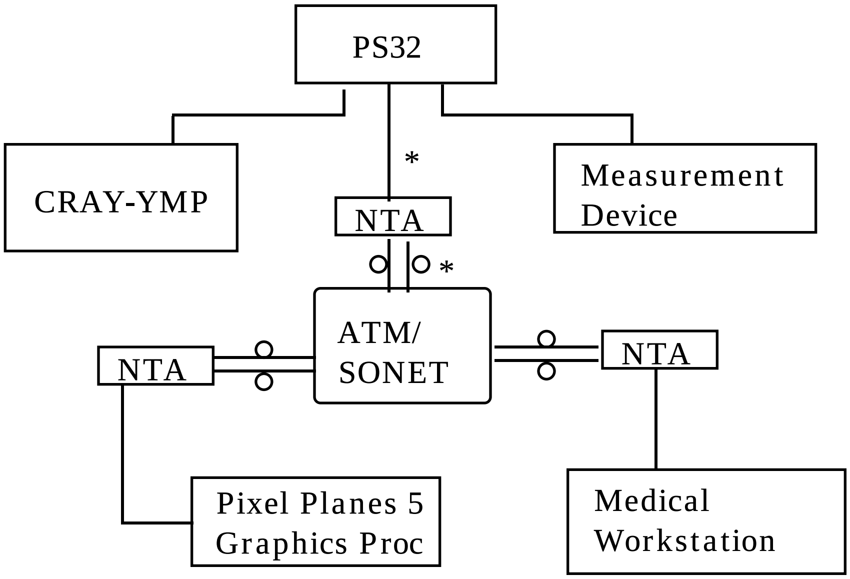

* Traffic Data Collection Sites

Figure 1

The Ethernet data were measured from very busy network segments at North Carolina State University and from MCNC in Research Triangle Park, North Carolina. The University is clearly an academic environment, while the MCNC data includes some academic and some industry traffic.

The token ring data is from an all industry environment at IBM in Research Triangle Park. The assistance of John Bartoles of ISSC is gratefully acknowledged.

2.0 Experimental Framework

2.1 Data Collection

Our philosophy of data collection is aimed at acquiring statistically significant samples during known busy periods. This traffic sampling technique is based on the sampling techniques used for many years in the telephony field. Traffic on any communications network is inherently non-stationary, as it is related to work and/or play patterns of human beings. Sampling during known busy periods is done acquire a "worst-case" traffic characterization which can be assumed to be stationary. The statistical techniques used for data analysis are formulated on the basis of stationarity.

Data traffic is recorded by a network analysis device as a series of interarrival times and packet lengths (where appropriate). The data recorded from the pre-UNI side (i.e. before terminal adaptation) of the VISTANET network were collected by a network analysis device designed by MCNC called HILDA. This device collects data from a 800 Mbps parallel HIPPI interface. The pre-UNI data consists of both interarrival times and packet lengths. Post-UNI data was collected from VISTAnet using a data collection system custom designed by Bell

South to record cell interarrival times at the OC-12 rate [TALM94]. During the collection period all arrivals were recorded.

Token ring data were recorded using a network sniffer, and Ethernet data with an Excelan EX5000 Series LANalyzer.

Experimental interarrival time density functions were calculated from the data by binning all interarrival times between 0 and 1 ms long, between 1 and 2 ms long, etc. Then serial correlation coefficients (sometimes called autocorrelation coefficients) were calculated from the data for a number of lags less than or equal to the square root of the number of packets/cells recorded [COX66]. This view of the data separates out the two aspects of a stochastic process in terms of amplitude and temporal variation to allow us to (theoretically) look at how each aspect affects network behavior.

2.2 Stochastic Models

After the experimental data have been processed, the data are fit to theoretical stochastic models. Many types of models exist, and the choices for candidate models are dictated by the uses to which the resultant models are to be put. Analytical queuing solutions and real-time monitoring algorithms versus complex network simulations may allow for quite different models in terms of complexity.

The candidate model initially considered for the data described in this paper is a three-state Markovian arrival process. The process description has been slightly modified to reduce the number of parameters in the model to six. Previous work with Ethernet data has shown that four parameter models, such as a two-state MMPP or two-state MAP, are unable to characterize correlation found in the data [BOTT92]. It was hoped that the increased number of parameters would allow for an adequate representation of correlation.

