On the Self-Similar Nature of Ethernet Traffic
(extended version)

Will E. Leland
Bellcore
445 South Street
MRE 2L-245
Morristown, NJ 07960-6438
wel@bellcore.com

Murad S. Taqqu
Department of Mathematics
Boston University
Boston, MA 02215
murad@bu-ma.bu.edu

Walter Willinger†
Bellcore
445 South Street
MRE 2P-372
Morristown, NJ 07960-6438
+1 201 829-4297
walter@bellcore.com

Daniel V. Wilson
Bellcore
445 South Street
MRE 2L-257
Morristown, NJ 07960-6438
+1 201 829-4375
dvw@bellcore.com

†All correspondence concerning this paper should be directed to Daniel V. Wilson
or alternatively to Walter Willinger if Dan is not available.

March 3, 1993

Abstract

We demonstrate that Ethernet local area network (LAN) traffic is statistically *self-similar*, that none of the commonly used traffic models is able to capture this *fractal* behavior, that such behavior has serious implications for the design, control, and analysis of high-speed, cell-based (B-ISDN) networks, and that aggregating streams of such traffic typically intensifies the self-similarity ("burstiness") instead of smoothing it.

Intuitively, the critical characteristic of this self-similar traffic is that there is no natural length of a "burst": at every time scale ranging from a few milliseconds to minutes and hours, similar-looking traffic bursts are evident.

Our conclusions are supported by a rigorous statistical analysis of hundreds of millions of high quality Ethernet traffic measurements collected between 1989 and 1992, coupled with a discussion of the underlying mathematical and statistical properties of self-similarity and their relationship with actual network behavior. We also consider some implications for congestion control in B-ISDN and present traffic models based on self-similar stochastic processes. Self-similar traffic models provide simple, accurate, and realistic descriptions of traffic scenarios encountered during B-ISDN deployment.

1. INTRODUCTION

1.1 SUMMARY OF MAIN RESULTS

The main objectives of this paper are (i) to establish in a statistically rigorous manner the *self-similarity* characteristic or *fractal* nature of the high time-resolution Ethernet traffic measurements of Leland and Wilson (1991), (ii) to illustrate some of the most striking differences between self-similar models and the standard models for packet traffic currently considered in the literature, and (iii) to demonstrate some of the serious implications of self-similar network traffic for the design, control, and performance analysis of high-speed, cell-based communications systems.

Intuitively, self-similar phenomena display structural similarities across all (or at least a very wide range of) time scales. In the case of Ethernet LAN traffic, self-similarity is manifested in the absence of a natural length of a "burst"; at every time scale ranging from a few milliseconds to minutes and hours, bursts consist of bursty subperiods separated by less bursty subperiods. We also show that the degree of self-similarity (defined via the *Hurst parameter*) typically depends on the utilization level of the Ethernet and can be used to measure "burstiness" of LAN traffic. The term "self-similar" was coined by Mandelbrot. He and his co-workers (Mandelbrot and van Ness (1968), Mandelbrot and Wallis (1969a, 1969b), Mandelbrot and Taqqu (1979)) brought self-similar processes to the attention of statisticians, mainly through applications in such areas as hydrology and geophysics. For further applications and references on the probability theory of self-similar processes, see Mandelbrot (1971, 1983) and the extensive bibliography by Taqqu (1985). For an early application of the self-similarity concept to communications systems, see the seminal paper by Mandelbrot (1965).

Self-similar traffic behaves very differently than voice traffic or that predicted by packet traffic models currently considered in the literature, none of which are able to capture the fractal feature exhibited by the measurements. For example, our analysis of the Ethernet data shows that the generally accepted argument for the "Poisson-like" nature of aggregate traffic, namely, that aggregate traffic becomes smoother (less bursty) as the number of traffic sources increases, has very little to do with reality. In fact, the burstiness (degree of self-similarity) of LAN traffic typically intensifies as the number of active traffic sources increases, contrary to commonly held views.

Because of the growing market for LAN interconnection services, LAN traffic is rapidly becoming one of the major traffic contributors for broadband integrated services digital networks (B-ISDNs), the target information networks for future high-speed, high-bandwidth communications systems. Another expected major contributor to B-ISDN traffic is variable-bit-rate (VBR) video service. Since VBR traffic has recently been shown to display the same fractal-like property as LAN traffic (see Beran et al. (1992)), self-similar models provide simple, accurate, and realistic descriptions of traffic scenarios encountered during B-ISDN deployment.

In light of this new understanding of the nature of broadband traffic, we address in this paper also some of the serious implications of self-similar traffic on issues related to design, control, and performance analysis of high-speed, cell-based networks. As one specific example, we consider the area of congestion management and show that the nature of congestion produced by self-similar network traffic differs drastically from that predicted by traffic models currently considered in the literature and is far more complex than has been typically assumed in the past. As a result, proposed congestion control schemes for B-ISDN that work well under conventional traffic conditions typically perform less than satisfactorily under a self-similar traffic environment.

Finally, we present traffic models based on self-similar stochastic processes and mention a second, very recently proposed modeling approach based on deterministic nonlinear *chaotic maps*. We also address problems related to parameter estimation for self-similar traffic models and illustrate methods for easy and quick generation of synthetic traffic from these models.

---

§ The research was supported at Boston University by ONR grant N00014-90-J-1287

1.2 THE ETHERNET COMMUNICATIONS NETWORK

While readers not familiar with Ethernet communications systems may find the following description useful in combination with our analysis of the Ethernet traffic measurements presented in Section 4, network experts may safely skip this part and Section 1.3 below.

The Ethernet communications network is a broadcast, multiaccess system for local computer networking with distributed control. Ethernets have been in operation for over a decade and are one of the most popular local area network (LAN) technologies in use today, as well as one of the most successful. Among their attractive features are ease of maintenance and administration, ease of network reconfigurations (stations can be moved - disconnected from one point and reconnected at another - without the need to take down the network), access by a single, passive medium that is shared by all host stations, and the absence of a central controller allocating access to the channel. One of the major disadvantages of Ethernet local networks is their by now relatively slow speed of 10 Mbits/sec.

The approach used in the Ethernet system is the following (we closely follow the discussion in Shoch and Hupp (1980)). In the absence of a central controller allocating access to the channel, a *random access* procedure is used in which each station decides independently when to transmit its packet. If an Ethernet station wants to send a packet, a *carrier sense* technique is first applied and forces it to defer transmission if some other station is transmitting a packet. It will wait until the packet has passed before transmitting its own packet. Collisions can nevertheless occur, because two or more stations can sense the channel to be idle and start transmission simultaneously. However, since each sender continues to monitor the channel during transmission, it can provide collision detection when the signal on the channel does not match its own output. In that case, each station interrupts transmission, uses a collision consensus enforcement procedure to ensure that all other colliding stations have seen the collision, and then stops. Retransmission is then scheduled for some later time, but in order to avoid repeated collisions, each station waits for a random period of time before retransmitting. Furthermore, to avoid overloading the channel (i.e., making the system unstable), a *binary exponential backoff* algorithm is employed and guarantees that the range of the retransmission interval increases properly in times of heavy load of the system. Together, these mechanisms represent the Ethernet random access protocol, also referred to as *carrier sense multiple access with collision detection* or CSMA/CD.

1.3 THE HISTORY OF ETHERNET TRAFFIC MEASUREMENTS AND MODELING

One of the early myths about Ethernets was that "they work in practice, but allegedly not in theory" (see Boggs et al. (1988)). In particular, under certain assumptions (small packet length, large number of hosts) Ethernets have been shown theoretically (Metcalfe and Boggs, (1976)) to saturate at an offered load of 37%. Early Ethernet measurement studies by Shoch and Hupp (1980) and Boggs et al. (1988), showed, however, that this theoretical result has very little to do with reality; emphasising user-oriented measures of Ethernet behavior such as throughput and delay, they demonstrated that the Ethernet is capable of supporting its nominal capacity under realistic conditions and, moreover, that it allocates the bandwidth fairly. These studies typically used time measurements with low resolution (on the order of one second) and high loss rates (as much as 5% to 9% of all packets could not be monitored). In this context, a packet loss occurs each time the monitoring device fails to record an Ethernet packet; by artificially generating pathological packet streams, it is possible to precisely determine the recording capabilities of the measuring device (i.e., accuracy of the time stamps, frequency of lost packets). More recent studies, such as Feldmeier (1986), Gusella (1990), and Jain and Routhier (1986), have used higher time-resolution Ethernet measurements but still report losses between 1% to 5% of all Ethernet packets.

While early measurement studies of actual Ethernet implementations were motivated by the apparent discrepancy between theory and reality, the motivation for more recent studies such as the one discussed in this paper is different. Aggregate Ethernet LAN traffic provides a natural application to be served by the broadband integrated services digital network (B-ISDN), the high-speed, high-bandwidth packet switched public communication network of the future. B-ISDN is expected to offer a wide variety of new services exhibiting many different traffic characteristics and requiring many different quality-of-service specifications. In contrast to other expected applications such as video and medical imaging, LAN interconnection (i.e., interconnecting different local area networks through a connectionless service) is expected to become an immediate and major B-ISDN application, because the need for LAN interconnection already exists (due to the large number of LANs in operation). The expected high statistical multiplexing gain (due to observed high peak-to-mean ratios in LAN traffic) makes LAN interconnection an ideal candidate to offer as public service. In order to understand the

interactions between LANs and the proposed B-ISDN interconnection networks, a detailed knowledge of the behavior of the complete traffic between individual LANs over many time scales - ranging from milliseconds to seconds to minutes - is needed. Early measurement studies of Ethernet LAN traffic are inappropriate for this purpose because of their low time resolution and their inability to record packets during traffic bursts accurately and without losses. In the absence of realistic B-ISDN prototypes, measuring, analyzing, and understanding high time-resolution behavior of aggregate LAN traffic is essential to the design and performance analysis of LAN interconnection schemes and for the evaluation of proposed stochastic models of broadband traffic.

In this paper, we use very high quality, high time-resolution LAN traffic data collected by Leland and Wilson (1991). Using a custom-built Ethernet monitor described in the next section of the paper, they were able to record hundreds of millions of Ethernet packets without loss (irrespective of the traffic load), and with recorded time-stamps accurate to within 100 μs. The data were collected between August 1989 and February 1992 on several Ethernet LANs at the Bellcore Morristown Research and Engineering Center. Leland and Wilson (1991) present a preliminary statistical analysis of this unique high-quality data and comment in detail on the presence of "burstiness" across an extremely wide range of time scales: traffic "spikes" ride on longer-term "ripples", that in turn ride on still longer term "swells", etc. This *self-similar* or *fractal*-like behavior of aggregate Ethernet LAN traffic is very different both from conventional telephone traffic and from currently considered formal models for packet traffic (e.g., pure Poisson or Poisson-related models such as Poisson-batch or Markov-Modulated Poisson processes (Heffes and Lucantoni (1986)), packet-train models (Jain and Routhier (1986)), fluid flow models (Anick et al. (1982), etc.). As already observed in Ramaswami (1988), these differences require a new look at modeling traffic and performance of broadband networks.

1.4 OUTLINE OF THE PAPER

The paper is organized as follows. In Section 2, we first describe the available Ethernet traffic measurements and comment on the changes of the Ethernet population, applications, and environment during the measurement period from August 1989 to February 1992. (Network experts may again safely skip various parts of this section.) In Section 3, we give the mathematical definition of self-similarity, identify classes of stochastic models which are capable of accurately describing the self-similar behavior of the traffic measurements at hand, and discuss the implications of self-similarity from a modeling perspective.

Section 4 describes our statistical analysis of the Ethernet data, with emphasis on testing for self-similarity. We illustrate our statistical methods with a variety of different sets of Ethernet traffic data, taken at different times during the measurement period, with quite different user populations and gross traffic rates. We typically deal with time series with hundreds of thousands of observations and are, therefore, in the unique situation to rely on statistical results known to hold asymptotically (in the number of observations).

