

Measuring ATM Traffic Cell-by-Cell: Experiences and Preliminary Findings from BAGNet

*W. Willinger¹, S. Devadhar², A.T. Heybey³, R. Sherman⁴,
M. Sullivan⁵, and J.R. Vollaro³*

*¹AT&T Labs-Research, ²Pacific Bell, ³Bellcore, ⁴California Institute
of Technology, ⁵Juno Online Services, L.P.*

*¹180 Park Avenue, Florham Park, NJ 07932-0971, U.S.A., Email:
walter@research.att.com;*

*²2600 Camino Ramon, San Ramon, CA 94583, U.S.A., Email:
sydevad@srv.pacbell.com;*

*³445 South Street, Morristown NJ 07960-6438, U.S.A., Email:
{ath, jrv}@bellcore.com;*

*⁴Division of Humanities and Social Sciences 228-77,
Pasadena, CA 91125, U.S.A., Email: sherman@amdg.caltech.edu;*

*⁵120 W 45th Street, 15th Floor, New York, NY 10036, U.S.A., Email:
sullivan@staff.juno.com.*

Abstract

This paper reports on our recent experience with cell-level ATM traffic measurements from the Bay Area Gigabit Testbed (BAGNet). In particular, we present a custom-built experimental prototype ATM monitoring device that is capable of non-intrusive and lossless recording of ATM traffic on OC-3 links, cell-by-cell, for extended periods of time. Using this monitor on BAGNet, we obtain a data set that is the first of its kind, consisting of complete (i.e., every bit of every ATM cell) traffic traces collected at different locations in the network (e.g., access line, trunk), over a period of about two weeks, and under a variety of network usage (e.g., normal operation, scheduled experiments). These measurements provide a wealth of information about network traffic at all protocol levels, but pose new challenges for data analysis and modeling. We present qualitative results from a preliminary exploratory analysis of some of the collected traces that illustrate the behavior of actual ATM traffic at the level of individual ATM cells (e.g., cell delay variation, “fractal” nature).

Keywords

ATM network, traffic recorder, ATM traffic measurements, fractal traffic.

1 INTRODUCTION

The technological advances that have made Gigabit and Terabit networks feasible need to be matched by capabilities to operate, manage and control these networks. Currently, our ability to manage traffic flows in these high-speed networks lags considerably behind our capability to deploy the necessary technologies, in part because of fundamental difficulties in analyzing and managing the empirically observed complex nature of bursty network traffic flows. Our approach for tackling these problems is experimental in nature and is firmly rooted in reality, i.e., based on traffic measurements from real networks that carry real services and applications generated by real users. Hence, our starting point is the collection of high resolution traffic traces from existing networks and emerging high-speed testbeds.

In the case of existing low to medium-speed mature packet networks (e.g., CCSN/SS7, ISDN, Ethernet LANs, MANs, WANs), this experimental approach has recently lead to a paradigm shift in our understanding of the nature of traffic in modern broadband networks, as documented for example, in Leland *et al.* (1994), Paxson *et al.* (1995), Willinger *et al.* (1997), and Crovella *et al.* (1995). In particular, the empirically observed “self-similar” or “fractal” behavior of measured network traffic is a new phenomenon that (i) is in sharp contrast to traditional traffic modeling assumptions, (ii) has already been shown to have practical implications for a wide range of traffic management and traffic engineering problems, and (iii) is currently the subject of numerous research efforts. However, when turning from low and medium-speeds to truly high-speeds (i.e., OC-3 and beyond), the hurdles and challenges for “measurement-driven” modeling and engineering work are daunting. For example, there currently exists no commercially available recording device that allows for the collection of complete ATM traces (cell-by-cell, lossless, non-intrusive), on OC-3 links, over multi-hour periods. (For recent measurements of a few seconds of actual ATM traffic from the Finnish University and Research Network FUNET, using the commercially available HP Broadband Series Test Equipment, see Molnar *et al.* (1996).) Moreover, complete ATM traffic traces collected at the cell-level over extended periods in time yield data sets that typically extend into the Terabyte range. Storing, processing, and managing such enormous data sets, as well as their effective analysis and efficient modeling requires high-performance computing capabilities and novel special-purpose software tools. In this paper, we demonstrate how advances in data recording and storage technologies have helped us to overcome these obstacles; in particular, we (i) present a custom hardware monitor for collecting ATM traffic on OC-3 links, (ii) report on the use of this monitor in an ATM traffic measurement experiment on BAGNet resulting in an unique data set of actual ATM traffic measurements, and (iii) summarize preliminary findings from an exploratory analysis of a small subset of the obtained ATM traffic traces. Note that the load levels in a typical experimental Giga-

bit testbed environment are generally low, and the main trunks are typically utilized by only a few active (permanent) virtual circuits. On the other hand, the measured traffic is “real”, that is, generated by real users running real applications.

In the absence of commercially available ATM traffic recorders that suffice for our purpose, we first describe our custom-built experimental prototype of a programmable recording device that allows for non-intrusive and loss-less recording of complete traffic traces on 155 Mbps (OC-3 speed) ATM links for multi-hour periods. The collected traces consist of complete (header plus payload) copies of all ATM cells seen on the ATM link during the measurement period (typically between hours and days, depending on the link load) augmented with timestamp information that is accurate to within 50 nanoseconds. We used this recording device as part of two ATM traffic measurement experiments that took place in the second half of September 1995 and in the first half of April 1996, respectively, and involved the Bay Area Gigabit Network (BAGNet). For a detailed overview of BAGNet, see Berc *et al.* (1996). We describe here the data set resulting from the September’95 ATM traffic measurement experiment. This data set is the first of its kind in the sense that it consists of complete cell-level ATM traces that contain all the information needed for analyzing ATM traces at all layers of interest (e.g., cell, AAL5, IP, TCP). Moreover, the traffic traces were collected over a two-week period, at different points in the network (e.g., access lines, trunks) and under a variety of different network usage (e.g., normal operation, showcase applications, scheduled experiments). Finally, we illustrate some findings from a preliminary exploratory analysis of a small subset of this data set. We emphasize that the analysis presented in this paper is preliminary and largely qualitative in nature, and concentrates exclusively on ATM cell-level characteristics; a more quantitative assessment of the statistical properties of these traces and the perhaps more intriguing analysis of the behavior of higher-level protocols in a “real-world” ATM network are currently in progress and will appear elsewhere. We study actual *Netperf* applications at the ATM cell level and provide empirical evidence that while some of the more stringent requirements for cell delay variation (CDV), discussed at length in the Standards bodies, are of theoretical interest, they are typically not implemented in first-generation ATM products. Moreover, our qualitative analysis lends empirical support to an observation made in Paxson *et al.* (1995) (slightly paraphrased) that “... *we will soon see broadband traffic which is perforce self-similar, simply due to the source characteristics of its individual connections*”. Specifically, we consider ATM traces consisting mainly of motion JPEG video applications, illustrate their dynamics below as well as above the frame level (i.e., 20-25 msec), and relate our findings to earlier work characterizing VBR video traffic at the frame level (e.g., see Garrett *et al.* (1994), Beran *et al.* (1995), and Huang *et al.* (1995).)