For a model to be sufficient, both dimensions of a stochastic process must be fit simultaneously. This can be accomplished by a variety of fitting techniques. Moment matches that include either a combination of interarrival moments and correlation coefficients and/or moments of the counting process can account for temporal and amplitude variation. Other techniques aim to extract parameters from some function of the arrival process such as entropy [DUFF94] or work [GIBB95]. The technique used for this paper is to fit the experimental interarrival density function and serial correlation coefficients to the equations for the theoretical model to determine the six model parameters.

We chose the equation fitting technique for a variety of reasons. We are very interested in how different aspects of an arrival process affect the behavior of a network. There is some evidence that the shape of the interarrival density function plays an important role in queuing behavior [BUX79]. In addition, work is in progress to demonstrate the role of correlation in queuing behavior. Preliminary results indicate that an inaccurate representation of degree of correlation can drastically affect queuing losses and/or delays for simple queuing models [ZHI92] and ATM switch models [JOU95]. One of the uses we hope to make of the current traffic measurements is to explore how correlation extending across many lags versus correlation of large magnitude affect various aspects of network behavior.

3.0 Traffic Data

Four types of data are presented in this paper: one set of data was collected from the user side of the user/network interface of the VISTAnet network and one set from the network side of the same interface. The different sets of data were taken when the same application was running on the network, but not simultaneously. Two sets of Ethernet data and one set of token ring data are also described. The data are presented in the form of packet/cell interarrival time density functions, serial correlation coefficients, and packet length distributions, as appropriate.

3.1 User-side ATM Data

Seven different sets of data were recorded each containing in the vicinity of 10,000 packets. Figure 2 shows an example of a packet interarrival density function from the user-side of the network interface. The speed of the HIPPI interface where data were recorded is 800 Mbps.
Each of the data sets exhibit a similar form for their interarrival time density functions. Keeping in mind that these are single application measurements, the structure is interesting. Essentially, most of the energy in the density function is placed at a few likely interarrival times.

Serial correlation coefficients for the same data set are pictured in figure 3. Recall that serial correlation coefficients are defined as:

$$\rho(k) = \frac{E[(X_i - E(X))(X_{i+k} - E(X))]}{\sigma_x^2}$$

where X_i is the ith interarrival time and σ_x^2 is the variance of interarrival times.

The periodic nature of the graph is believed to be due to the way the Cray buffers data for output. Clearly, single application data is heavily dependent on its