Finally, in Section 5 we discuss the significance of self-similarity for traffic engineering, and for operation, design, and control of B-ISDN environments. Among the implications discussed are (i) infinite variance source models for individual Ethernet users, (ii) inadequacies of commonly used notions of "burstiness", and (iii) better understanding of the nature of congestion for broadband network traffic. We conclude with a discussion of two different approaches for modeling self-similar network traffic.

2. TRAFFIC MEASUREMENTS

In this section, we give a brief description of the monitor used to gather the data for the analysis presented later, and discuss the network environment at the Bellcore Morris Research and Engineering Center (MRE) where the traffic measurements were collected. Our description of the types of traffic on the networks under examination and of the changes in the network during the data collection period will show that the MRE environment is rather typical of a research or software development environment. A longer discussion of the capabilities of the original monitoring system, including extensive testing of its capacity and accuracy can be found in Leland and Wilson (1991).

2.1 THE TRAFFIC MONITOR

There is a wide range of commercial and custom-built equipment available for monitoring or taking data from LANs. At one end of the spectrum, one finds equipment that processes packets as the packets are seen by the

monitor, usually doing at least some form of data reduction as each packet is seen. We term this *on-the-fly* analysis. This mode typically requires considerable *a priori* knowledge about what is interesting about a given network. At the other end of the spectrum is equipment that does absolutely no analysis of the LAN traffic, but simply records some or all of the contents of each packet as it is seen on the network. Since data reduction and analysis are done after the traffic data is gathered, we term this *retrospective* analysis. Most commercially available LAN monitoring equipment is more geared toward the *on-the-fly* rather than the *retrospective* analysis. Although the typical "LAN analyzer" is capable of some long-term real time analysis of a network and of some traffic capture and retrospective analysis, and although facilities for choosing the type packets to be captured are usually provided, the number of packets such an analyzer can capture is limited in number and rate over the long term. As one moves towards the *retrospective* mode, some aspects of the traffic can be recorded more accurately, most notably the arrival times for each packet, and the number of packets that can be captured in one run becomes significantly larger.

The monitoring system used to collect the data for the present study is custom built, records all packets seen on the Ethernet under study with accurate timestamps, and will do so for very long runs without interruption. Since we require the traffic analyst to make no *a priori* decisions as to what they are searching, other than how much of each packet is to be saved, our monitor is very much at the *retrospective* end of the spectrum. The monitor was custom-built by one of the authors (Wilson) in 1987/88 and has been in use to the present day with one upgrade. The following paragraph describes the updated version of the monitor; the original version is described at length in Leland and Wilson (1991). There is only one major difference between the two versions.

The heart of the monitor is a 68030 microprocessor-based single board computer (SBC) with on-board Ethernet hardware that is used as a high performance Ethernet interface. The SBC is programmed to gather the packets off the network being studied as they arrive, place them in buffers, and record a timestamp and status information for each packet. The SBC is connected to a Sun SPARCstation 1™ workstation via a set of bus adapters (the original monitor had the SBC sharing the bus with a Sun-3™ processor). A program running on the Sun processor polls the SBC periodically to see if a buffer is full; if so it is copied to disk. Another lower priority process running on the workstation copies the data on disk to one of two 8mm digital tape drives on the system. The memory on the SBC and the disk each form a layer of buffering between the relatively fast network and the slow 8mm tape drives, allowing the capture of packets from long periods of high network load without loss. The program driving the 8mm tape drives allows the traces to span multiple tapes without loss of data; over 27 million packets worth of data will fit on a single 2.3 Gbyte 8mm tape.

For each packet seen on the Ethernet under study, the monitor records a timestamp consisting of a 48-bit integer number of 4 μsec intervals since the start of the trace. This timestamp represents the arrival time of the end of the packet rather than the time the packet was placed on the Ethernet; the latter can easily be calculated. The monitor also records the packet length, the status of the Ethernet interface (which contains information about whether the packet is well-formed or whether packets were lost since data on the last one was taken), and finally the first 60 bytes of data in each packet. The system delivers timestamps accurate to within 20 μsec for the arrival time of each packet and was used to gather the most recent data set used in this study. (The older system delivered timestamps accurate to 100 μsec and was used in gathering all the other data sets used.) As we will show in Section 4, the high-accuracy timestamps of the Ethernet packets produced by this monitor are crucial for our statistical analyses of the data.

The use of a high-performance Ethernet interface (the SBC) is the reason the data is timestamped with high accuracy. There are a number of software programs and packages for use on workstations that use the workstation's Ethernet interface for monitoring, but these systems will not timestamp the packets as accurately as our monitor. The main reasons are that there is resource contention inherent in the I/O system of the workstation, there are devices that are handled at a higher priority than the Ethernet interface (such as the system clock), and sections of the operating system can cause processing of I/O to be deferred for significant periods. By using a separate high performance Ethernet interface only loosely coupled to the workstation, we avoid such problems.

2.2 THE NETWORK ENVIRONMENT AT BELLCORE

The MRE environment is probably typical of a research or software development environment where workstations are the primary machines on people's desks. It is also typical in that much of the original installation was well thought out and planned but then grew in a manner that can be described as haphazard at best. For the purposes of this study, this haphazard growth is not necessarily a liability, as we are able to study the traffic on a

network that is evolving over time. Table 1 gives a summary description of the traffic data analyzed later in the paper. We consider 4 sets of traffic measurements, each one representing between 20 and 40 consecutive hours of Ethernet traffic and each one consisting of tens of millions of Ethernet packets. The data were collected on different intracompany LAN networks at different periods in time over the course of approximately 4 years (August ’89, October ’89, January ’90, and February ’92). In order to fully understand where in the Bellcore network data was collected and what types of traffic were seen during each of the four measurement periods, it is necessary to first introduce some network-related terms. (Network experts may safely skip the next subsection.)

2.2.1 A FEW DEFINITIONS

The following terms are useful for describing and understanding the changes on a Ethernet LAN that is evolving over time: *router*, *bridge*, *repeater*, and *Ethernet segment*. A *router* uses a table to know where to distribute packets. The table may be static (e.g. typed in by an administrator) or dynamically updated via explicit distribution of routing information via the network. Many current routers allow filtering of packets based on the origin or destination address of the packet or the service the packet is meant to access. A router works on protocol layer addresses (e.g. Internet addresses) and port numbers, not Media Access Controller (MAC) addresses. Hence a router must "understand" all the protocols (such as the Internet protocol (IP), Apple’s Ethertalk™, or Novell’s IPX protocol) being sent through it. On the other hand, a *bridge* works on MAC addresses only and does not care about the upper layer protocols being used.

The algorithm a bridge uses is simple. Bridges can have any number of ports; we will use a 3 port bridge as an example. For a three-port bridge, we have ports 1, 2, and 3. When the bridge sees a packet from source host A on port 1, it makes a note in a table that A’s packet came in on port 1 and sends the packet out through all the remaining ports. When it sees a reply packet on port 3 from host B to destination host A, the bridge knows from the notes it has made that hosts A’s packets need only be sent out through port 1, *and* that hosts B’s packets need only be sent out through port 3. In this way a bridge "learns" the topology of the network, setting up a spanning tree so there is only one path between any two hosts. Bridges generally contain software to detect and handle loops in the topology of the network and to time out old entries in their tables, but this is not pertinent to the work being presented in this paper. A much more detailed explanation of a bridge may be found in Sincoskie (1986).

A *repeater* is a device that amplifies and relays Ethernet signals at the physical layer. It is often used to extend cables beyond the lengths dictated in the Ethernet specifications. Lastly, the terms *Ethernet segment* or just *segment* will be used in this paper to denote a cable or set of cables coupled by repeaters; every packet seen at one point of the segment is seen at every other point. The term *"logical Ethernet segment"* will sometimes be used to denote a group of Ethernet segments coupled by bridges. The important difference between the last two terms is that due to the learning nature of an Ethernet bridge, one will not see the exact same set of packets on one point of the logical Ethernet segment as is observed on another point of the logical segment. Finally, the term *network* is usually used to denote a group of Ethernet or other types of physical segments connected together with a combination of bridges and repeaters, and connected to other networks via routers. The hosts on a network usually have a protocol layer "network number" in common. An *internet* is a set of coupled networks usually covering an area that can range in size from a campus to a continent or larger.

2.2.2 WORKGROUP NETWORK TRAFFIC DATA

Four data sets will be considered in this paper. A summary description of these data sets is given in Table 1. The first two sets of traffic measurements, taken in August and October of 1989 (see first two rows in Table 1), were from an Ethernet network serving a laboratory of researchers engaged in everything from software development to prototyping new services for the telephone system. The traffic was mostly from services that used the Internet Protocol (IP) suite for such capabilities as remote login or electronic mail, and the Network File System™ (NFS) protocol for file service from servers to workstations. There were some unique services, though; for example, the audio of a local radio station was μ-law encoded and distributed over the network during portions of the day. While it is not our intent to provide here a detailed description of the particular MRE network segments under study, some words about the types of traffic on them are appropriate.

A reasonably accurate snapshot of the network configuration at the time of collection of the earliest data set being used (August ’89) does exist (see Figure 2.2.1): there were about 140 hosts and routers connected to this intra-laboratory network at that time, of which 121 spoke up during the 27 hour monitoring period. This network

Traces of Ethernet Traffic Measurements					
Measurement Period		data set	total number of bytes	total number of packets	Ethernet utilization
AUGUST 1989	total (27.45 hours)		11,448,753,134	27,901,984	9.3%
Start of trace: Aug. 29, 11:25am	low hour (6:25am-7:25am)	AUG89.LB AUG89.LP	224,315,439	652,909	5.0%
End of trace: Aug. 30, 3:10pm	normal hour (2:25pm-3:25pm)	AUG89.MB AUG89.MP	380,889,404	968,631	8.5%
	busy hour (4:25pm-5:25pm)	AUG89.HB AUG89.HP	677,715,381	1,404,444	15.1%
OCTOBER 1989	total (20.86 hours)		14,774,694,236	27,915,376	15.7%
Start of trace: Oct. 5, 11:00am	low hour (2:00am-3:00am)	OCT89.LB OCT89.LP	468,355,006	978,911	10.4%
End of trace: Oct. 6, 7:51am	normal hour (5:00pm-6:00pm)	OCT89.MB OCT89.MP	827,287,174	1,359,656	18.4%
	busy hour (11:00am-12:00am)	OCT89.HB OCT89.HP	1,382,483,551	2,141,245	30.7%
JANUARY 1990	total (40.16 hours)		7,122,417,589	27,954,961	3.9%
Start of trace: Jan. 10, 6:07am	low hour (Jan. 11, 8:32pm-9:32pm)	JAN90.LB JAN90.LP	87,299,639	310,038	1.9%
End of trace: Jan. 11, 10:17pm	normal hour (Jan. 10, 9:32am-10:32am)	JAN90.MB JAN90.MP	182,636,845	643,451	4.1%
	busy hour (Jan. 11, 10:32am-11:32am)	JAN90.HB JAN90.HP	711,529,370	1,391,718	15.8%
FEBRUARY 1992	total (47.91 hours)		6,585,355,731	27,674,814	3.1%
Start of trace: Feb. 18, 5:22am	low hour (Feb. 20, 1:21am-2:21am)	FEB92.LB FEB92.LP	56,811,435	231,823	1.3%
End of trace: Feb. 20, 5:16am	normal hour (Feb. 18, 8:21pm-9:21pm)	FEB92.MB FEB92.MP	154,626,159	524,458	3.4%
	busy hour (Feb. 18, 11:21am-12:21am)	FEB92.HB FEB92.HP	225,066,741	947,662	5.0%

Table 1. Qualitative description of the sets of Ethernet traffic measurements used in the analysis in Section 4.2.

consisted of two cable segments connected together with a bridge, implying that not all the traffic on the network as a whole was visible from our monitoring point. During the period this data was collected, among the 25 most active hosts were two DEC™ 3100 fileservers, one Sun-4™ fileserver, six Sun-3™ fileservers, two VAX™ 8650 minicomputers, and one CCI Power 6™ minicomputer. At that time, the less active hosts were mainly diskless Sun-3 machines and a smattering of Sun-4's, DEC 3100's, personal computers, and printers.