The rest of the paper is organized as follows. In Section 2, we describe the

design and performance of our ATM traffic recorder. Section 3 gives a brief overview of BAGNet, its topology and user community, and summarizes the traffic traces that were collected from BAGNet using our recording device. In Section 4, we present preliminary findings from an exploratory analysis of a small subset of the collected data, focusing on ATM cell-level characteristics. We conclude in Section 5 with some remarks about current and future work concerning this unique traffic data set.

2 MEASURING TRAFFIC AT OC-3 LINK SPEEDS

Commercially available ATM analyzers as well as special-purpose ATM recording devices used in the past (e.g., the HiPPI Link Data Analyzer HILDA developed by MCNC Winkelstein *et al.* (1991)) are severely limited in (i) how many contiguous cells they can capture, and (ii) how many bytes from each cell they can record. These limitations have restricted past ATM traffic data analysis work both in scope and in the level of detail and have stood in the way of checking theoretical assumptions about ATM traffic against actual data. In contrast, our approach seeks to (i) capture complete (i.e., header plus payload of every cell) ATM traces over extended periods of time and (ii) analyze the complete data set at all layers of interest. In the ATM context, capturing complete cells is of crucial importance because higher-level protocol information is typically part of the payload of ATM cells. In this section, we describe the design of our own experimental prototype of an ATM traffic recorder that is suitable for our approach and overcomes the limitations of existing equipment. We also report on our experience using this recorder in the field.

2.1 Hardware Configuration of ATM recorder

The hardware configuration of our ATM recording device consists of 3 main parts: an ATM/HiPPI interface that currently resides on a single circuit module, a high-speed tape drive that uses a removable, industry standard digital cassette that can store 96 gigabytes of data, and a control computer (Sun SparcStation) that starts the recorder and coordinates the HiPPI transfers. HiPPI is a switched, 32 bit parallel interface standard used in high-performance computer networks. It features an 800 Mbps transfer rate and is used here to maintain compatibility with the HiPPI-based high-performance computing environment where the post-processing and analysis of the data is performed.

The recording device is connected to an operational fiber-based (SONET/OC-3) ATM network through an optical splitter that needs to be installed on a designated ATM link prior to monitoring. The splitter taps 10 percent of the optical power from the link for this purpose. The ATM recorder converts

the received data into a byte stream and marks the first byte of each ATM cell. The “start of cell” marker causes the time stamp circuit to insert a 48 bit time stamp and framing flags into the byte stream preceding each cell. The inserted time stamp has a resolution of 50 nanoseconds. Subsequently, the byte stream (including time stamps and framing flags) gets converted to HiPPI format and is buffered in the HiPPI interface prior to transfer to the HiPPI compatible tape drive. Note that the control computer is not in the recorder’s data path; it simply coordinates the ATM/HiPPI transfer. While this design feature is crucial to the recorder’s high throughput, it prevents us from processing or filtering the data during collection (of course, one could build extra hardware to perform the latter tasks). An additional feature of the current architecture of the ATM recorder is that the ATM/HiPPI interface resides on a single circuit module. As a result, only a single OC-3 ATM link can be monitored at a given time. At full link load, 4-5 hours of contiguous ATM cells can be stored on a single 96 Gbytes cartridge.

To support simultaneous recording of multiple fibers, a future version of the recorder will allow multiple ATM receivers to be multiplexed on a single HiPPI interface. This design change will increase the capabilities of the recorder but will cause the potential bandwidth of multiple ATM links to exceed the bandwidth of the single HiPPI interface. However, since today’s ATM links are rarely fully loaded, the HiPPI interface will typically be able to cope with current aggregate loads, handling short bursts at peak rates by buffering on the receiver and HiPPI interfaces; to ensure adequate performance of the recorder under future higher link loads, extensive stress testing of the equipment will be necessary.

2.2 Field Experience with the ATM Recorder

Prior to field deployment, the equipment underwent numerous tests in a lab environment to ensure that after connection to an ATM switch via an optical splitter (i) the recorder does not generate any signals that are fed into the switch ports, and (ii) the reduced signal strength has no effect on switch function. These tests demonstrated that the recorder is a genuinely passive recording device.

Although the ATM recorder is designed for field deployment, our first-hand experience based on transporting the equipment across the country resulted in some re-designs that should yield a more “field hardened” version of the current monitor. For example, when post-processing the data collected with the current version of the ATM recorder on BAGNet (see below), we detected errors in a number of the recorded time stamps. The problem turned out to be hardware related, i.e., a register in the time stamp circuitry became lose as a result of improper shipping and handling of the equipment, and was failing intermittently. Fortunately, these failures corrupted only the most

significant 24 bits of the 48 bit time stamp (which do not change often), and the failures occurred relatively infrequent. Thus, it was possible to completely and accurately restore the corrupted time stamps in a single run through the data.

3 BAGNET TRAFFIC MEASUREMENTS

The traffic recording device described in Section 2 was used in an ATM traffic measurement experiment on the Bay Area Gigabit Network (BAGNet) during the latter half of September 1995. BAGNet is an ATM testbed provided by Pacific Bell under a CalREN (California Research and Educational Network) grant and connects 15 commercial, educational, and research organizations in the San Francisco Bay Area by a 155 Mbps ATM metropolitan area network. Below, we first discuss briefly its topology, architecture, user community, and dominant (showcase) applications (for more information about BAGNet, see Berc *et al.* (1996)). We then report on the traffic measurement experiment itself and summarize relevant information concerning the unique data set resulting from this experiment.