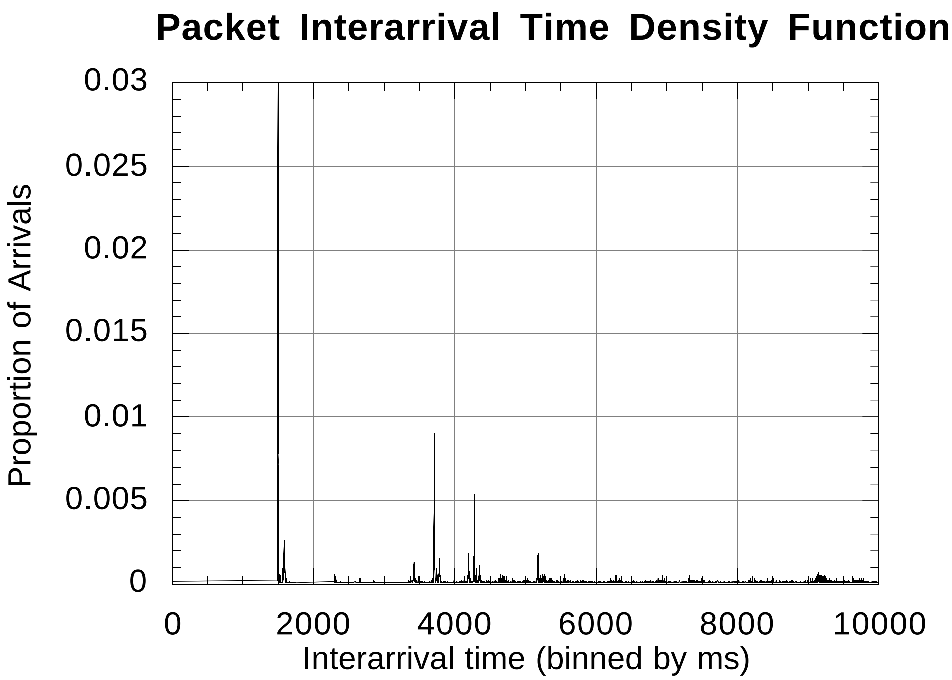

Figure 2

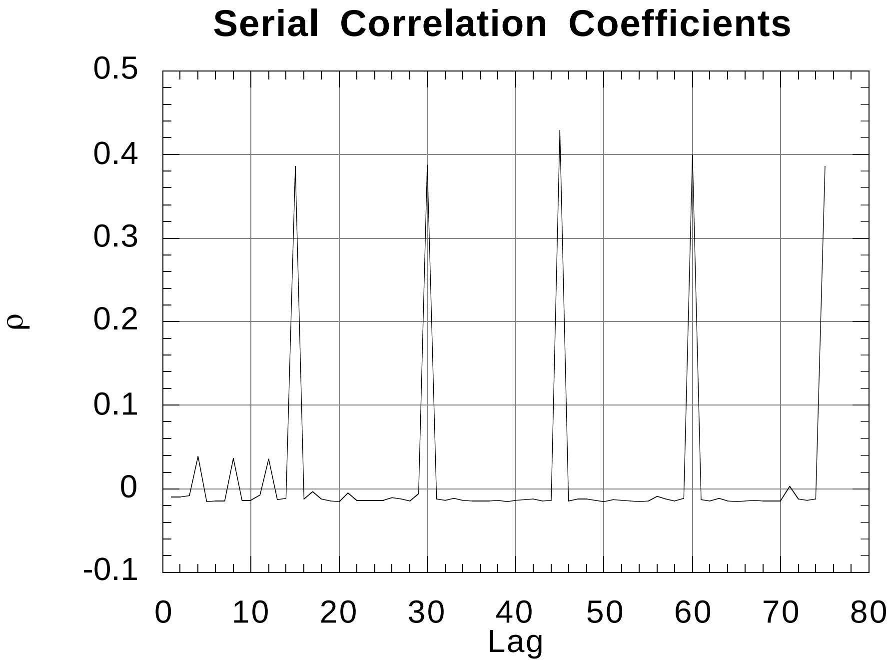

Figure 3

platform. Nevertheless, significant correlation is observed of a degree much higher than that seen in previous Ethernet data [BOTT92].

The data clearly do not fit any of the Markovian models (which all have some form of hyperexponential interarrival density function). However, these data may be represented by a much simpler process which is more deterministic.

The packet length distribution for the same data is depicted in figure 4.

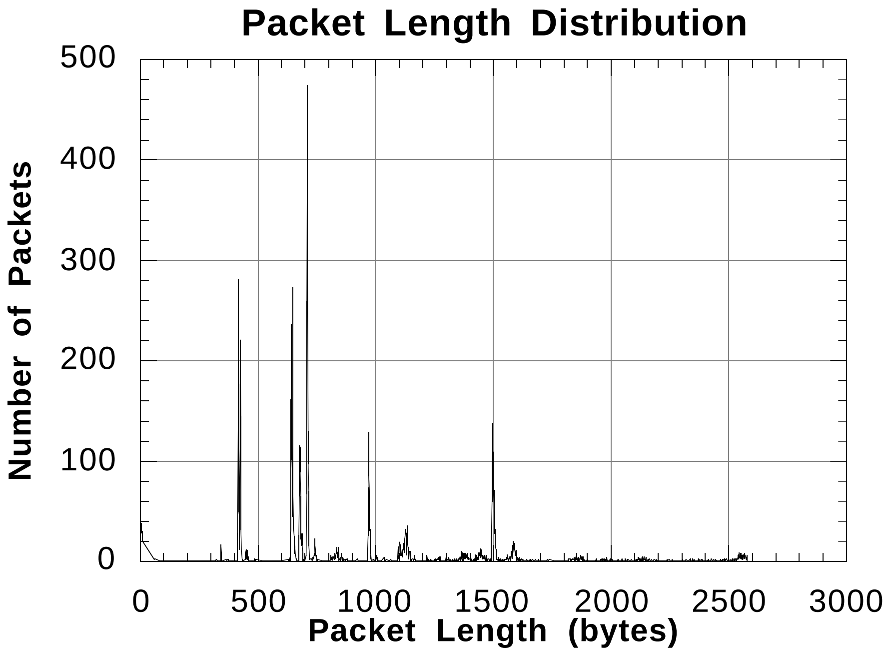

Figure 4

3.2 Network-side Data

Numerous data were collected from the network side of the ATM adapter. (The data files are available by anonymous ftp from kira.mcnc.org under pub/vista.net.) Traffic at this point in the network has been broken into 53 byte ATM cells, and the link is inherently slotted. The link speed is at the OC-12c rate. Because packet lengths are much greater than 53 bytes and the adaptation of the packets is done in real time, cells tended to arrive in large bursts. The cell interarrival distribution is thus heavily weighted toward a large proportion of very short interarrival times (i.e. 99% are below 1 ms). We thus calculated a burst interarrival density function, an example of which is shown in figure 5.