During the latter part of 1989 when the first two data sets were collected, a revolution was taking place on this network. The older Sun-3 class workstations were rapidly replaced with RISC-based workstations such as the SPARCstation-1™ and DEC 3100. Many of the new workstations were "dataless" (where the operating system is stored on a local disk but user data on a server) instead of "diskless" (where all files for the user and for the operating system are stored on a remote server). Because of the increased computing power of the machines connected to this segment, the network load increased appreciably, in spite of the trend towards dataless workstations. Note, for example, that the "busy hour" from the October ’89 data set is indeed busy: 30.7% utilization as compared to 15.1% during the August ’89 busy hour; similar increases can also be observed for the low and normal hours. Not long after this data was taken, this logical Ethernet segment was again segmented by adding yet a third cable and a bridge, and moving some user workstations and their fileserver to that new cable.

The above network has always been isolated from the rest of the Bellcore world by one or more routers. The other sides of these routers were connected to a large corporate internet consisting of many Ethernet segments and T-1 point-to-point links connected together with bridges at that time. Less than 5% of the total traffic on this workgroup network during either of the traces went out to either the rest of Bellcore or outside of the company.

2.2.3 WORKGROUP AND EXTERNAL TRAFFIC

The third data set, taken in January 1990 (row 3 in Table 1), came from an Ethernet cable that linked the two wings of the MRE facility that were occupied by a second laboratory (see Figure 2.2.2). At the time this data set was collected, this second laboratory comprised about 160 people, engaged in work similar to the first laboratory. This particular segment was unique in that it was also the segment serving Bellcore’s link to the outside Internet world. Thus the traffic on this cable was from several sources: (i) two very active file servers directly connected to the segment; (ii) traffic (file service and remote login) between the two wings of this laboratory, (iii) traffic between the laboratory and the rest of Bellcore, and (iv) traffic between Bellcore as a whole and the larger Internet world. This last type of traffic we term *external* traffic, and in 1990 could come from conversations between machines in any part of Bellcore and the outside world. This Ethernet segment was specifically monitored to capture this *external* traffic. We will be considering the aggregate and external traffic from this data set separately. This segment was separated from both the Bellcore internet and the two wings of the laboratory by bridges, and from the outside world by a vendor-controlled router programmed to pass anything with a Bellcore address as source or destination. In contrast to the two earlier data sets, over 1200 hosts spoke up during the 40 hour monitoring period on this segment.

2.2.4 BELLCORE AND THE MORRIS WORM

Between 1990 and 1992, the topology of the Bellcore internet changed considerably, driven significantly by factors having nothing to do with traditional measures such as segment overload and response time. Network security became a major concern due to the effects, both explicit and political, from the release of the "Morris worm". This "worm" is described in detail in Spafford (1989).

Before the release of the Morris worm, any non-Bellcore host could communicate with any host inside of Bellcore; there were no restrictions on external traffic of any kind (other than the bandwidth of the 1.5 Mbit T-1 link to the regional IP connectivity provider). Not long after the January 1990 data set was collected, a Bellcore-controlled security gateway router (usually called a *firewall*) was placed between Bellcore and the outside world. Because the flat Bellcore internet was becoming unmanageable due to broadcast "storms" and growth (from about 600 hosts total in 1988 to several thousand in 1992), routers were added between major workgroups and the Bellcore internet throughout the corporation. The router linking Bellcore to the outside world was placed directly on the Ethernet backbone of the MRE building rather than on a cable within the network domains of either of the laboratories discussed above. This backbone is part of the large Bellcore internet rather than a part of any particular workgroup or laboratory. The two laboratories mentioned previously were also connected to the MRE segment of the Bellcore internet via routers.

The firewall router was programmed to heavily restrict conversations initiated from *outside* of Bellcore. Telnet, remote login, and electronic mail are allowed only to a small number of machines within the corporation. Other services, such as the *finger* service which allows one to get a person’s real name, location, and so forth given their login name, are not allowed into Bellcore at all. By default, most of the rest of the standard IP-based Internet services are blocked at the security firewall.

The last data set, from February 1992 (see row 4 in Table 1), was taken from the building-wide Ethernet backbone in MRE after all the security measures had been put into place (see Figure 3.2.3). This cable carried all traffic going between laboratories within MRE, traffic from other Bellcore buildings destined for MRE, and all traffic destined for locations outside of Bellcore. Some hosts were still directly connected to this company-wide network in early 1992, but the trend to move them from the Bellcore internet to workgroup cables connected to the Bellcore internet via routers continues to the present. Because this cable had very little host to file server traffic, the overall traffic levels were much lower than for the other three sets. On the other hand, the percentage of remote login and mail traffic was higher. This cable also carried the digitized radio traffic between the two laboratories under discussion. The most radical difference between this data set and the others is that the traffic is primarily router to router rather than host to host. In fact, about 600 hosts spoke up during the measurement period (down from about 1200 active hosts during the January '90 measurement period), and the five most active hosts were routers.

3. SELF-SIMILAR STOCHASTIC PROCESSES

For a period of 27 consecutive hours of monitored Ethernet traffic (internal and external traffic) from the August 1989 measurements (first row in Table 1), Figure 3.1 depicts a sequence of simple plots of the packet arrival rate (i.e., number of packets per time unit) for 5 different choices of time units. Starting with a time unit of 100 seconds (a), each subsequent plot is obtained from the previous one by increasing the time resolution by a factor of 10 and by concentrating on a randomly chosen subinterval (as indicated by the darker shade). Recall that the time unit corresponding to the finest time scale is 10 milliseconds (e). Note that in order to avoid the visually irritating quantization effect associated with the finest resolution level, plot (e) depicts a "jittered" version of the number of packets per 10 milliseconds, i.e., a small amount of noise has been added to the actual arrival rate. With the possible exception of plot (a) (the presence of a daily cycle is clearly visible), all other plots look intuitively and in a distributional sense very "similar" to one another and are distinctively different from white noise (i.e., an independent and identically distributed sequence of random variables). This scale-invariant or "self-similar" feature of Ethernet traffic is drastically different from both conventional telephone traffic and from currently considered stochastic models for packet traffic, and motivates the use of self-similar stochastic processes for traffic modeling purposes. The presentation below of the concept of self-similar processes closely follows Cox (1984).

3.1 DEFINITION

Let $X = (X_t: t = 0, 1, 2, ...)$ be a *covariance stationary* (sometimes called *wide-sense stationary*) stochastic process; that is, a process with constant mean $\mu = E[X_t]$, finite variance $\sigma^2 = E[(X_t - \mu)^2]$, and an autocorrelation function $r(k) = E[(X_t - \mu)(X_{t+k} - \mu)]/E[(X_t - \mu)^2]$ $(k = 0, 1, 2, ...)$ that depends only on k. In particular, we assume that X has an autocorrelation function of the form

$$r(k) \sim k^{-\beta} L_1(k), \quad as\ k \to \infty, \tag{3.1.1}$$

where $0 < \beta < 1$ and L_1 is slowly varying at infinity, that is, $\lim_{t\to\infty} L_1(tx)/L_1(t) = 1$ for all $x > 0$ (examples of such slowly varying functions are $L_1(t) = const$, $L_1(t) = \log(t)$). For each $m = 1, 2, 3, \cdots$, let $X^{(m)} = (X_k^{(m)}: k = 1, 2, 3, ...)$ denote a new time series obtained by averaging the original series X over non-overlapping blocks of size m. That is, for each $m = 1, 2, 3, \cdots$, $X^{(m)}$ is given by

$$X_k^{(m)} = 1/m\ (X_{km-m+1} + \cdots + X_{km}), \quad k = 1, 2, 3, \cdots \tag{3.1.2}$$

Note that for each m, the aggregated time series $X^{(m)}$ defines a covariance stationary process; let $r^{(m)}$ denote the corresponding autocorrelation function.

The process X is called *(exactly second-order) self-similar* with self-similarity parameter $H = 1 - \beta/2$ if the corresponding aggregated processes $X^{(m)}$ have the same correlation structure as X, i.e.,

$$r^{(m)}(k) = r(k), \quad \text{for } all\ m = 1, 2, \cdots \quad (k = 1, 2, 3, ...). \tag{3.1.3}$$

In other words, X is exactly self-similar if the aggregated processes $X^{(m)}$ are indistinguishable from X - at least with respect to their second order properties. X is called *(asymptotically second-order) self-similar* with self-similarity parameter $H = 1 - \beta/2$ if

$$r^{(m)}(1) \to 2^{1-\beta} - 1, \quad as\ m \to \infty, \tag{3.1.4}$$
$$r^{(m)}(k) \to 1/2\ \delta^2(k^{2-\beta}), \quad as\ m \to \infty \quad (k = 2, 3, ...),$$

where $\delta^2(f)$ denotes the second central difference operator applied to a function f, i.e., $\delta^2(f(k)) = f(k+1) - 2f(k) + f(k-1)$. Thus, an asymptotically self-similar process has the property that for large m, the corresponding aggregated time series $X^{(m)}$ have a fixed correlation structure, solely determined by β; moreover, due to the asymptotic equivalence (for large k) of differencing and differentiating, $r^{(m)}$ agrees asymptotically with the correlation structure of X given by (3.1.1).

Intuitively, the most striking feature of (exactly or asymptotically) self-similar processes is that their aggregated processes $X^{(m)}$ possess a nondegenerate correlation structure as $m \to \infty$. This behavior is in stark contrast to the more conventional stochastic models, all of which have the property that their aggregated processes $X^{(m)}$ tend to second order pure noise (as $m \to \infty$), i.e.,

$$r^{(m)}(k) \to 0, \quad as\ \ m \to \infty, \ \ (k = 1, 2, 3, ...). \tag{3.1.5}$$

The above-mentioned intuition is best illustrated with the sequence of plots in Figure 3.1; if the original time series X represents the number of Ethernet packets per 10 milliseconds (plot (e)), then plots (a) to (d) depict segments of the aggregated time series $X^{(10000)}$, $X^{(1000)}$, $X^{(100)}$, and $X^{(10)}$, respectively. All of the plots look "similar" and distinctively different from pure noise.

Note that we have chosen the above definitions of self-similarity over the mathematically more convenient definition of a *self-similar* continuous-time stochastic process $X = (X_t: t \geq 0)$ with stationary increments, namely, for all $a > 0$,

$$X_{at} = a^H X_t, \tag{3.1.6}$$

where equality is understood in the sense of equality of the finite-dimensional distributions, and the exponent H is the self-similarity parameter. Definitions (3.1.4) and (3.1.5) have the advantage that they do not obscure the connection with standard time series theory, and they reflect the fact that we are mainly interested in large m's (time "scales"); here we are less concerned about deviations from self-similarity for $m \to 0$. From a modeling perspective, the crucial point is that both the discrete-time and the continuous-time definitions involve a wide range of time scales.

3.2 STOCHASTIC MODELING OF SELF-SIMILAR PHENOMENA

Representations that view a process with a correlation structure of the form (3.1.1) as a (finite approximation of a) continuous sum of Gauss-Markov processes are often interpreted as suggesting the presence of a multilevel hierarchy of underlying mechanisms to account for self-similarity (see Mandelbrot and Taqqu (1979)). For examples using such explicit hierarchies in a physical context, see Cassandro and Jona-Lasinio (1978), and, in a textile application, see Cox and Townsend (1948). In general, however, it is very difficult to demonstrate the physical existence of a multilevel of hierarchies of underlying mechanisms as well as to show why its presence should contribute to self-similar behavior. As a result, formal mathematical models have been introduced that yield elegant representations of the self-similarity phenomenon, mostly though without providing any physical interpretation. Below, we present two such models, the exactly self-similar fractional Gaussian noise and the class of asymptotically self-similar fractional autoregressive integrated moving-average (ARIMA) processes. Then we also present a construction (originally introduced by Mandelbrot (1969) and later extended by Taqqu and Levy (1986) and Levy and Taqqu (1987)) of self-similar processes based on aggregating many simple renewal reward processes exhibiting inter-renewal times with infinite variances. This last construction can be viewed as an attempt to "explain" self-similarity in terms of representing some underlying physical process; it is particularly appealing in the context of high-speed packet traffic (see Section 5).