3.1 BAGNet Topology and User Community

The BAGNet topology is depicted in Figure 1. The BAGNet backbone ATM switches are located in Pacific Bell's Central Offices (COs) in Oakland and Palo Alto and connect the 15 sites in a network about 75km in diameter. The switches use Virtual Circuit Identifiers (VCIs) to route individual ATM cells through the network. In BAGNet's case these are permanent VCs (PVCs), i.e., the ATM switches and ATM hosts have been pre-configured with the routing information. The resulting PVC ATM mesh connects 60 hosts at the 15 sites. All links in the network are OC-3 links and provide an ATM Forum conforming User-Network Interface (UNI) to the customer. Each site is minimally required to support Classical IP over ATM (see for example, Laubach (1994)) and is equipped with high-performance workstations that can communicate over the network at up to 155 Mbps. Service is provided on a best effort basis, with no guarantees on the Quality of Service (QoS) parameters such as Cell Loss Ratio (CLR), Cell Transfer Delay (CTD) and Cell Delay Variation (CDV).

The purpose of CalREN's support for BAGNet is to promote research and development in high-speed communications and related applications, including the development, deployment, and demonstration of existing and new applications that utilize modern high-performance, long-distance communications. BAGNet is demonstrating high-speed communications using the TCP/IP protocol suite, with high-speed motion JPEG compressed video playing a key role. BAGNet's showcase application is the teleseminar application

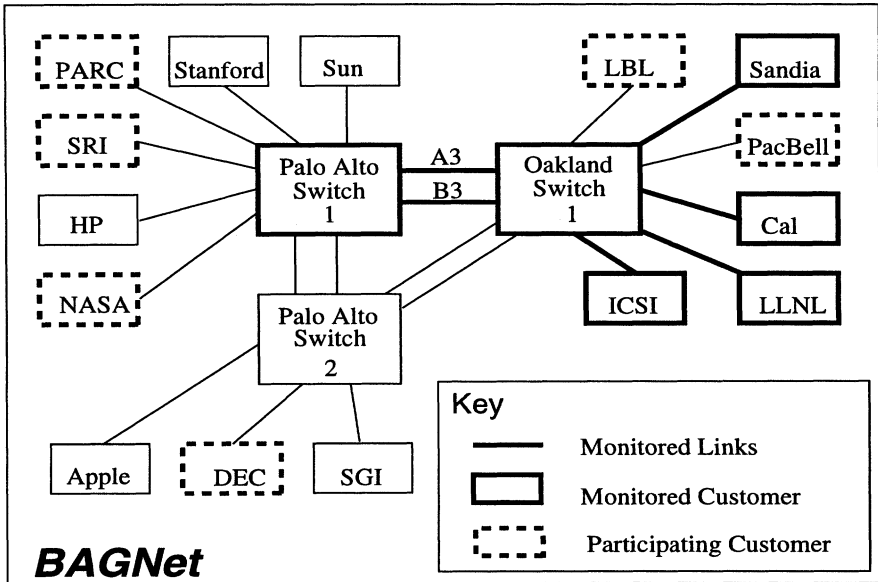


Figure 1 BAGNet Topology and User Community

(initially developed for use in the Internet and collectively referred to as the *MBONE* tools), but there are numerous other applications that are experimented with on BAGNet on a regular basis, e.g., network performance tests utilizing *Netperf*, high-rate data transfer experiments, collaborative science and simulation environments, and distance learning.

3.2 ATM Traffic Data Collection

Our brute force approach to ATM traffic data collection consisting of capturing every bit from every ATM cell over extended periods of time has obvious implications for data privacy and confidentiality. However, BAGNet was an ideal site for this ATM traffic measurement experiment. Being primarily a research-oriented group of participants, there already existed a genuine desire to learn about the behavior of ATM traffic generated as a result of a mixture of traditional and novel ATM-specific applications, and concerns about data confidentiality could easily be mitigated by reasonable guarantees for limited access to the complete traffic traces and for sensible “data scrubbing” prior to their distribution to a wider audience.

After soliciting the participation of the BAGNet participants, optical split-

ters were installed on OC-3 ATM links in the Oakland CO on two trunk ports and four customer ports. The two trunks between the OKLD and PLAL switches in the Oakland CO and Palo Alto CO, respectively, are connected to ports "A3" and "B3" on the OKLD switch and are therefore referred to as the "A3 trunk" and "B3 trunk," respectively. Starting September 11, 1995, the OC-3 lines were monitored. After a week of verification and making preliminary observations on various lines, actual data collection began on September 18, 1995. Lawrence Livermore National Labs (LLNL) access was monitored during this initial period. In the week of September 25, 1995, a coordinated network-wide measurement experiment was performed in cooperation with the BAGNet users. During this period, we have very detailed information about the network activities of the individual BAGNet participants and the state of the network. Finally, from September 29 until October 6, 1995, the "A3" trunk was monitored.

During the the course of this experiment, traffic traces totalling about 200 Gbytes were collected on BAGNet. We separated the data into three groups corresponding to the three main monitoring periods (i.e., earlier part of September, week of September 25, and late September/early October), and focused mainly on measurements obtained during the middle period. For this period, a detailed experiment schedule exists that identifies the host IP addresses of the participating workstations and the VCs used for the connections, describes the main applications that were run (e.g., Mpeg2 stream to Lockheed and WWW Hypertext to LLNL, on Sept. 27, between 8:00 and 9:00 am), and also provides information about the background traffic (e.g., ping, MPEG video) that was present at the time of the experiments.

4 EXPLORATORY DATA ANALYSIS: PRELIMINARY RESULTS

The traffic traces described in Section 3 were transferred from the tape drive of the ATM data recorder to a HiPPI-based high-performance computing environment that is available to us for post-processing of the data and for data analysis. We report here on preliminary findings from an exploratory analysis of two data sets, data set #1 and data set #2. The data shown below have been extracted from the raw ATM traces using *Tribeca* (see Sullivan (1995)), an experimental database management system optimized for network traffic analysis. *Tribeca* supports a query language that allows for filtering, aggregation, and grouping operations to entire traces or to specific segments of the traces. We emphasize that the analysis in this paper concentrates exclusively on ATM cell-level behavior. Higher-level protocol characteristics and performance properties will be reported elsewhere at a later date.