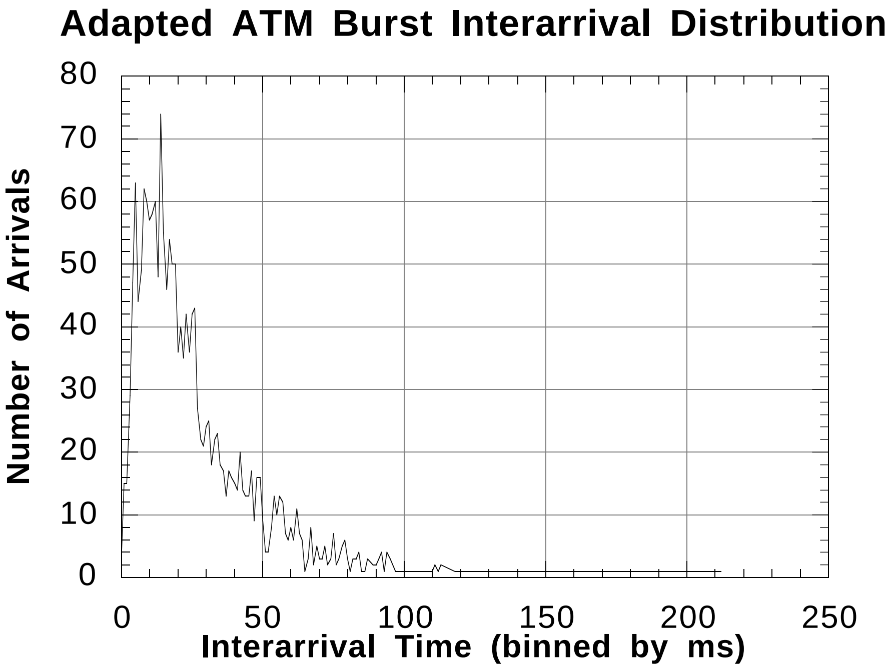

Figure 5

The burst length distribution is shown in figure 6.

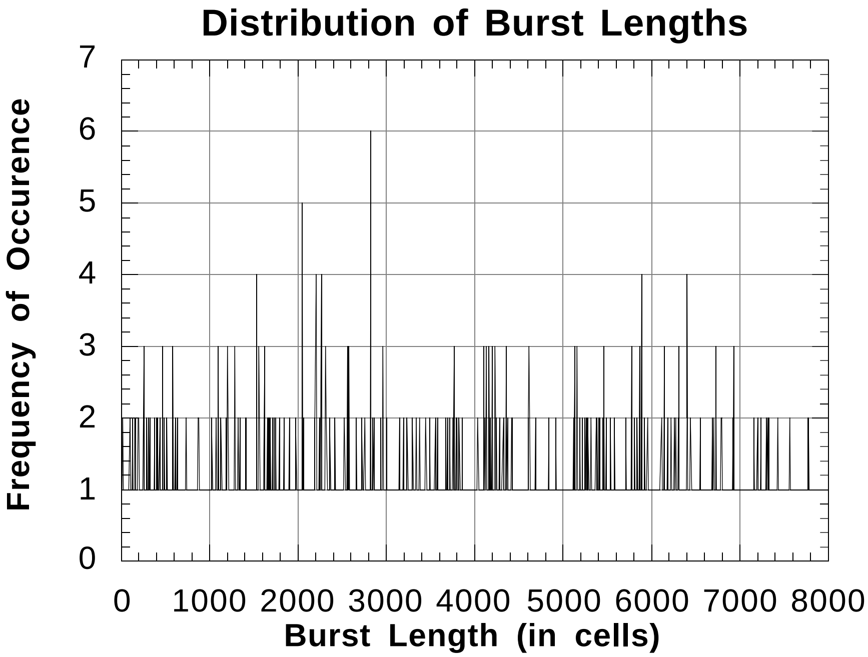

Figure 6

A burst is defined as a series of contiguous cells. The first empty slot marks the end of a burst, and the next cell to arrive marks the beginning of the next burst. The distribution of burst lengths is surprisingly uniform.

The burst interarrival density function was fit to a MAP-3 model. The resulting fit is quite poor. The MAP-3 interarrival density function has the form of a sum of three exponentials. Instead, a lognormal density function was found to be an excellent fit; see figure 7.