3.2.1 FRACTIONAL GAUSSIAN NOISE

A *fractional Gaussian noise* $X = (X_k: k = 0, 1, 2, ...)$ is a stationary Gaussian process with mean $\mu = E[X_k]$, variance $\sigma^2 = E[(X^k - \mu)^2]$, and autocorrelation function

$$r(k) = 1/2\ (|k+1|^{2H} - |k|^{2H} + |k-1|^{2H}), \quad k = 1, 2, 3, \cdots \tag{3.2.1}$$

It is easy to see that, asymptotically, $r(k) \sim H(2H-1)|k|^{2H-2}$ $(k \to \infty,\ 0 < H < 1)$; also, simple calculations

show that the resulting aggregated processes $X^{(m)}$ ($m = 1, 2, 3, ...$) all have the same distribution as X, for all $0 < H < 1$. Thus, by (3.1.1) and (3.1.3), fractional Gaussian noise is exactly second-order self-similar with self-similarity parameter H, as long as $1/2 < H < 1$.

Fractional Gaussian noise was introduced by Mandelbrot and Van Ness (1968) and has been especially popular in hydrological modeling (e.g., see Mandelbrot and Wallis (1969a, 1969b), and Lawrence and Kotegoda (1977)). Despite its rigid correlation structure (3.2.1), fractional Gaussian noise can be viewed as an idealization or as a reasonable first approximation of a more complex structure because if one were to subject certain long-range dependent processes (see Section 3.3) to a special type of central limit theorem, one would obtain in the limit fractional Gaussian noise. Methods for estimating the three unknown parameters μ, σ^2, and H are known and will be addressed in Section 4.

3.2.2 FRACTIONAL ARIMA(p,d,q) PROCESSES

A *fractional ARIMA(p,d,q) process*, where p and q are non-negative integers and d is real, is defined to be a stochastic process $X = (X_k: k = 0, 1, 2, ...)$ with a representation given by

$$\Phi(B)\Delta^d X_k = \Theta(B)\varepsilon_k, \tag{3.2.2}$$

where $\Phi(B) = 1 - \phi_1 B - \cdots - \phi_p B^p$ and $\Theta(B) = 1 - \theta_1 B - \cdots - \theta_q B^q$ are polynomials in the backward-shift operator $B\,X_k = X_{k-1}$, $\Delta = 1 - B$ denotes the differencing operator, and Δ^d is the fractional differencing operator defined by $\Delta^d = (1 - B)^d = \sum_k \binom{d}{k} (-B)^k$ with $\binom{d}{k}(-1)^k = \Gamma(-d + k) / (\Gamma(-d)\Gamma(k + 1))$ and $(\varepsilon_k: k = 0, 1, 2, ...)$ is a white noise process. It has been shown in Granger and Joyeux (1980) that for $d \in (-1/2, 1/2)$, X is stationary and invertible, and its autocorrelations satisfy $r(k) \sim ak^{2d-1}$ *as* $k \to \infty$, where a is a finite positive constant independent of k. Moreover, Cox (1984) showed that the aggregated time series $X^{(m)}$ of a fractional ARIMA(p,d,q) process satisfy (3.1.4) for $-1/2 < d < 1/2$. Thus, relations (3.1.1) and (3.1.4) hold and X is asymptotically second-order self-similar with self-similarity parameter $d + 1/2$, for all $0 < d < 1/2$.

Fractional ARIMA(p,d,q) processes were introduced by Granger and Joyeux (1980) and Hosking (1981). Their correlations at high lags are similar to those of an ARIMA($0,d,0$) process with the same value of d. ARIMA($0,d,0$) are the simplest and most fundamental of the fractionally differenced ARIMA processes and have a representation of the form $\Delta^d X_k = \varepsilon_k$ ($k = 0, 1, 2, ...$), i.e., an infinite order autoregressive representation. The corresponding infinite order moving average representation $X_k = \Delta^{-d}\varepsilon_k$ shows that ARIMA($0,d,0$) processes are obtained by subjecting white noise to fractional differencing of order $-d$. Upon setting $H = d + 1/2$, both fractional Gaussian noise and the ARIMA($0,d,0$) process have correlations which behave asymptotically as k^{2d-1} (with different constants of proportionality). From the point of view of time series modeling, one of the main advantages of the ARIMA($0,d,0$) family over fractional Gaussian noise processes is that the former can be combined in a natural way with the established class of Box-Jenkins models (Box and Jenkins (1976)), resulting in the family of ARIMA(p,d,q) processes. Fractional ARIMA processes are much more flexible with regard to the simultaneous modeling of the short-term and long-term behavior of a time series than fractional Gaussian noise, mainly because the latter, having only the three parameter μ, σ^2, and H, is not capable of capturing the wide range of low-lag correlation structures encountered in practice. This flexibility can already be observed when considering the simplest processes of the ARIMA(p,d,q) family, namely the two-parameter models ARMA($1,d,0$) and ARMA($0,d,1$) (see Hosking (1981)).

3.2.3 SELF-SIMILARITY THROUGH AGGREGATION

Let $U_0, U_1, U_2, \cdots$ be a sequence of i.i.d. integer-valued random variables ("inter-renewal times") with

$$P[U \geq u] \sim u^{-\alpha} h(u), \quad as \;\; u \to \infty,\; 1 < \alpha < 2, \tag{3.2.3}$$

where h is slowly varying at infinity. For example, the stable (Pareto) distribution with parameter $1 < \alpha < 2$ satisfies the "heavy-tail behavior" or "longtailedness property" (3.2.3). Furthermore, let $W_0, W_1, W_2, \cdots$ be an i.i.d. sequence ("rewards") with $E[W] = 0$, $E[W^2] < \infty$, independent of the U's. In order to obtain a stationary (delayed) renewal sequence $(S_k)_{k \geq 0}$, we take S_0 independent of the U_k's and distributed as

$$P[S_0 = u] = \mu^{-1} P[U \geq u + 1],\;\; u = 0, 1, 2, ..., \tag{3.2.4}$$

where $\mu = E[U]$; then for $k \geq 1$, we set $S_k = S_0 + \sum_{j=1}^{k} U_j$. The renewal reward process

$W = (W(t)\colon t = 0,\ 1,\ 2,\ \ldots)$ is then given by

$$W(t) = \sum_{k=0}^{t} W_k I_{(S_{k-1},\ S_k]}(t) \tag{3.2.5}$$

with $I_A(\cdot)$ denoting the indicator function of the set A. Notice that W is stationary in the sense that its finite-dimensional distributions are invariant under time shifts. A sample path of W is simply a sequence of points located on segments parallel to the x-axis at a level W_k when $t \in (S_{k-1},\ S_k]$.

By aggregating M i.i.d. copies $W^{(1)}$, $W^{(2)}$, ..., $W^{(M)}$ of W, we obtain the model of interest, namely the process W^* given by

$$W^*(T,\ M) = \sum_{t=1}^{T} \sum_{m=1}^{M} W^{(m)}(t). \tag{3.2.6}$$

with $W^*(0,\ M) = 0$. In the economic framework originally considered in Mandelbrot (1969), T denotes time, M is a model index (level of aggregation), and W^* models a commodity price. Such a modeling approach reflects the belief that many microeconomic variables are aggregates of a very large number of micro-variables. Mandelbrot (1969) and Taqqu and Levy (1986) determined the combined effect of aggregating an increasing number of copies of W and accumulating them over larger and larger time intervals. In particular, they showed that for T and M both large with $T << M$, W^* behaves like fractional Brownian motion, that is, properly normalized, $W^*(T,\ M)$ converges (in the sense of convergence of the finite-dimensional distributions) to the integrated version of fractional Gaussian noise, i.e., a mean-zero Gaussian process $B_H = (B_H(s)\colon s \geq 0)$, $1/2 < H < 1$, with correlation function

$$R(s,\ t) = 1/2\ (s^{2H} + t^{2H} - |\ s - t\ |^{2H}). \tag{3.2.7}$$

B_H has stationary increments and is self-similar (in the sense of (3.1.6)) with self-similarity parameter H. For more details concerning fractional Brownian motion, see Mandelbrot and Van Ness (1968), and Mandelbrot and Taqqu (1979). As an immediate consequence of Taqqu and Levy's result, we have that for T and M both large with $T << M$, the increment process of W^* behaves like fractional Gaussian noise.

Producing self-similarity by aggregating more and more i.i.d. copies of the rather elementary renewal reward process W over longer and longer time periods relies crucially on the heavy tail behavior (3.2.3) of the inter-renewal times U. Because of property (3.2.3), the process W can assume with high probability the same value (namely W_k) over long periods of time. Aggregating a large number of copies of W results in an overall sum that becomes more and more Gaussian, and evaluating these copies of W over longer and longer time periods introduces significant dependence, attributes shared by fractional Brownian motion. For a construction that is similar in nature but uses certain AR(1)-processes instead of the renewal reward processes W and results in asymptotic self-similarity (in the limit as M and T increase), see Granger (1980). These constructions are of practical interest since they provide an intuitive explanation for the occurrence of self-similarity in high-speed network traffic (see Section 5.1). A rather different construction of self-similar processes as limits of integrated shot noise processes is discussed in Giraitis and Surgailis (1991); however, its interpretation in the context of telecommunication systems is less obvious.

3.3 PROPERTIES OF SELF-SIMILAR PROCESSES

3.3.1 LONG-RANG DEPENDENCE AND THE HURST EFFECT

A stochastic process satisfying relation (3.1.1) is said to exhibit *long-range dependence* (see, for example, Beran (1992), Cox (1984), Kuensch (1986), or Taqqu (1987)). Thus, processes with long-range dependence are characterized by an autocorrelation function that decays hyperbolically as the lag increases. Moreover, it is easy to see that (3.1.1) implies $\sum_k r(k) = \infty$. This non-summability of the correlations captures the intuition behind long-range dependence, namely that while high-lag correlations are all individually small, their cumulative effect is of importance and gives rise to features which are drastically different from those of the more conventional, i.e., short-range dependent processes. The latter are characterized by an exponential decay of the autocorrelations, i.e., $r(k) \sim \rho^k,\ \ as\ \ k \to \infty$ $(0 < \rho < 1)$, resulting in a summable autocorrelation function $0 < \sum_k r(k) < \infty$.

When working in the frequency domain, long-range dependence manifests itself in a spectral density that obeys a power-law behavior near the origin. In fact, equivalently to (3.1.1) (under weak regularity conditions on the slowly varying function L_1), there is long-range dependence in X if

$$f(\lambda) \sim \lambda^{-\gamma} L_2(\lambda), \quad as \;\; \lambda \to 0, \tag{3.3.1}$$

where $0 < \gamma < 1$, L_2 is slowly varying at 0, and $f(\lambda) = \sum_k r(k) e^{ik\lambda}$ denotes the spectral density function. Thus, from the point of view of spectral analysis, long-range dependence implies that $f(0) = \sum_k r(k) = \infty$, that is, it requires a spectral density which tends to $+\infty$ as the frequency λ approaches 0 ("1/*f*-noise"). On the other hand, short-range dependence is characterized by a spectral density function $f(\lambda)$ which is positive and finite for $\lambda = 0$.

From our earlier discussion, it follows that both fractional Gaussian noise processes (with $1/2 < H < 1$) and fractional ARIMA(p,d,q) processes (with $0 < d < 1/2$) exhibit long-range dependence. The parameters H and d, respectively, measure the degree of long-range dependence and can be estimated from empirical records (see Section 4). Heuristically, long-range dependence manifests itself in the presence of cycles of all frequencies and orders of magnitude, displays features suggestive of nonstationarity, and has been found to be relevant in economics (Lo (1991), Mandelbrot (1969), Sowell (1990)), in hydrology and geology (Mandelbrot and Van Ness (1968), Mandelbrot and Wallis (1969a, 1969b)), and in telecommunication (Mandelbrot (1965), Beran et al. (1992)) - see also Beran (1992), Cox (1984), Hampel (1987), and Taqqu (1985) and the references therein.