Figure 2 Time series plots of the number of ATM cells per second for the total traffic on the link (top), and for VC 39-only traffic (bottom).

4.1 Netperf over ATM

Running a simple *Tribeca* filtering query on data set #1 for identifying the active VCs and determining their respective traffic volume revealed one dominant VC (VCI=39), generating more than 95% of the total traffic on the access line from LLNL to the Oakland CO. The primary application running on this VC was *netperf*, a network performance tool that is capable of measuring transactional and stream UDP and TCP throughput between two hosts on the network. In this case, the intent of running *netperf* was to validate cell pacing (at speeds of 1 Mbps, 5 Mbps, 10 Mbps, 25 Mbps, 40 Mbps, and ABR, respectively, for five minutes each, sending 8KB UDP/IP datagrams) performed by a commercially available ATM interface.

Using conventional time series plots (number of ATM cells per seconds), Figure 2 (top) depicts about 25 minutes worth of the total traffic observed on the access line from LLNL to the Oakland CO. The bottom plot in Figure 2 shows the traffic on VC 39 for the same time period. These plots clearly identify the cell pacing experiment on VC 39 as the main application, with relatively little background traffic present (mainly pings). The cell pacing experiment starts halfway into the plot, and four different cell pacing rates (1, 5, 10 and 25 Mbps) are clearly visible.

To obtain information about the *netperf*-generated traffic at the level of

individual ATM cells, we show in Figure 3 four different plots involving the interarrival times of contiguous ATM cells during a period when the cell pacing rate was constant, namely 1 Mbps (similar observations hold for the other rates). For a series of 15,000 consecutive ATM cells, Plot (a) shows for every cell (horizontal axis) the corresponding cell interarrival time (vertical axis). The plot suggests that with a few exceptions, the ATM interface emits the cells emitted by the *netperf* application in a more or less periodic fashion (i.e., constant interarrival times, with minor variations). However, by ignoring the outer band of large interarrival times in Plot (a), it is easily seen (Plot (b)) that the ATM cell interarrival times generated by the ATM interface card are not constant, but are instead tightly clustered around approximately 9-10 distinct values.

Additional insights into the cell-level properties of measured ATM traffic can be obtained by studying *phase plots* of time series of interarrival times; in the case at hand, a phase plot is obtained by plotting the interarrival time of the $(i+1)$ -st cell against the interarrival time of the i -th cell. The structure revealed by the phase plot in Figure 3(c) (using the data from Plot (b)) clearly displays the possible combinations of successive ATM cell interarrival times. Moreover, by connecting the points in Plot (c) by lines indicating the temporal order of occurrence (see Plot (d)), we add to the static plot in Figure 3(c) a time component and obtain thus a temporal phase plot that reveals considerable structure concerning ATM cell interarrival times corresponding to the *netperf* application. One conclusion from these phase plots (and similar ones for the remaining cell pacing rates) is that successive interarrival times of ATM cells generated by *netperf* tend to be negatively correlated (i.e., large interarrival times are typically followed by shorter ones), and that the dynamic features exhibited by the phase plots remain unchanged under different cell pacing rates. However, at this stage, it is unknown whether this feature is a result of the cell pacing's scheduling algorithm or of the ATM interface card itself. In any case, these plots provide empirical evidence that although some of the more stringent cell delay variation requirements, discussed in Standards bodies such as ITU or the ATM Forum, are of theoretical interest, first-generation ATM products are generally less interested in implementing these requirements; as shown here, these requirements may, in fact, already be violated before the cells enter the network.

4.2 Motion JPEG over ATM

In the case of data set #2, there are 3 dominant VCs (VCI=486, 490, and 501), generating more than 95% of the total traffic recorded on the trunk from the Oakland CO to the Palo Alto CO. The experiment scheduled during this measurement period asked all BAGNet participants to transmit motion JPEG video at as high a rate as possible. During this period, the ATM traffic

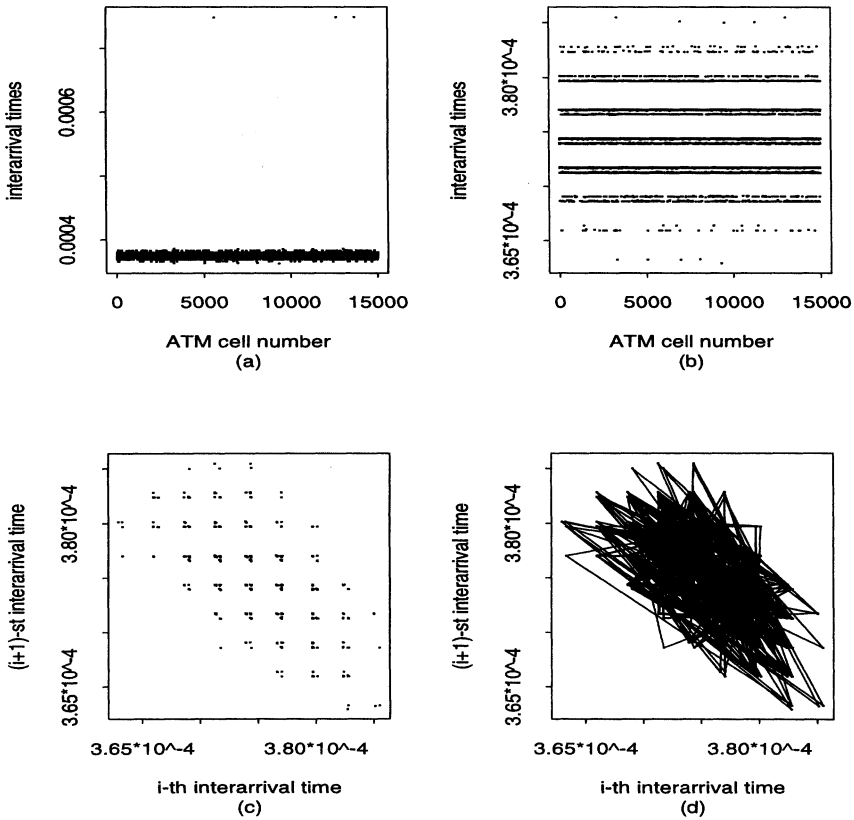


Figure 3 *Netperf*-generated traffic at the ATM cell level: (a) time series of 15,000 interarrival times, (b) close-up of (a), (c) static phase plot, (d) dynamic phase plot.

recorder monitored the trunk on the senders' side. Since only 3 active VCs were observed, 2 out of the 5 potential participants failed to take part in this experiment.