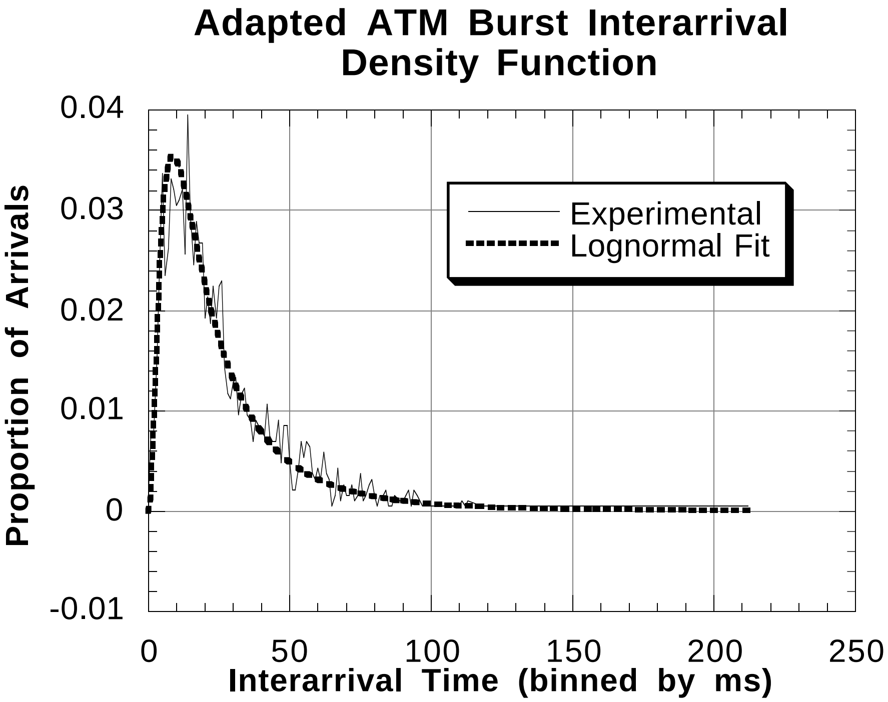

Figure 7

The serial correlation coefficients for burst interarrival times are pictured in figure 8.

A significance test for the data indicates a value of ρ over 0.05 is significant [COX66]. This implies that correlation between burst interarrival times is not significant. Correlation at the cell level is, on the contrary, as high as 0.9 at small lags. This is due to the tendency of cells to arrive in relatively large bursts.

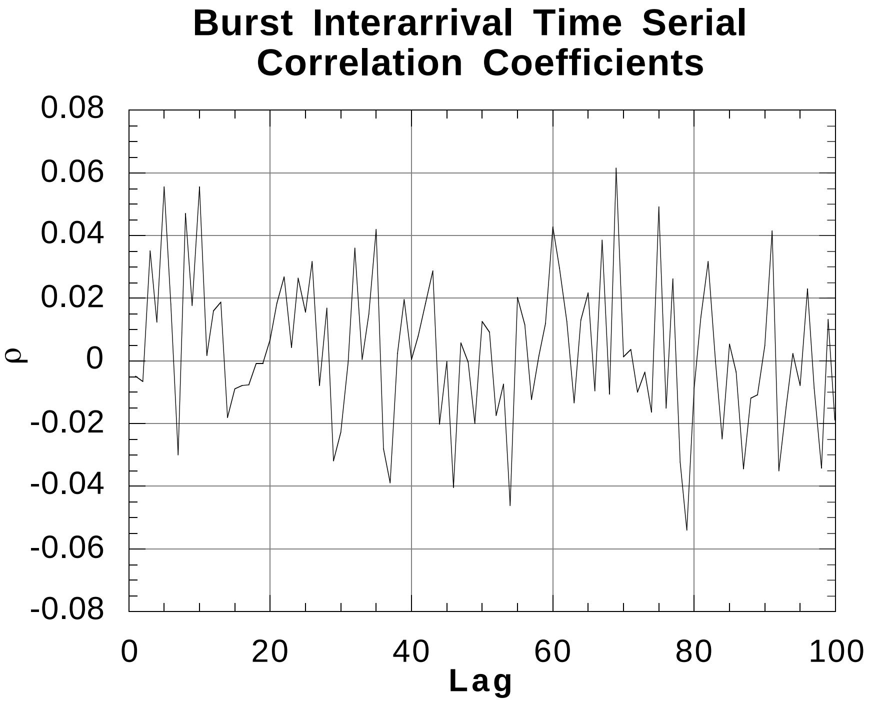

Figure 8

The shape of the burst interarrival density function and the results for cell and burst correlation imply that a simple two-state burst arrival renewal process may adequately represent this data. It also means, of course, that the Markovian-based models such as MAP and MMPP are not good fits for this data.