Historically, the importance of self-similar processes as defined in Section 3.1 lies in the fact that they provide an elegant explanation and interpretation of an empirical law that is commonly referred to as *Hurst's law* or the *Hurst effect*. Briefly, for a given set of observations $(X_k: k = 1, 2, ..., n)$ with sample mean $\bar{X}(n)$ and sample variance $S^2(n)$, the *rescaled adjusted range* or the *R/S statistic* is given by

$$R(n)/S(n) = 1/S(n) \; [\max(0, W_1, W_2, ..., W_n) - \min(0, W_1, W_2, ..., W_n)], \tag{3.3.2}$$

with $W_k = (X_1 + X_2 + \cdots + X_k) - k\bar{X}(n)$, $k = 1, 2, ..., n$. Hurst (1951, 1955) found that many naturally occurring time series appear to be well represented by the relation

$$E[R(n)/S(n)] \sim cn^H, \quad as \;\; n \to \infty, \tag{3.3.3}$$

with *Hurst parameter* H "typically" about 0.73, and c a finite positive constant that does not dependent on n. On the other hand, if the observations X_k come from a short-range dependent model, then Mandelbrot and Van Ness (1968) showed that

$$E[R(n)/S(n)] \sim dn^{0.5}, \quad as \;\; n \to \infty, \tag{3.3.4}$$

(d a finite positive constant, independent of n). The discrepancy between (3.3.3) and (3.3.4) is generally referred to as the *Hurst effect* or *Hurst phenomenon*.

3.3.2 SLOWLY DECAYING VARIANCES

From a statistical point of view, the most salient feature of self-similar processes as defined in Section 3.1 is that the variance of the arithmetic mean decreases more slowly than the reciprocal of the sample size; that is, it behaves like $n^{-\beta}$ for some $\beta \in (0, 1)$, instead of like n^{-1} for the processes whose aggregated series converge to second-order pure noise. For our discussion below, we assume for simplicity that the slowly varying functions L_1 and L_2 in (3.1.1) and (3.3.1), respectively, are asymptotically constant. Cox (1984) showed, in fact, that a specification of the autocorrelation function satisfying (3.1.1) (or equivalently, of the spectral density function satisfying (3.3.1)) is the same as a specification of the sequence $(var(X^{(m)}): m \geq 1)$ with the property

$$var(X^{(m)}) \sim am^{-\beta}, \quad as \;\; m \to \infty, \tag{3.3.5}$$

where a is a finite positive constant independent of m, and $0 < \beta < 1$; in fact, the parameter β is the same as in (3.1.1) and is related to the parameter γ in (3.3.1) by $\beta = 1 - \gamma$. On the other hand, for covariance stationary processes whose aggregated series $X^{(m)}$ tend to second-order pure noise (i.e., (3.1.5) holds), it is easy to see that the sequence $(var(X^{(m)}): m \geq 1)$ satisfies

$$var(X^{(m)}) \sim bm^{-1}, \quad as \;\; m \to \infty, \tag{3.3.6}$$

where b is a finite positive constant independent of m.

The consequences of the slow decaying variances $var(X^{(m)})$ for classical statistical tests and confidence or prediction intervals can be disastrous (e.g., see Beran (1992) and Hampel et al. (1986)), since usual standard errors (derived for conventional models) are wrong by a factor that tends to infinity as the sample size increases!

3.3.3 PARSIMONIOUS MODELING

Since we are always dealing with finite data sets, it is in principle not possible to decide whether the asymptotic relationships (3.1.1), (3.1.3), (3.1.4), etc. hold or not. For processes that are not self-similar in the sense that their aggregated series converge to second-order pure noise (property (3.1.5)), the correlations will eventually decrease exponentially, continuity of the spectral density function at the origin will eventually show up, the variances of the aggregated processes will eventually decrease as m^{-1}, and the rescaled adjusted range will eventually increase as $n^{0.5}$. For finite sample sizes, distinguishing between these asymptotics and the ones corresponding to self-similar processes is, in general, problematic.

In the present context of Ethernet measurements, we typically deal with time series with hundreds of thousands of observations and are, therefore, in the unique situation to investigate "close-to" asymptotic behavior of quantities such as the rescaled adjusted range or the variance of the aggregated processes. Moreover, with such sample sizes, parsimonious modeling becomes a necessity due to the large number of parameters needed when trying to fit a conventional process to a "truly" self-similar model. Modeling, for example, long-range dependence with the help of ARMA processes is equivalent to approximating a hyperbolically decaying autocorrelation function by a sum of exponentials. Although always possible, the number of parameters needed will tend to infinity as the sample size increases, and giving physically meaningful interpretations for the parameters becomes more and more difficult. In contrast, the long-range dependence component of the process can be modeled (by a self-similar process) with only one parameter! Finally, from a modeling perspective, it would be very unsatisfactory to use for a single empirical time series two different models, one for a short sequence, another one for a long sequence.

4. STATISTICAL ANALYSIS OF ETHERNET TRAFFIC MEASUREMENTS

In this section, we establish in a statistically rigorous manner the self-similar nature of Ethernet traffic and of some of its major components (e.g., external traffic, external TCP traffic). To this end, we describe the necessary statistical and graphical tools, and illustrate their use on a variety of different sets of Ethernet traffic measurements taken during a four-year period from 1989 to 1992. The significance of our findings from a traffic engineering point of view will be discussed in Section 5. For further evidence for the omnipresence of self-similarity and long-rang dependence in real-time traffic measurements of high-speed telecommunications systems, see Beran et al. (1992), Erramilli and Willinger (1993), and Garrett and Willinger (1993).

4.1 STATISTICAL METHODS FOR TESTING FOR SELF-SIMILARITY

From a theoretical point of view, slowly decaying variances, long-range dependence, and a spectral density of the form (3.3.1) are different manifestations of one and the same property of the underlying covariance stationary process X, namely that X is (asymptotically or exactly) self-similar. Subsequently, we can approach the problem of testing for and estimating the degree of self-similarity form three different angles: (1) analysis of the variances of the aggregated processes $X^{(m)}$, (2) time-domain analysis based on the R/S-statistic, and (3) periodogram-based analysis in the frequency-domain. This subsection provides a brief description of the corresponding statistical and graphical methods (for more details, see Beran et al. (1992) and the references therein); their use in analyzing the Ethernet data will be illustrated in subsections 4.2 and 4.3 below.

4.1.1 VARIANCE-TIME PLOTS

We have seen in Section 3.3.2 that for self-similar processes, the variances of the aggregated processes $X^{(m)}$ ($m = 1, 2, 3, \cdots$) decrease linearly (for large m) in log-log plots against m with slopes arbitrarily flatter than -1 (see (3.3.5). On the other hand, none of the short-range dependent processes (e.g., MMPP, fluid models, ARMA models) yield a power-law for the variances of the form (3.3.5); it can be approximated for some transient period of time by short-range dependent models with a large number of parameters, but the variance of $X^{(m)}$ will eventually decrease linearly in log-log plots against m with a slope equal to -1 (see (3.3.6)). The so-called *variance-time plots* are obtained by plotting $\log(\mathrm{var}(X^{(m)}))$ against $\log(m)$ ("time") and by fitting a simple least squares line through the resulting points in the plane, ignoring the small values for m. Values of the estimate $\hat{\beta}$ of the asymptotic slope between -1 and 0 suggest self-similarity, and an estimate for the degree of self-similarity is given by $H = 1 - \hat{\beta}/2$.

Clearly, variance-time plots are not reliable for empirical records with small sample sizes. However, as we will demonstrate below, with sample sizes of the magnitude of the Ethernet traffic data sets, "eyeball tests" such as the variance-time plots become highly useful and give a rather accurate picture about the self-similar nature of the underlying time series and about the degree of self-similarity.

4.1.2 R/S ANALYSIS

The objective of the R/S analysis of an empirical record is to infer the degree of self-similarity H (Hurst parameter) in relation (3.3.3) for the self-similar process that presumably generated the record under consideration. In practice, R/S analysis is based on a heuristic graphical approach (originally described in detail in Mandelbrot and Wallis (1969b)) that tries to exploit as fully as possible the information in a given record. The following graphical method has been used extensively in the past. Given a sample of N observations $(X_k : k = 1, 2, 3, \cdots, N)$, one subdivides the whole sample into K non-overlapping blocks and computes the rescaled adjusted range $R(t_i, n)/S(t_i, n)$ for each of the new "starting points" $t_1 = 1, t_2 = N/K + 1, t_3 = 2N/K + 1, \cdots$ which satisfy $(t_i - 1) + n \le N$. Here, the R/S-statistic $R(t_i, n)/S(t_i, n)$ is defined as in (3.3.2) with W_k replaced by $W_{t_i + k} - W_{t_i}$ and $S^2(t_i, n)$ is the sample variance of $X_{t_i + 1}, X_{t_i + 2}, \cdots, X_{t_i + n}$. Thus, for a given value ("lag") of n, one obtains many samples of R/S, as many as K for small n and as few as one when n is close to the total sample size N. Next, one takes logarithmically spaced values of n, starting with $n \approx 10$. Plotting $log(R(t_i, n)/S(t_i, n))$ versus $log(n)$ results in the *rescaled adjusted range plot* (also called the *pox diagram of R/S*). When the parameter H in relation (3.3.3) is well defined, a typical rescaled adjusted range plot starts with a transient zone representing the nature of short-range dependence in the sample, but eventually settles down and fluctuates in a straight "street" of a certain slope. Graphical R/S analysis is used to determine whether such asymptotic behavior appears supported by the data. In the affirmative, an estimate $\hat{H}$ of the self-similarity parameter H is given by the street's asymptotic slope (typically obtained by a simple least squares fit) which can take any value between 1/2 and 1.

With regard to the effectiveness of R/S analysis as a function of the sample size, similar comments as in Section 4.1.1 apply. For practical purposes, the most useful and attractive feature of the R/S analysis is its relative robustness against changes of the marginal distribution. This feature allows for practically separate investigations of the self-similarity property of a given empirical record and of its distributional characteristics.

4.1.3 PERIODOGRAM-BASED ANALYSIS WITH AGGREGATION

While variance-time plots and pox plots of R/S are very useful tools for identifying self-similarity (in a mostly heuristic manner), the absence of any results for the limit laws of the corresponding statistics make them inadequate when a more refined data analysis is required (e.g., confidence intervals for the degree of self-similarity H, model selection criteria, and goodness of fit tests). In contrast, a more refined data analysis is possible for maximum likelihood type estimates (MLE) and related methods based on the periodogram. In particular, for Gaussian processes $X = (X_k : k = 0, 1, 2, ...)$, Whittle's approximate MLE has been studied extensively (see Whittle (1953), Beran (1986), Fox and Taqqu (1986), and Dahlhaus (1989)) and is defined as follows. Let $f(x; \theta) = \sigma_\varepsilon^2 f(x; (1, \eta))$ be the spectral density of X with $\theta = (\sigma_\varepsilon^2, \eta) = (\sigma_\varepsilon^2, H, \theta_3, ..., \theta_k)$, where $H = (\gamma + 1)/2$ (with γ as in (3.3.1)) describes the degree of self-similarity and $\theta_3, ..., \theta_k$ model the short-range dependence structure of the process. As the scale parameter, we use the variance σ_ε^2 of the innovation ε in the infinite AR-representation of the process, i.e., $X_j = \sum_{i \ge 1} \alpha_i X_{j-i} + \varepsilon_j$, with $\sigma_\varepsilon^2 = \text{var}(\varepsilon_j)$. Note that this implies

$$\int \log(f(x; (1, \eta)))dx = 0. \tag{4.1.1}$$

The Whittle estimator $\hat{\eta}$ of η minimizes

$$Q(\eta) = \int_{-\pi}^{\pi} \frac{I(x)}{f(x; (1, \eta))} dx \tag{4.1.2}$$

where $I(\cdot)$ denotes the *periodogram* of X defined by

$$I(x) = \frac{1}{2\pi n} \left| \sum_{j=1}^{n} X_j e^{ijx} \right|^2, \tag{4.1.3}$$

and the estimate of σ_ε^2 is given by

$$\hat{\sigma}_\varepsilon^2 = \int_{-\pi}^{\pi} \frac{I(x)}{f(x; (1, \hat{\eta}))} dx. \tag{4.1.4}$$

Then $n^{1/2}(\hat{\theta} - \theta)$ is asymptotically normally distributed if $(X_j)_{j \geq 1}$ can be written as an infinite moving average process. For Gaussian processes, $\hat{\theta}$ has the same asymptotic distribution as the MLE and is asymptotically efficient.