Figure 4 displays the time series plots of the number of ATM cells per time unit for about 23 minutes for the total traffic (top panel), VC 486-only traffic (second panel), VC 490-only traffic (third panel), and VC 501-only traffic (bottom panel). The bottom three plots depict three different types of motion JPEG and illustrate a range of diverse dynamic behaviors (at this time, we have no information about scene types and/or particular JPEG implementations used during the experiment). The top plot shows the aggregate traffic consisting of the three motion JPEG traces and a small




Figure 4 Time series plots of the number of ATM cells per second for (a) total traffic on the link, (b) VC 486-only traffic, (c) VC 490-only traffic, and (d) VC 501-only traffic.

amount of background traffic (mainly pings). Intuitively, this plot illustrates that even though the total traffic load on the monitored link is only about 5%, the traffic exhibits variations and fluctuations over a wide range of time scales, i.e., traffic bursts appear on small time scales (“spikes”) as well as on large time scales (“swells”); we refer to this empirical feature as the *self-similar* or *fractal* nature of broadband network traffic.

Figure 5 represents a visual illustration of the fractal characteristic of the

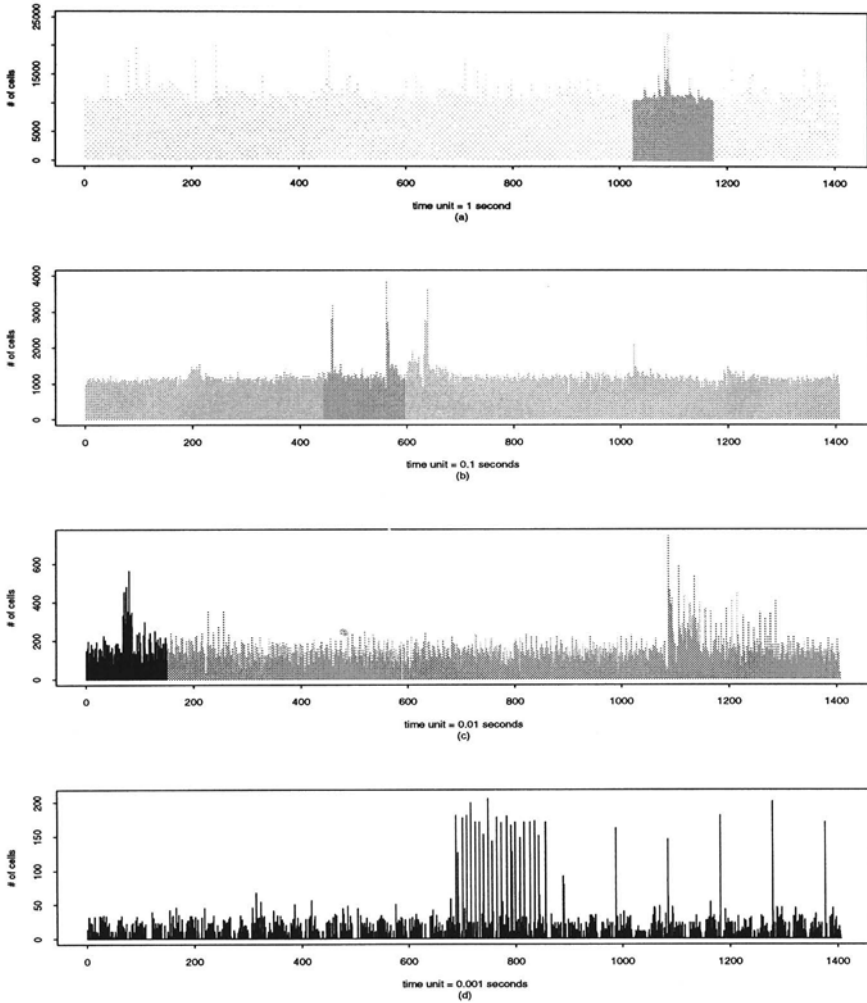


Figure 5 Time series plots of (aggregate) ATM traffic on four different time scales.

measured (aggregate) ATM trace. Plots (a)-(d) depict a plot sequence of time series of ATM cell counts (number of cells per time unit) for 4 different choices of time units, namely 1 second in Plot (a), 100 milliseconds in Plot (b), 10 milliseconds in Plot (c) and 1 millisecond in Plot (d). Subintervals viewed on a smaller time scale are indicated by a darker shade in each plot. Although visually, the plots are not necessarily “similar” to each other, they do exhibit the crucial fractal characteristic, namely fluctuations and variations over a wide

range of time scales. In other words, the ATM trace at hand does not show the same “clean” self-similarity property observed in the Ethernet LAN traces considered in Leland *et al.* (1994). However, the trace exhibits a behavior that is consistent with a fractal traffic model; for example, using by now standard techniques (for more details, see for example Leland *et al.* (1994) and Paxson *et al.* (1995)), the trace can be shown to be consistent with long-range dependence and asymptotic self-similarity. For this data set, the absence of *exact* statistical self-similarity is to be expected since many (independent) sources need to be aggregated before a second-order statistical characteristic such as self-similarity suffices to accurately describe the dynamics of network traffic (for a set of precise conditions, see Erramilli *et al.* (1996)). Recall, that this data set is the aggregate of only three (independent) sources.