3.3 Ethernet Data

Two sets of Ethernet data are presented: one from North Carolina State University and one from MCNC. Each interarrival density function was fit to a MAP-3 model. The resulting fits are shown in figures 9 and 10. Both networks operate at 10 Mbps.

Serial correlation coefficients for the two data sets are shown in table 1.

Clearly, the Ethernet data present much less degree of correlation than the ATM data. Initial fits of interarrival density function simultaneous with correlation coefficients indicate that the MAP-3 process can fit up to the first ten correlation coefficients while maintaining a good interarrival density fit. This should be adequate to model this data, since lags with significant correlation are included.

Packet length distributions for the Ethernet data are strictly bimodal.

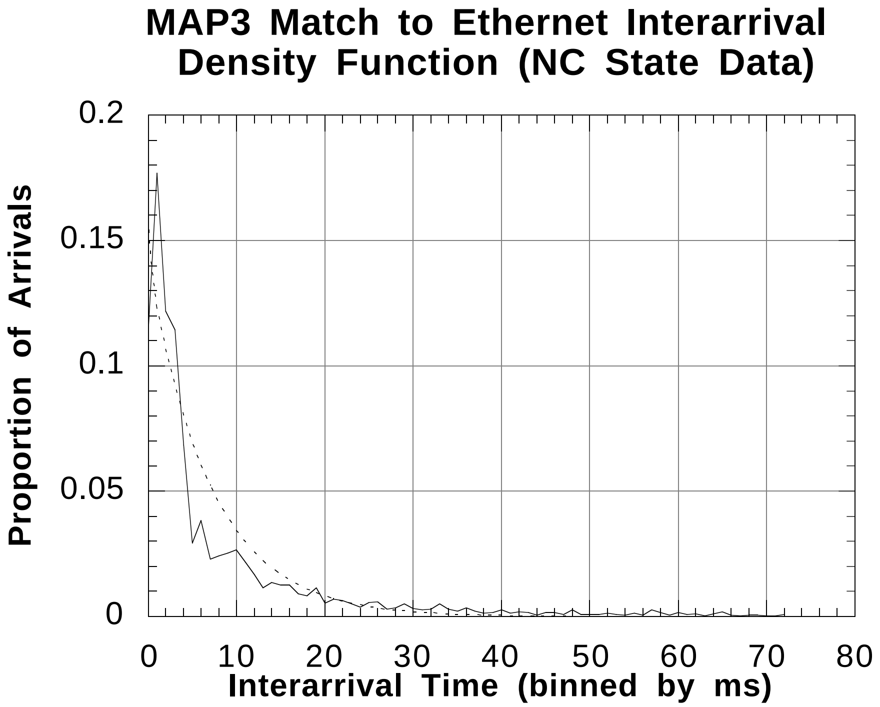

Figure 9

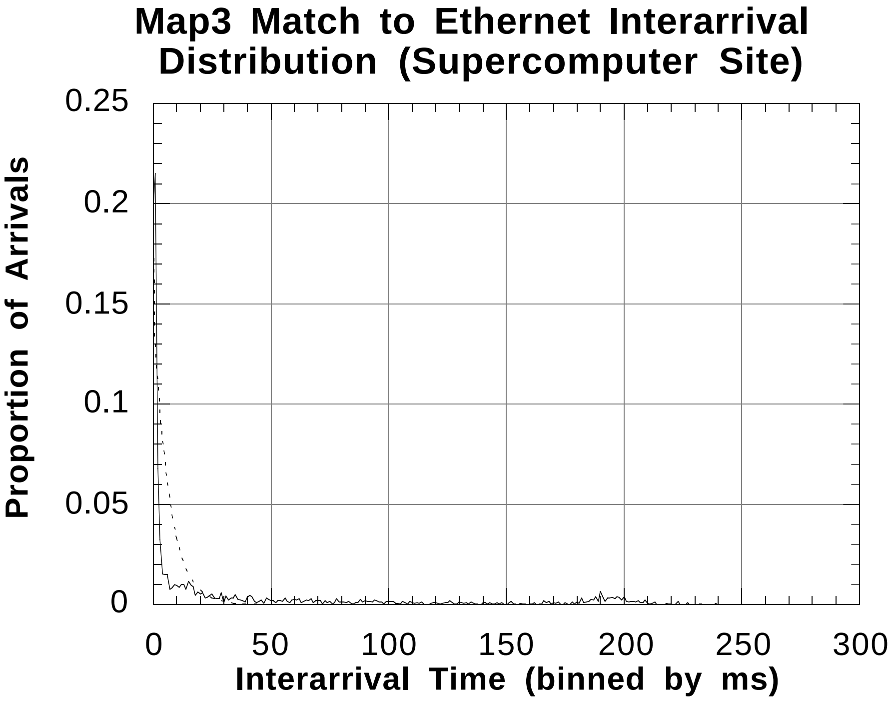

Figure 10

1	0.05	-0.09
2	0.07	0.16
3	0.03	0.06
4	0.05	0.08
5	0.04	0.05
6	0.03	0.08
7	0.05	0.05
8	0.02	0.08
9	0.003	0.04
10	0.04	0.06
5% Significance Level	0.07	0.09

Table 1

3.4 Token Ring Network Data

The token ring data was recorded from a network under approximately 64% utilization. The accuracy of the time stamps is somewhat less than desirable at only 1 ms. Each of the other data sets was recorded with at least 1 microsecond accuracy. The speed of the token ring network was 16 Mbps. Figure 11 shows the experimental token ring packet interarrival density function fit to a MAP-3 interarrival density function. Figure 12 shows the packet length distribution.

The MAP-3 fit is not ideal but is not bad either. More work is required to fit this data better, and greater measurement accuracy is desired.

Figure 13 shows the serial correlation coefficients calculated for this data. The simple significance test gives a threshold of 0.01 for this data. The nearly exponential decline is typical of Markovian arrival processes.