In this context, two problems of robustness due to (i) deviations from Gaussianity, and (ii) deviations from the assumed model spectrum are commonly encountered. Transforming the data so as to obtain approximately the desired marginal (normal) distribution is generally considered a viable method to overcome (i). For problem (ii), there are several proposals in the literature, including estimating H from the periodogram ordinates at low frequencies only or to bound the influence of $I(x)$ at high frequencies. In the presence of large data sets, an alternative and more direct method for tackling (ii) uses the method of aggregation (see Section 3.1): If $(X_i)_{i \geq 1}$ is a Gaussian process satisfying (3.3.1), then the aggregated processes $X^{(m)} (m \geq 1)$ defined by

$$X_j^{(m)} = m^{-H} L^{-1/2}(m) \sum_{i = (j-1)m + 1}^{mk} (X_i - E[X_i]), \quad j \in \{1, 2, ..., [n/m]\}, \tag{4.1.5}$$

converge (in distribution) to fractional Gaussian noise as $m \to \infty$ ($L(\cdot)$ is a slowly varying function at infinity). The same holds true if $X_i = \mu + G(Y_i)$, where $(Y_i)_{i \geq 1}$ is a Gaussian process satisfying (3.3.1), $E[G(Y_i)] = 0$, $E[G^2(Y_i)] < \infty$, and $E[G(Y_i)Y_i] \neq 0$. Hence, for sufficiently large m, fractional Gaussian noise is a good model for $X^{(m)}$ so that we can apply a maximum likelihood type estimator for fractional Gaussian noise.

Combined, Whittle's approximate MLE approach and the aggregation method give rise to the following operational procedure for obtaining confidence intervals for the self-similarity parameter H. For a given time series, consider the corresponding aggregated processes $X^{(m)}$ with $m = 100, 200, 300, \cdots$, where the largest m-value is chosen such that the sample size of the corresponding series $X^{(m)}$ is not less than about 100. For each of the aggregated series, estimate the self-similarity parameter H via a discretized version of (4.1.2) (replace the integral by a Riemann sum), where $f(x; (1, H^{(m)}))$ denotes the spectral density of fractional Gaussian noise. This procedure results in estimates $\hat{H}^{(m)}$ of H and corresponding, say, 95%-confidence intervals of the form $\hat{H}^{(m)} \pm 1.96\hat{\sigma}_H$, where $\hat{\sigma}_H^2$ is given by the central limit theorem for $\hat{\theta}$ mentioned earlier. Finally, we plot the estimates $\hat{H}^{(m)}$ of H together with their 95%-confidence intervals versus m. Such plots will typically vary a lot for small aggregation levels, but will stabilize after a while and fluctuate around a constant value, our final estimate of the self-similarity parameter H. Among the possible choices for the corresponding confidence interval, we obviously choose the one with the smallest value for m, because the size of the confidence intervals increases in m (the more we aggregate, the less observations we have).

4.2 THE SELF-SIMILAR NATURE OF ETHERNET TRAFFIC

For each of the 4 measurement periods described in Table 1, we identified what are considered "typical" low-, medium-, and high-activity hours. By limiting our statistical analysis to such hour-long traces of "typical" network traffic scenarios, we largely avoid any difficulties related to the potential non-stationary behavior of the time series under consideration due to the presence of clearly visible diurnal cycles. (See Klemes (1974) for a revealing discussion about the stationarity/non-stationarity assumption and the Hurst phenomenon; for a more detailed look at the stationarity assumption, see also Section 4.2.3 below.) With the resulting data sets, we are able to investigate features of the observed traffic (e.g., self-similarity) that persist across the network as well as across time, irrespective of the utilization level of the Ethernet. Although only one LAN could be monitored at any one time (making it impossible to study correlations in the activity on different LAN's) and all data were collected from LAN's in the same company (making it not representative for all LAN traffic), some of the characteristics uncovered by our analysis of the data in Table 1 are likely to be universal for LAN traffic and, moreover, are likely to be inherent of packet traffic in many high-speed networks of the future.

4.2.1 ETHERNET TRAFFIC OVER A 27-HOUR PERIOD

Figure 3.1 gives a "pictorial" proof of the self-similar nature of the 27 hour-long segment of the August '89 trace of Ethernet traffic (number of packets per time unit) described in the top part of Table 1. That is, for time scales ranging from milliseconds to hundreds of seconds (covering 5 orders of magnitude), the resulting traffic (measured in terms of the total number of packets per time unit) looks the same and is distinctly different from white

noise. The same holds true if instead of the number of packets per time unit, we look at the number of bytes per time unit. Here we proceed in a more rigorous manner in order to test the self-similarity property of the August '89 "snapshot" of Ethernet traffic. More precisely, using the graphical tools described in the previous section, namely, variance-time plots, pox plots of R/S, and periodogram plots, we will analyze the 3 subsets AUG89.LB, AUG89.MB, and AUG89.HB of the August '89 trace that correspond to a typical "low hour", "normal hour", and "busy hour" traffic scenario, respectively (see Table 1). Each sequence contains 360000 observations, and each observation represents the number of bytes sent over the Ethernet every 10 milliseconds. Also note that with time series of length 360000, tests for self-similarity based on the asymptotic slope of the variance-time curve or of the R/S plot become rather powerful, and the corresponding estimates for the degree of self-similarity can be expected to be very accurate.

As an illustration of the usefulness of the graphical tools for "detecting" self-similarity in an empirical record, Figure 4.2.1 depicts the variance-time curve (a), the pox plot of R/S (b), and the periodogram plot (c) corresponding to the sequence AUG89.MB, the number of bytes in every 10 millisecond interval during a "normal hour" of the August '89 trace of Ethernet traffic. The variance-time curve (Figure 4.2.1 (a)) which has been normalized by the corresponding sample variance shows an asymptotic slope that is distinctly different from -1 (dotted line) and is easily estimated to be about -.40, resulting in an estimate $\hat{H}$ of the Hurst parameter H of about $H \approx .80$. Estimating the Hurst parameter directly from the corresponding pox plot of R/S (Figure 4.2.1 (b)) leads to a practically identical estimate; the value of the asymptotic slope of the R/S plot is clearly between 1/2 and 1 (lower and upper dotted line, respectively), with a simple least-squares fit resulting in $\hat{H} \approx .79$. Finally, looking at the periodogram plot corresponding to the time series AUG89.MB, we observe that although there are some pronounced peaks in the high-frequency domain of the periodogram, the low-frequency part is characteristic for a power-law behavior of the spectral density around zero. In fact, by fitting a simple least-squares line using only the lowest 10% of all frequencies, we obtain a slope estimate $\hat{\gamma} \approx .64$ which results in a Hurst parameter estimate $\hat{H}$ of about .82. Thus, together the 3 graphical methods suggest that the sequence AUG89.MB is self-similar with self-similarity parameter $H \approx .80$.

Figure 4.2.1 illustrates yet another characteristic feature of the Ethernet measurements during a normal traffic hour of the August '89 trace. In general, variance-time plots or pox plots of R/S depict clearly visible initial transient segments, i.e., aggregation levels for which the asymptotic behavior has not yet taken over. Typically, the presence of such initial segments indicates a particular form of the short-range dependence structure of the underlying time series or, viewed from a frequency perspective, a special shape of the spectrum at the high frequencies (the latter is easily recognized by looking at the corresponding periodogram plots). In contrast, plots (a) and (b) in Figure 4.2.1 show practically no transient phase and depict an asymptotic behavior which is already present for very small levels of aggregation; similarly, the periodogram plot (c) gives no indications of any special feature of the spectrum at high frequencies (except, of course, for the presence of some clearly visible periodic components). These observations agree intuitively with the already noted similarity among the different plots in Figure 3.1 and suggest that the normal hour Ethernet traffic of the August '89 data is (exactly) self-similar in the sense of (3.1.3) rather than asymptotically self-similar in the sense of (3.1.4). In order to make the intuitive connection between Figures 3.1 (more precisely, the corresponding version with packets replaced by bytes) and 4.2.1 more rigorous, Figure 4.2.2 shows the estimates of the Hurst parameter H for selected aggregated time series derived from the sequence AUG89.MB, as a function of the aggregation level m. That is, for aggregation levels $m = 1, 5, 10, 50, 100, 500, 1000$, we plot in Figure 4.2.2 the Hurst parameter estimate $\hat{H}^{(m)}$ for the aggregated time series $X^{(m)}$ against (the logarithm of) the aggregation level m. As in Figure 4.2.1, the estimates $\hat{H}^{(m)}$ are based on the variance-time curves ("○"), the pox plots of R/S ("∗"), and the periodogram plots ("□") corresponding to the aggregated series $X^{(m)}$. Notice that the estimates are extremely stable and practically constant over the depicted range of aggregation levels $1 \leq m \leq 1000$. Thus, in terms of their second-order statistical properties, the aggregated series $X^{(m)} (m \geq 1)$ can be considered to be identical and produce, therefore, realizations which have similar overall structure and look very much alike.

In addition to the sequence AUG89.MB, we also analyzed the sequences AUG89.LB and AUG89.HB, i.e., the time series representing the number of bytes in every 10 millisecond interval during a typical "low hour" and "busy hour" of the August '89 trace of Ethernet traffic, respectively. While both series behave very much like the sequence AUG89.MB, the resulting Hurst parameter estimates $\hat{H}$ differ slightly; $\hat{H} \approx .75$ for AUG89.LB, and $\hat{H} \approx .85$ for AUG89.HB. This difference shows that although Ethernet traffic over approximately a 24-hour period is self-similar, the degree of self-similarity (expressed in terms of the Hurst parameter) depends on the utilization

level of the Ethernet and increases as the utilization increases (e.g., low H-values during night hours, high H-values for late morning and early afternoon hours). Finally, we also analyzed the sequences AUG89.LP, AUG89.MP, and AUG89.HP, i.e., the time series representing the total number of Ethernet packets rather than the total number of bytes in every 10 millisecond interval. Not surprisingly, the packet traffic is also self-similar, but with slightly larger H-values than the corresponding bytes/time unit data. For a more detailed analysis of Ethernet traffic at the packet level, see Section 4.2.2 below.

4.2.2 ETHERNET TRAFFIC OVER A 4-YEAR PERIOD

As discussed in Section 2, Ethernet LANs are generally known to change significantly during the course of a few years. The purpose of this section is to analyze additional data sets (see Table 1), similar to the August ’89 ones, but taken at different points in time and at different physical locations within the network, and to show that Ethernet traffic is self-similar, irrespective of when and where the data were collected in the Bellcore Morristown network during the 4-year period August ’89 - February ’92. The August ’89 and October ’89 data are examples of Ethernet traffic measurements taken at the same physical location within the network but during a period which saw significant changes in the computing environment (replacing old workstations by faster ones). On the other hand, the post-89 data sets (i.e., January ’90 and February ’92) are examples of Ethernet traffic measurements taken at a different physical location (than the ’89 data) within the network, during a period of significant re-configurations of the network due to the installment of additional bridges, routers, and gateways. Most importantly, as we move from the August ’89 traffic data to the February ’92 measurements, the observed network traffic moves from almost exclusively host-to-host traffic to primarily router-to-router traffic. Thus, the collection of Ethernet traffic measurements described in Table 1 enables us to examine in detail the nature of Ethernet traffic across time as well as across the network under consideration.