For a different and more striking way of visualizing the nature of traffic generated by individual sources or particular applications (such as VBR video) at the ATM cell level, we apply a simple yet powerful exploratory data analysis tool called *textured dot strip plot* or simply *textured plot*, originally proposed in Tukey *et al.* (1990) (see also Willinger *et al.* (1997) for an application of this technique in the context of network traffic data analysis). Intuitively, the idea of textured plots is to display one-dimensional data points (e.g., arrival times of single ATM cells) in a strip in an attempt to show all data points individually. Thus, if necessary, the points are displaced vertically by small amounts that are partly random, partly constrained. The resulting textured dot strip facilitates a visual assessment of changing patterns of data intensities in a way other better-known techniques such as histogram plots, one-dimensional scatterplots, or box-plots are unable to provide, especially in the presence of extreme values. To illustrate the effectiveness of textured plots for visualizing and assessing the dynamics of ATM traffic generated by an individual source or application, we display in Figure 6 the textured plot associated with the motion JPEG application on VC 490 on 5 different time scales. Note that each point in each plot represents the time of arrival of an ATM cell. More precisely, the textured plot in the top panel in Figure 6 (Plot (a)) represents the initial 23 seconds of motion JPEG on VC 490; Plot (b) takes a closer look (i.e., at the time scales of a few seconds) at the first four seconds of the top plot and reveals the distinct frame structure underlying JPEG video that is not visible in the first plot. This frame structure and a more refined block-within-frame structure become much more apparent when we magnify the time scale further (i.e., 100 milliseconds in Plot (c), 10 milliseconds in Plot (d)); furthermore, the plots identify the frame size for the motion JPEG transmitted on VC 490 to be about 20-25 milliseconds. Finally, Plot (e) at the bottom of Figure 6 shows motion JPEG at the ATM cell level for the duration of slightly more than a single block, and identifies the block size within a frame for the motion JPEG transmitted on VC 490 to be between 0.5-1.0 milliseconds. The textured plots in Figure 6 give a comprehensive view of (a certain implementation of) motion JPEG over ATM; in particular, we

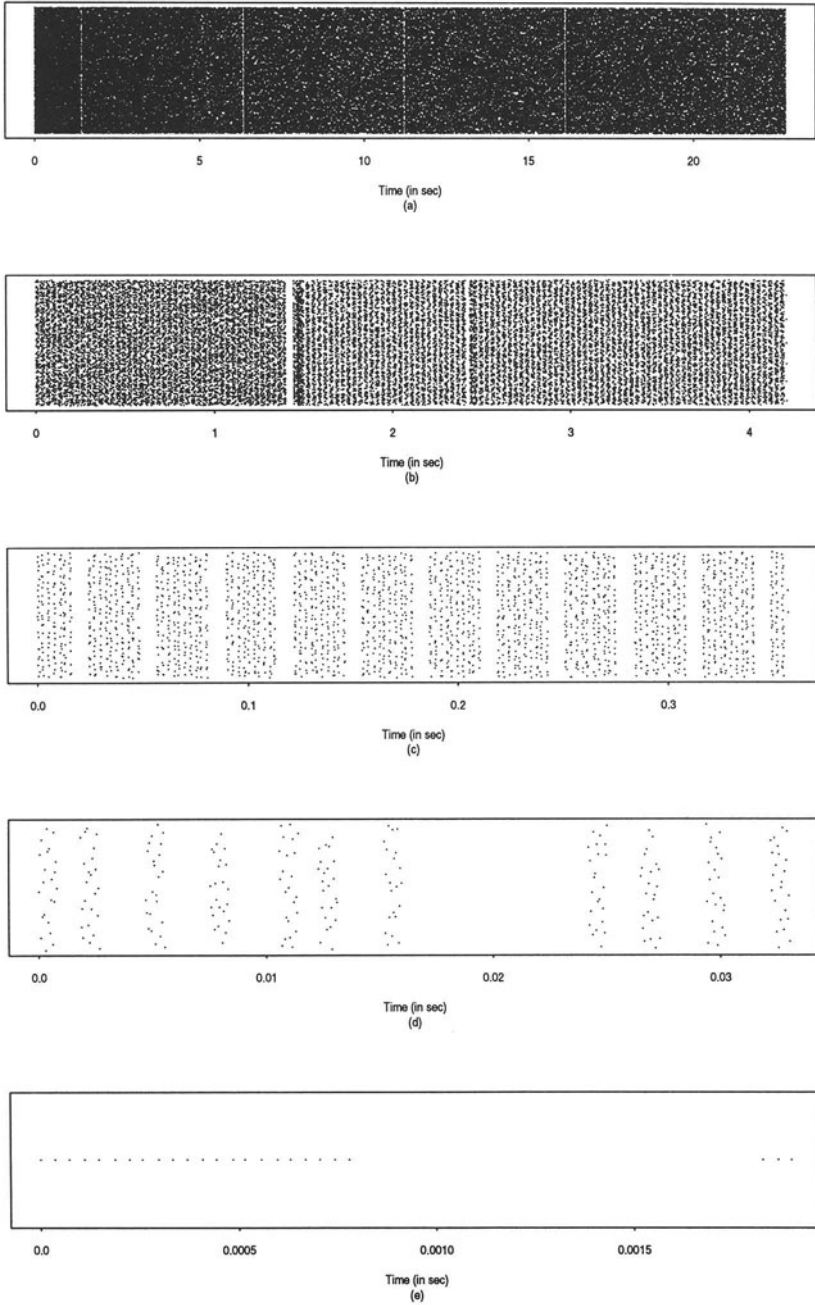


Figure 6 Textured plots of ATM traffic (motion JPEG, VCI=490) on 5 different time scales.