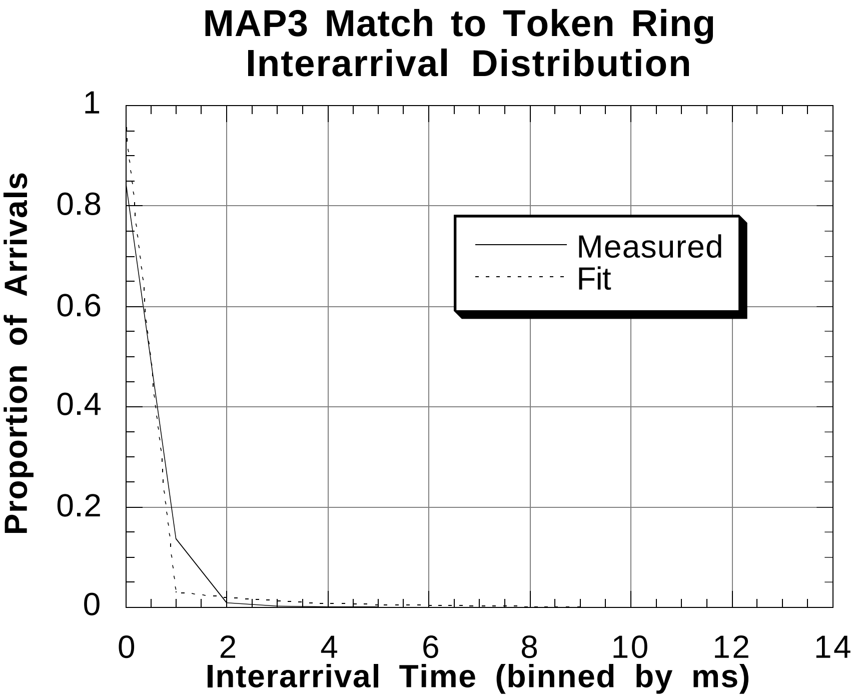

Figure 11

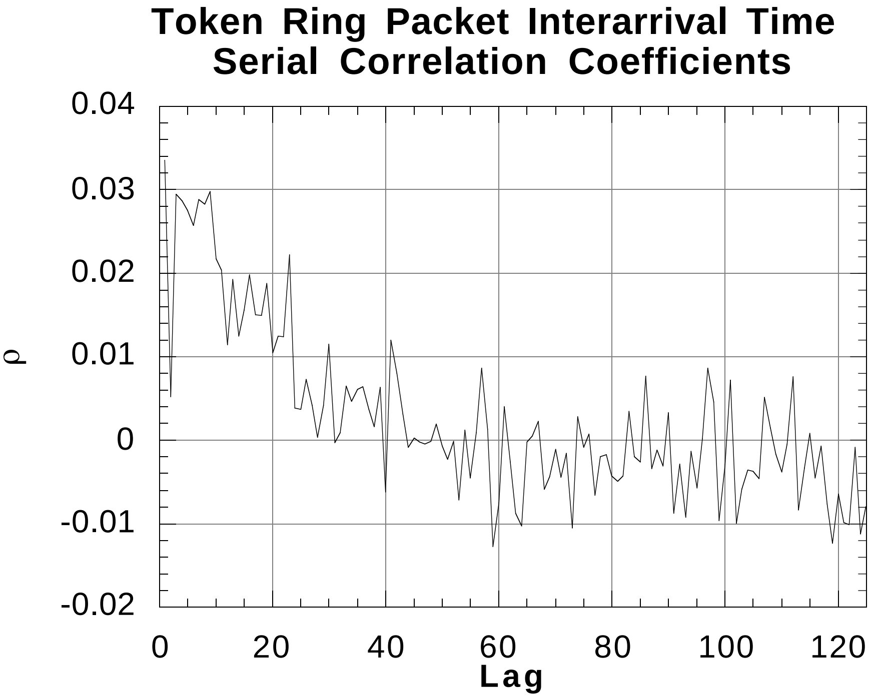

Figure 13

4.0 Conclusions and Future Work

The data collected from the VISTAnet network are from a single metacomputing application. Single application traffic has a tendency to be less well behaved from a modeling perspective, so we are encouraged to see the results from these measurements.

For the cases where "significant" correlation is observed, we seek an answer to the question, "How much correlation is significant from a network perspective?" As a result, in parallel with future measurement work, we are continuing experiments to identify how both degree and persistence of correlation affect various network protocols/policies, switch designs, etc. Preliminary results indicate that correlation can play an important role.

5.0 References

[BOTT91] Laura J. Bottomley, Arne A. Nilsson, and Alan J. Blatecky, "Traffic Measurements on a Working Wide Area Network," ITC13, Copenhagen, June 1991.

[BOTT92] Laura J. Bottomley, "Traffic Measurements on a Working Wide Area Network," PhD Dissertation, North Carolina State University, February 1992.

[BUX79] Bux, W., "Single Server Queues with General Interarrivals and Phase-Type Service Times," ITC9, Torremolinos, 1979.

[COX66] Cox, D. R. and P. Lewis, The Statistical Analysis of Series of Events, John Wiley and Sons Inc., 1966.

[CROS95] Crosby, Simon, et. al., "Bypassin Modelling: an Investigation of Entropy as a Traffic Descriptor in the Fairisle ATM Network," Proc. of the 12th IEE UK Teletraffic Symposium, Old Winsor, March 1995.