In contrast to Section 4.2.1, our analysis below will result in estimates of the self-similarity parameter H together with their respective 95%-confidence intervals. While for these large data sets graphical tools like variance-time plots and pox plots of R/S are important for detecting self-similarity and for crudely estimating the degree of self-similarity, it is, in general, not possible to use these techniques when deciding if there is statistically significant evidence for self-similarity and how significant the resulting estimate of the self-similarity parameter is. As discussed in Section 3.1.3, such a refined analysis is possible if maximum likelihood type estimates (MLE) or related estimates based on the periodogram are used instead of the mostly heuristic graphical estimation methods illustrated in the previous section. The latter have the advantage that they seem to perform well regardless of the shape of the marginal distribution; i.e., they give good estimates of H even for skewed and heavy-tailed distributions. On the other hand, MLE methods and related approaches based on the periodogram assume Gaussian (or, at least, approximate Gaussian) marginals and require, therefore, an initial check of the time series against possible deviations of the empirical marginal distribution from Gaussianity. In practice, one can often obtain the desired marginal distribution by transforming the original data.

In order to obtain the desired normal-looking marginals for the time series analyzed in this section, we transformed some of the original time series using the log-transformation $Y = \log(X)$. Note that an important property of the log-transformation is that asymptotically, the estimates of the self-similarity parameter H are the same for the original and the transformed data. As a result, the marginals of the data sets analyzed below appear to be approximately normal, with some more or less apparent deviations from normality in the tails. Typically, these deviations are less pronounced for the time series representing the total number of packets in every 10 millisecond interval than for those representing the total number of bytes during that same time interval. Also, the deviations from normality in the tails become less pronounced as the level of aggregation m for the aggregated time series $X^{(m)}$ increases.

Plots (a) - (d) of Figure 4.2.3 show the result of the MLE-based estimation method when combined with the method of aggregation. For each of the 4 sets of traffic measurements described in Table 1, we use the time series representing the packet counts during normal traffic conditions (i.e., AUG89.MP in (a), OCT89.MP in (b), JAN90.MP in (c), and FEB92.MP in (d)), and consider the corresponding aggregated time series $X^{(m)}$ with $m = 100, 200, 300, ..., 1900, 2000$. Note that while the original time series represent the number of Ethernet packets per 10 milliseconds, these aggregated sequences give the packet counts per 1 second, 2 seconds, ..., 19 seconds, and 20 seconds, respectively. We plot the Hurst parameter estimates $\hat{H}^{(m)}$ of H obtained from the aggregated series $X^{(m)}$, together with their 95%-confidence intervals, against the aggregation level m. Using the MLE-type method, for each time series $X^{(m)}$, the estimate $\hat{H}^{(m)}$ is obtained via a discretized version of (4.1.2), and

the corresponding 95%-confidence interval is given by $\hat{H}^{(m)} \pm 1.96\,\hat{\sigma}_{H^{(m)}}$, where $\hat{\sigma}^2_{H^{(m)}}$ is given by the central limit theorem result mentioned in Section 4.1.3. In (4.1.2) and (4.1.4), $f(\cdot\,;\theta)$ denotes the spectral density of fractional Gaussian noise. Figure 4.2.3 shows that for the packet counts during normal traffic loads (irrespective of the measurement period), the values of $\hat{H}^{(m)}$ are quite stable and fluctuate only slightly in the 0.85 to 0.95 range throughout the aggregation levels considered. The same holds for the 95%-confidence interval bands, indicating strong statistical evidence for self-similarity of these 4 time series with degrees of self-similarity ranging from about 0.85 to about 0.95. Note that the confidence interval bands widen as the aggregation level m increases; the larger m the fewer observations there are in $X^{(m)}$. The relatively stable behavior of the Hurst parameter estimates $\hat{H}^{(m)}$ for the different aggregation levels m also confirms our earlier finding that Ethernet traffic during normal traffic hours can be considered to be exactly self-similar in the sense of (3.1.3) rather than asymptotically self-similar in the sense of definition (3.1.4). Plots (a) - (d) of Figure 4.2.3 suggest that this property holds irrespective of when and where the Ethernet was monitored. For exactly self-similar time series, determining a single point estimate for H and the corresponding 95%-confidence interval is straightforward and can be done by visual inspection of plots such as the ones in Figure 4.2.3 (see below). Notice that in each of the four plots in Figure 4.2.3, we added two lines corresponding to the Hurst parameter estimates obtained from the pox diagrams of R/S and the variance-time plots, respectively. Typically, these lines fall well within the 95%-confidence interval bands which confirms our earlier argument that for these long time series considered here, graphical estimation methods based on R/S or variance-time plots can be expected to be very accurate.

In addition to the 4 normal hour packet data time series, we also applied the combined MLE/aggregation method to the other traffic data sets described in Table 1. Figure 4.2.4 (a) depicts all Hurst parameter estimates (together with the 95%-confidence interval corresponding to the choice of m discussed earlier) for each of the 12 packet data time series, while Figure 4.2.4 (b) summarizes the same information for the time series representing the number of bytes per 10 milliseconds during a typical low, normal, and busy hour for each of the four measurement periods. We also included in these summary plots the Hurst parameter estimates obtained via the variance-time plots ("○") and R/S analysis ("∗") in order to indicate the accuracy of these essentially heuristic estimators when compared to the statistically more rigorous Whittle estimator ("•"). Concentrating first on the packet data, i.e., Figure 4.2.4 (a), we see that despite the transition from mostly host-to-host workgroup traffic during the August ’89 and October ’89 measurement periods, to a mixture of host-to-host and router-to-router traffic during the January ’90 measurement period, to the predominantly router-to-router traffic of the February ’92 data set, the Hurst parameter corresponding to the typical normal and busy hours, respectively, are comparable, with slightly higher H-values for the busy hours than for the normal traffic hours. This latter observation might be surprising in light of conventional traffic modeling where it is commonly assumed that as the number of sources (Ethernet users) increases, the resulting aggregate traffic becomes smoother and smoother. In contrast to this generally accepted argument for the "Poisson-like" nature of aggregate traffic, our analysis of the Ethernet data shows that, in fact, the aggregate traffic tends to become less smooth (or, more bursty) as the number of active sources increases (see also our discussion in Section 5.1). In fact, while there were about 120 hosts that spoke up during the August ’89 or October ’89 busy hour, we heard from an order of magnitude more hosts (about 1200) during the January ’90 high traffic hour; the comparable number of active hosts during the February ’92 busy hour was around 600. The major difference between the early (pre-1990) measurements and the the later ones (post-1990, i.e., January ’90 and February ’92) can be seen during the low traffic hours. Intuitively, low period router-to-router traffic consists mostly of machine-generated packets which tend to form a much smoother arrival process than low period host-to-host traffic which is typically produced by a smaller than average number of actual Ethernet users, i.e., researchers working late hours. Next, turning our attention to Figure 4.2.4 (b), i.e., the Hurst parameter estimates for the bit rates, we observe that as in the case of the packet data, the degree of self-similarity H increases as we move from low to normal to high traffic hours. Moreover, while there is practically no difference between the two post-1990 data sets, the two pre-1990 sets clearly differ from one another but follow a similar pattern as the post-1990 ones. The difference between the August ’89 and October ’89 measurements can be explained by the transition from diskless to "dataless" workstations that occurred during the latter part of 1989 (see Section 2.2.2). Except during the low hours, the increased computing power of many of the Ethernet hosts causes the Hurst parameter to increase and gives rise to a bit rate that matches more or less the self-similar feature of the corresponding packet process. Also note that the 95%-confidence intervals corresponding to the Hurst parameter estimates for the low traffic hours are typically wider than those corresponding to the estimates of H for the normal and high traffic hours. This widening indicates that Ethernet traffic during low traffic periods is asymptotically self-similar rather than exactly self-similar.

We also notice in Figure 4.2.4 that some of the analyzed time series result in estimated Hurst parameters close to 1, i.e., their corresponding 95%-confidence intervals include the value $H = 1$. When finding an H-estimate close to 1, it is generally advised to analyze the time series further in order to ensure that the observed high degree of self-similarity is genuine and cannot be explained by elementary arguments such as differencing (see Mandelbrot and Taqqu (1979)) or shifting means (see Klemes (1974)). In order to illustrate our more detailed analysis concerning the stationarity property of the time series in Figure 4.2.4, we will concentrate on the sequences JAN90.HP and FEB92.HP since both of these data sets show a high H-value and will be used again in the section below.

Visual inspection of the time series JAN90.HP and FEB92.HP (see Figure 4.2.5 which depicts the aggregated process $X^{(100)}$ for each sequence, i.e., the number of Ethernet packets per second during the high traffic hours of the January '90 and February '92 measurement periods, respectively) and comparisons with traces of fractional Gaussian noise with $H \approx 0.9$ (see, for example, plots in Mandelbrot and Wallis (1969a), Mandelbrot and Taqqu (1979), and Hampel (1987)) show no obvious signs of non-stationarity; the mean seems to be changing with time but the overall mean appears constant and although, locally, there clearly exist spurious trends and cycles of varying frequences, these "typical" features of nonstationarity are characteristic of stationary long-range dependent processes. Also, the variance-time plots as well as the pox diagrams of the adjusted range R (without rescaling by S) of the two time series yield slope estimates (not shown) that are consistent with the observed high H-values. As discussed in Beran et al. (1992), this consistency is a strong indication that the given time series cannot be regarded as nonstationary due to a lack of differencing. Finally, we briefly illustrate a recently developed method (see Beran and Terrin (1992)) to test for nonstationarity of a given time series due to changes in H over time.

The problem of deciding whether inhomogeneities in H over time are real (due to actual changes in the dependence structure in the data) or are due to randomness is very delicate because even optimal estimates of H turn out to vary considerably when calculated for disjoint parts of a long-range dependent time series (see Beran and Terrin (1992)). In order to assess quantitatively how much the estimates of H can vary when estimated from different portions of the data, Beran and Terrin (1992) obtain the joint asymptotic distribution of the Whittle estimates (see (4.1.2)) of H based on disjoint parts of the data. More precisely, in order to test the hypothesis $H_0 : H \equiv H_0$ for X_i, $1 \le i \le N$, against the alternative $H_a : H \neq H_0$, i.e., among the sequences $(X_{N_{j-1}+1}, \ldots, X_{N_j})_{j=1,\ldots,k}$, $N_0 = 0 < N_1 < \cdots < N_k = N$ with corresponding H-parameters $H_1, \ldots, H_k$, there exists at least one pair $j \neq i$ such that $H_j \neq H_i$, define

$$\bar{H} = N^{-1} \sum_{j=1}^{j=k} (N_j - N_{j-1})\, \hat{H}_j \tag{4.2.1}$$

and the test-statistic

$$T_{1,2,\ldots,k} = \sum_{j=1}^{j=k} (\hat{H}_j - \bar{H})^2 [\hat{\sigma}_j^2 (N_j - N_{j-1})]^{-1} \tag{4.2.2}$$

where $\hat{\sigma}_j^2$ is given by a generalization of the central limit theorem for Whittle's estimator (see Beran and Terrin (1992)). They then show that under the null hypothesis H_0, $T_{1,2,\ldots,k}$ is asymptotically χ^2-distributed with $k-1$ degrees of freedom. Hence we reject H_0 at the level of significance α, if $T_{1,2,\ldots,k} > \chi^2_{k-1;\alpha}$, where $\chi^2_{k-1;\alpha}$ is the upper $(1-\alpha)$-quantile of the χ^2-distribution with $k-1$ degrees of freedom.