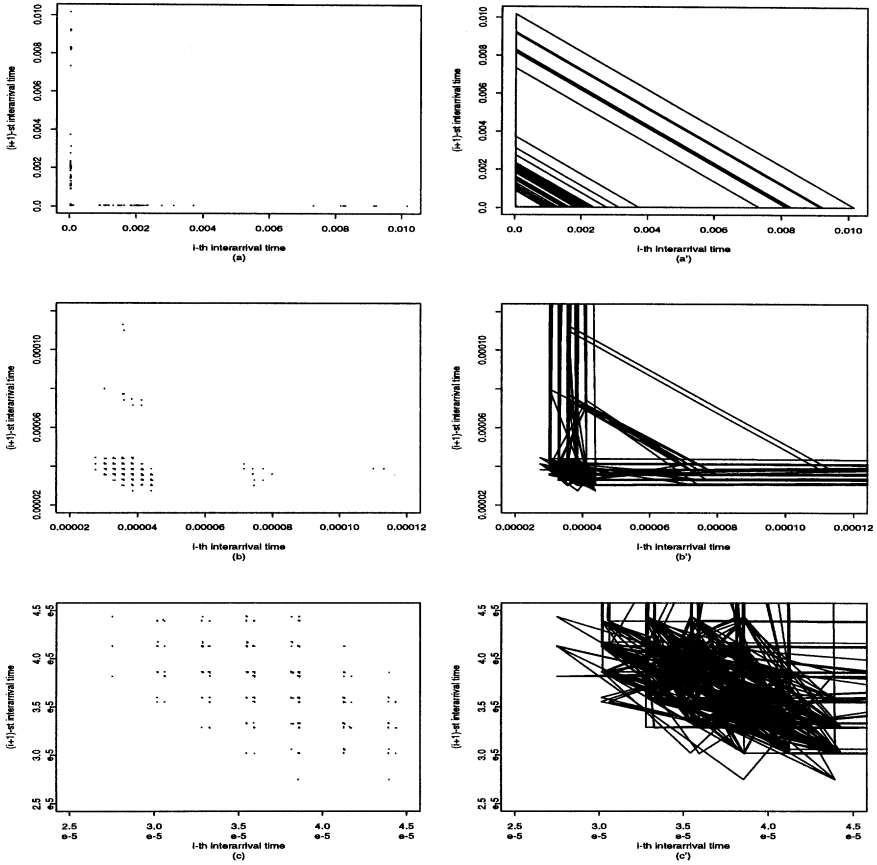


Figure 7 Phase plots (left) and temporal phase plots (right) for the interarrival times of 5000 ATM cells.

obtain for the first time insight into the nature of VBR video traffic at the ATM cell level – previous studies have typically been restricted to frame-level information.

Further details about the dynamics of motion JPEG at the ATM cell level are provided by the static and temporal versions of *phase plots* of successive cell interarrival times in Figure 7. Plots (a)-(c) in Figure 7 (left column) were generated using interarrival times of the first 5000 ATM cells; while the top plot identifies roughly 3 distinct clusters along each axis (corresponding to interarrival times of cells within the same block, between blocks, and between frames, respectively), the middle plot reveals further details about the structure of the interarrival times of ATM cells belonging to the same frame, and

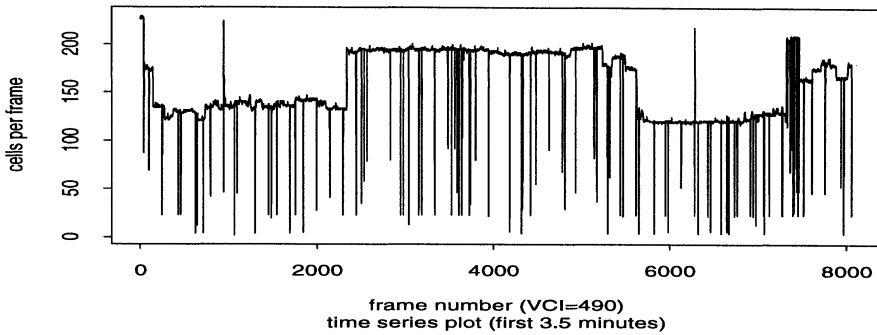


Figure 8 3.5 minutes of motion JPEG (VC 490) at the frame level.

when magnifying the apparent main cluster in the middle plot, Plot (c) in Figure 7 shows features that are reminiscent of the phase plots associated with the *netperf* application in Figure 3. The temporal phase plots in Plots (a')-(c') in Figure 7 (right column) provide further evidence for this observation and also indicates that successive cell interarrival times are, in fact, negatively correlated (i.e., larger interarrivals are typically followed by shorter ones).

Finally, we present some preliminary analysis that connects our empirical findings concerning motion JPEG at the ATM cell level with earlier VBR traffic characterization work at the frame level (see for example, Garrett *et al.* (1994) and Beran *et al.* (1995)). To this end, Figure 8 shows a time series plot obtained by identifying the frame boundaries in the raw ATM trace (for a duration of about 3.5 minutes worth of motion JPEG), counting the number of cells per frame, and plotting the resulting time series versus the frame number. The displayed trace resembles closely previously studied VBR traces (based, for example, on JPEG codecs), such as the sequence “Star Wars” (see Garrett *et al.* (1994)), or other motion picture sequences depicted in Beran *et al.* (1995). Moreover, a first analysis (not shown here) of the temporal dependence structure of the VC 490 motion JPEG sequence reveals the same long-range dependence feature that has been observed for the vast majority of previously measured VBR video traces and that is the underlying reason for the apparent fractal traffic characteristic observed in Figure 4, i.e., the bursty nature of the measured traffic over a wide range of time scales (ranging from milliseconds to seconds and beyond). Motion JPEG on VC 486 reveals a similar behavior, with slight variations, most likely due to different implementations of JPEG and possibly different ATM interface cards. VC 501 seems uninteresting from a traffic characterization perspective; after further investigations, it was found that user 501 forgot to turn on the VBR option and sent instead CBR-quality video.

Thus, a preliminary finding from this empirical work involving measured motion JPEG over ATM is that previously identified traffic characteristics

of VBR video at the frame level, e.g., long-range dependence, are preserved and extend below the frame level. At the same time, the ATM traces corresponding to the motion JPEG applications seen in this data set reveal a very definitive structure at the ATM cell level, and work is in progress that attempts to connect our improved cell-level understanding with previously obtained knowledge about VBR video at the frame and higher levels.

5 CONCLUSION

In the recent past, numerous studies that have relied crucially on the availability of high-resolution traffic traces from “live” networks, have provided significant and highly relevant new insight into the dynamics of modern communications networks. However, these studies were strictly limited to traditional legacy LANs and non-ATM networks. In this paper, we demonstrate that advances in data recording and storage technologies have made it possible to overcome this limitation. More specifically, we describe here an ATM traffic recorder that is capable of recording every bit from every cell on an OC-3 ATM link over extended periods in time; moreover, this is done in a fully non-intrusive manner, without losing a single bit, and by ensuring a time stamp accuracy of 50 nanoseconds for each ATM cell.