[DUFF94] Duffield, J. T., et. al., "Issues Raised by the Bellcore Data," Proc. of the 11th IEE UK Teletraffic Symposium, Cambridge, March 1994.

[FUCH70] E. Fuchs and P. Jackson, "Estimates of Distributions of Random Variables for Certain Computer Communications Traffic Models," Comm. of the ACM, Vol. 13, 1970.

[GIBB95] Gibbens, R. J. and F. P. Kelly, "A Decision-Theoretic Approach to Call Admission Control in ATM Networks," IEEE Journal on Selected Areas of Communications, 1995.

[HOLT93] Douglas Holtsinger, "Analysis of VISTAnet Traffic Measurements," CCSP Internal Report, North Carolina State University, 1993.
[JOU95] F. Jou and A. Nilsson, "ATM Traffic Measurement and Congestion Control Issues from VISTAnet Testbed," to be published, 1995.

[LELA93] W. Leland, et. al., "On the Self-Similar Nature of Ethernet Traffic," Proc. of SIGCOMM 1993.

[RIND95] A. Rindos, S. Woolet, I. Viniotis and K. Trivedi, "Exact Methods for the Transient Analysis of Nonhomogeneous Continuous-Time Markov Chains," 2nd Intnl.Workshop on the Numerical Solution of Markov Chains, Raleigh, NC, 1995.

[SHOC80] Shoch, J. F. and J. A. Hupp, "Measured Performance of an Ethernet Local Network," Comm. of the ACM, Dec. 1980.

[TALM94] K. Talman, "Asynchronous Transfer Mode Traffic Characterization on the VISTAnet Gigabit Network," Masters Thesis, Georgia Tech, April 1994.

[ZHI92] Zhi, Chi, private conversation, January 1992.