We also report on a traffic measurement experiment on BAGNet, where this ATM traffic recorder was successfully tested under field conditions and produced a data set of complete ATM traffic traces that is the first of its kind in the sense that it contains all information needed to study actual ATM traffic at all layers of interest. It is reasonable to expect that detailed analyses of this data set as well as of future traffic traces collected from other ATM testbeds will provide invaluable insights into problems related to the design, engineering and control of modern high-speed networks. To this end, future work that identifies the dynamics and performance of high-speed network traffic at the application level and across higher-level protocols will be of particular interest and importance.

Finally, a preliminary analysis (at the level of individual ATM cells) of a small subset of the traffic data collected from BAGNet has already provided strong evidence that traffic characteristics such as self-similarity and long-range dependence, which were originally observed in legacy LANs, will continue playing a crucial role in current and future ATM networks. The reason is simple and intuitively obvious: a significant fraction of the traffic seen on future ATM networks can be expected to be comprised of applications such as VBR video, LAN interconnection services and WWW; thus, these networks will soon carry traffic of which a substantial portion is perforce self-similar in nature, simply due to the source characteristics of the individual connections and applications. This property of modern broadband network traffic was originally hypothesized in Paxson *et al.* (1995), and our preliminary analysis of actual ATM traffic traces validates their claim empirically.

REFERENCES

- Berc, L., Devadhar, S. and Wiltzius, D. (1996) Experiences with ATM – BAG-Net. <http://www.llnl.gov/bagnet/article.html>.
- Beran, J., Sherman, R., Taqqu, M.S. and Willinger, W. (1996) Long-range dependence in variable-bit-rate video traffic. *IEEE Transactions on Communications*, **43**, 1566–1579.
- Crovella, M.E. and Bestavros, A. (1995) Self-similarity in world-wide-web traffic: Evidence and possible causes, in *Proceedings of the 1996 ACM Sigmetrics Conference*, Philadelphia, PA, 160–1696.
- Erramilli, A., Narayan, O. and Willinger, W. (1996) Experimental queueing analysis with long-range dependent packet traffic. *IEEE/ACM Transactions on Networking*, **4**(2), 209–223.
- Garrett, M.W. and Willinger, W. (1994) Analysis, modeling and generation of self-similar VBR video traffic, in *Proceedings of the ACM/Sigcomm '94*, London, UK, 269–280.
- Huang, C., Devetsikiotis, M., Lambadaris, I., and Kaye, A.R. (1995) Modeling and simulation of self-similar variable bit rate compressed video: A unified approach, in *Proceedings of the ACM/Sigcomm '95*, Cambridge, MA, 114–125.
- Laubach, M. (1994) Classical IP and ARP over ATM. RFC 1577, Internet Engineering Task Force.
- Leland, W.E., Taqqu, M.S., Willinger, W. and Wilson, D.V. (1994) On the self-similar nature of Ethernet traffic (Extended Version). *IEEE/ACM Transactions on Networking*, **2**, 1–15.
- Molnar, S. and Vidacs, A. (1997) On modeling and shaping self-similar ATM traffic. in *Teletraffic Contributions for the Information Age* (eds. R. Ramaswami and P.E. Wirth), Elsevier Science B.V., 1409–1430.
- Paxson, V. and Floyd, S. (1995) Wide area traffic: The failure of poisson modeling. *IEEE/ACM Transactions on Networking*, **3**, 226–244.
- Sullivan, M. (1995) Using Tribeca for network traffic analysis. *Bellcore Internal Document*, Morristown, NJ.
- Tukey, J.W. and Tukey, P.A. (1990) Strips displaying empirical distributions: Textured dot strips. *Bellcore Technical Memorandum*, Morristown, NJ.
- Willinger, W., Taqqu, M.S., Sherman, R. and Wilson, D.V. (1997) Self-similarity through high-variability: Statistical analysis of Ethernet LAN traffic at the source level. *IEEE/ACM Transactions on Networking*, **5**(1), 71–86.
- Winkelstein, D. and Stevenson, D. (1991) HiPPI link data analysis system: Test equipment for high-speed network analysis. in *Proceedings of TriComm '91*, Chapel Hill, NC.

6 BIOGRAPHIES

Walter Willinger received the Diplom (Dipl. Math.) degree from the ETH Zürich, Switzerland, and the M.S and Ph.D. degrees from the School of ORIE, Cornell University, Ithaca, NY, in 1984 and 1987, respectively. Before joining AT&T Labs-Research, Florham Park, NJ, he was a Member of Technical Staff at Bellcore Applied Research, Morristown, NJ. He is co-recipient of the 1996 IEEE W.R.G. Baker Prize Award and the 1994 IEEE Communications Society W.R. Bennett Prize Award for the paper titled "On the self-similar nature of Ethernet traffic."

Siddhartha Devadhar received a B.Tech. in Mechanical Engineering from the Indian Institute of Technology, Bombay, 1985, an M.E. in Robotics (Mech. Eng.) from Tulane University, New Orleans, 1987, and an M.S. in Computer Science from UCLA, Los Angeles, 1993. He is currently pursuing a Ph.D. in Network and Traffic Modeling at UCLA, while working at Pacific Bell, San Ramon, California. His major interests are inline measurement of broadband traffic, measurement based traffic modeling, trace driven simulations and performance analysis.

Andrew T. Heybey received his BS and MS in EECS from the Massachusetts Institute of Technology. He worked for several years at MIT before joining Bellcore in 1993. His interests include networks, operating systems, and computer architecture.

Robert Sherman received the Ph.D. degree in statistics from Yale University, New Haven, CT, in 1991. He is currently an Assistant Professor of Economics and Statistics at the California Institute of Technology, Pasadena, CA. Before that, he was a Member of Technical Staff at Bellcore Applied Research, Morristown, NJ, where he worked in the Statistics and Data Analysis Research Group.

Mark Sullivan received his Ph.D. in Computer Science from UC Berkeley in 1992. He did research in database management systems at Bellcore until late 1996. He is currently developing an electronic commerce system for Juno Online Services.

John R. Vollaro has 3 decades of experience in telecommunications research. He has been involved in the research and development of several network technologies, where he has always emphasized proof of concept through the construction of operational prototypes. Examples include the Spider and Datakit networks implemented at AT&T Bell Laboratories to demonstrate packet switching and virtual circuit techniques, and the Synton parallel processor at Bellcore, for which he holds 2 patents. His recent work includes the development of techniques used to capture and analyze traffic patterns from operational packet networks.