We applied this procedure for testing for a constant H-parameter to the two series JAN90.HP and FEB92.HP. More precisely, in order to reduce the amount of computation, we considered the corresponding aggregated time series $X^{(100)}$ of length 3600 depicted in Figure 4.2.5 representing the number of Ethernet packets per second during a high traffic hour in the January '90 and February '92 data, respectively. For example, for the sequence $X^{(100)}$ derived from the time series JAN90.HP, Figure 4.2.5 (a) also shows the different ways we partitioned the data into disjoint pieces and, regardless of the chosen partition, the null hypothesis is never rejected, with P-values consistently larger than about 0.3. For instance, when partitioning $X^{(100)}$ into 6 (5) non-overlapping 10-minute blocks starting at the beginning of the hour (5 minutes into the high hour), the estimates of H for these 6 (5) time blocks are .939, .975, .941, .984, .962, and .919 (.976, .917, .977, .986, and .946), and the values of the test-statistic is $T_{1,2,3,4,5,6} = 2.987$ ($T_{1,2,3,4,5} = 3.203$) with a P-value of 0.70 (.52). Similarly convincing evidence for a constant, yet high, H-value is obtained using the 6-minute, 15-minute, or 20-minute partitions of $X^{(100)}$ as indicated in Figure 4.2.5 (a). On the other hand, the same analysis performed for the aggregated series $X^{(100)}$ derived from the sequence FEB92.HP shows some statistical evidence that H is not constant over the whole high traffic hour.

Indeed, partitioning $X^{(100)}$ the same way as before into 6 (5) 10-minute intervals yields the following estimates of H for the 6 (5) segments: .980, .932, .948, .900, .786, and .920 (.973, .882, .946, .907, and .867). The resulting test-statistics are $T_{1,2,3,4,5,6} = 22.33$ with a P-value of 0.0005, and $T_{1,2,3,4,5} = 7.59$ with a P-value of 0.107. While the six terms in the sum $T_{1,2,3,4,5,6}$ (see (4.2.2)) are 4.696, 0.417, 1.334, 0.123, 15.681, and 0.080, the five terms in $T_{1,2,3,4,5}$ are 3.256, 1.053, 0.926, 0.060, and 2.292. This shows that rejection of the null hypothesis is mostly caused by H_5 and is in agreement with the visual impression of Figure 4.2.5 (b): the estimated long-range dependence for the 5th 10-minute segment seems to be clearly weaker than that of the first 40-minute interval. Thus, while both time series JAN90.HP and FEB92.HP yield an estimated H-value close to 1, the sequence JAN90.HP appears stationary and seems to be best described using a constant H-value of about .98. On the other hand, the sequence FEB92.HP shows some signs of nonstationarity due to changes in the degree of self-similarity after about 40 minutes into the high traffic hour, and modeling the whole hour using two different H-values ($H \approx .96$ for the first 40 minutes, and $H \approx .86$ for the last 20 minutes) seems more appropriate.

We performed similarly detailed studies of the stationarity aspect for the other time series described in Table 1. With the exception of the sets FEB92.MP and FEB92.MB, all sequences appear stationary and seem to be appropriately modeled using a constant H-value. The normal traffic hour of the February '92 data has the peculiar feature that for a stretch of about 20 minutes, 1-2 conversations running almost simultaneously produce about 50% of all traffic. In fact, after removing this clearly visible traffic component, the resulting traffic (in terms of packets and bytes) shows no apparent signs of nonstationarity and seems to be best described with a constant Hurst parameter of about .85.

4.3 SELF-SIMILAR FEATURES OF SOME IMPORTANT ETHERNET TRAFFIC COMPONENTS

The Ethernet traffic analyzed in Section 4.2 is also called *internal* traffic and consists of all packets on a LAN. Understanding internal traffic is not only of practical use for maintenance and administration of LAN's, but characteristics such as the observed self-similarity in the internal LAN Ethernet traffic are also likely to be inherent of packet traffic in many high-speed networks of the future since some of the projected major contributors to this traffic exhibit these characteristics. For instance, we analyze in this section *remote* or *external* Ethernet traffic, an important component of the internal traffic consisting of all those Ethernet packets that originate on one LAN but are routed to another LAN. That is, for the traffic measurements at hand, an *external* packet is defined to be an IP (Internet protocol) packet with a source or destination address that is not on any of the Bellcore networks. Thus, this external traffic can be viewed as representative for LAN interconnection services which are expected to contribute significantly to future broadband traffic. Together with our current knowledge about characteristics of variable-bit-rate (VBR) video traffic (see Beran et al. (1992), and Garrett and Willinger (1993)), another potentially major contributor to future broadband traffic, understanding external Ethernet traffic will provide a realistic view of the nature of broadband traffic.

4.3.1 EXTERNAL ETHERNET TRAFFIC

Table 2 summarizes the external Ethernet traffic data analyzed below. We consider the two most recent measurement traces i.e., the January '90 and February '92 data sets, and for ease of comparison, we analyze for both measurement periods the time series consisting of the number of external packets (bytes) per 10 milliseconds during the same low-, normal-, and high-hours of (internal) Ethernet traffic as considered in Table 1. Thus, for example, the data sets JAN90E.LP, JAN90E.MP, and JAN90E.HP in Table 2 represent the portions of external packets from the corresponding data sets in Table 1, namely JAN90.LP, JAN90.MP, and JAN90.HP, and are, therefore, not necessarily identical with a typical low-, normal-, or high-hour of external Ethernet traffic. The last column in Table 2 shows that external traffic (in terms of packets or bytes) makes up between 1-10% of the internal traffic during the low hours in January '90 and February '92, about 2-25% during the corresponding busy hours, and up to 56% during the February '92 normal hour. As a result, it is reasonable to expect external traffic to behave very similarly to the overall traffic analyzed in Section 4.2. Differences (if any) between the internal and external traffic can, in general, be attributed to NSF traffic between workstations and file servers which is missing completely in the external traffic.

Repeating the same laborious analysis of Section 4.2 to the data sets described in Table 2, we find that in terms of its self-similar nature, external traffic does not differ from the internal traffic studied earlier. More specifically, the Hurst parameters for the external traffic during normal and high (internal) traffic hours (or during

Traces of External Ethernet Traffic Measurements					
Measurement Period	Internal Traffic Data (see Table 1)	Data Set	Total Number of Bytes	Total Number of Packets	Percentage of Internal Traffic
JANUARY 1990					
Start of trace: Jan. 10, 6:07am	JAN90.LB	JAN90E.LB	1,105,876		1.27%
	JAN90.LP	JAN90E.LP		9,369	3.02%
	JAN90.MB	JAN90E.MB	16,536,148		9.05%
End of trace: Jan. 11, 10:17pm	JAN90.MP	JAN90E.MP		87,307	13.57%
	JAN90.HB	JAN90E.HB	13,023,016		2.00%
	JAN90.HP	JAN90E.HP		68,405	4.96%
FEBRUARY 1992					
Start of trace: Feb. 18, 5:22am	FEB92.LB	FEB92E.LB	2,319,881		4.08%
	FEB92.LP	FEB92E.LP		25,247	10.89%
	FEB92.MB	FEB92E.MB	86,283,283		55.80%
End of trace: Feb. 20, 5:16am	FEB92.MP	FEB92E.MP		270,636	51.60%
	FEB92.HB	FEB92E.HB	55,154,789		24.50%
	FEB92.HP	FEB92E.HP		202,367	21.35%

Table 2. Qualitative description of the sets of external Ethernet traffic measurements used in the analysis in Section 4.3.1.

previously identified stationary parts of the corresponding data sets) are only slightly smaller than the ones depicted in Figure 4.2.5. For instance, even though the portion of external packets during the high (internal) traffic hour of the January '90 data is only 2% of all the packets seen during this period, the data set JAN90E.HP seems to be well described by an H-value that changes from $H = 0.82$ for the first 30 minutes to $H = 0.94$ for the second 30 minutes; recall that the corresponding data set of internal traffic, i.e., the sequence JAN90.HP, has an estimated Hurst parameter of 0.98. A more significant change in the Hurst parameter occurs during the low traffic hours. While the internal traffic data (JAN90.LB, JAN90.LP, FEB92.LB, and FEB92.LP) yield a Hurst parameter of about 0.70, the sequences JAN90E.LB, JAN90E.LP, FEB92E.LB, and FEB92E.LP have $H \approx 0.55$, and the corresponding 95%-confidence intervals contain the value $H = 0.5$. These are the only cases in all the data sets considered in this paper, where an H-value of 0.5 (i.e., conventionally used short-range dependent models such as Poisson, batch-Poisson, or MMPPs) seems to describe the data accurately. For all other data sets described in Tables 1 and 2, the 95%-confidence intervals for the Hurst parameter estimates do not even come close to covering the value $H = 0.5$. As already mentioned in our discussion of Figure 4.2.5, the low hour traffic in the January '90 and February '92 data is mostly machine-generated and produces traffic that is typically smoother (i. e., less bursty) than traffic that is generated during the normal and busy hours by humans using their workstations. This argument applies even more so when considering low hour external traffic.

4.3.2 EXTERNAL TCP TRAFFIC

We also looked at the portion of external traffic using the Transmission Control Protocol (TCP) and IP. There were two main reasons for this. First, the traditional services offered by the Internet are for the most part based around TCP, which offers reliable delivery of data and protection against data loss due to lost or corrupted packets. These services include remote login, file transfer (including anonymous file transfer for making information and programs publicly available to any Internet user), electronic mail, and more recently the delivery of the electronic bulletin board known as Netnews. The second reason is that application programs using the TCP protocol have significantly *less* control over how their data is actually sent than do applications using the User Datagram Protocol (UDP) or their own protocol. The TCP protocol has significant control over how the user data is segmented and a great deal of control over the spacing of the packets as they are sent out. UDP is a "send and pray" service, which means the application must deal with lost or misordered data whereas TCP takes care of this for the application.

When investigating the external TCP traffic, we found that there was little point in doing a separate analysis. For instance, in the heavy traffic hour from the MRE backbone taken in 1992 (FEB92E.HP), 87% of the packets were TCP packets, and a plot of the external TCP traffic (e.g., number of external TCP packets per second) is practically indistinguishable from the plot in Figure 4.3.5 (b) that shows the number of external packets per second for the February '92 high traffic hour. Of those TCP packets of the FEB92E.HP data set, about 66% of the packets were for file transfer, 9% for remote login/TELNET, 11% for electronic mail, and 13% for netnews delivery. The 12% of non-TCP traffic simply had no effect on the results of our analysis for this data set; external TCP traffic is practically identical to the external traffic studied in Section 4.3.1, and our findings for the external traffic apply directly to external TCP traffic.

Note that the above external traffic was monitored after the firewall router was put in place between Bellcore and the external Internet world (see Section 2). While the firewall blocks some services from the outside, and restricts where others may go within Bellcore, it places few restrictions on those within the Bellcore security domain other than blocking the remote execution of applications based on the X window system. But looking at the JAN90E.HP data set we see an even lower percentage of non-TCP traffic than in the set discussed above, leading us to conclude that the effect of non-TCP external traffic is insignificant and that the nature of external TCP traffic is the same as that of the external traffic analyzed in Section 4.3.1.

5. SIGNIFICANCE OF SELF-SIMILARITY FOR TRAFFIC ENGINEERING

Our measured data show drastically different statistical properties than those predicted by the stochastic models currently considered in the literature. Almost all these models are characterized by an exponentially decaying autocorrelation function. As a result, they give rise to a Hurst parameter estimate of $H = .50$. In view of Figure 4.2.1, their corresponding variance-time curves would be (ultimately) parallel to the dotted line (with slope -1), their corresponding pox plots of R/S would be (ultimately) parallel to the lower dotted line (with slope 1/2), and their corresponding periodogram plots would be flat for the low frequencies (reflecting the fact that their spectral densities are bounded at low frequencies). In terms of the aggregation procedure described above, the theoretical models have the property that typically, after aggregating over non-overlapping blocks of about 10-100 observations, the aggregated series become indistinguishable from second-order pure noise (i.e., a sequence of independent and identically distributed random variables). Especially surprising in light of the traffic models currently considered in the literature is the observation that the aggregate traffic tends to become more bursty as the number of active Ethernet sources increases. In fact, our analysis of the Ethernet data shows that the generally accepted argument for the "Poisson-like" nature of aggregate traffic (namely, that aggregate traffic becomes smoother as the number of traffic sources increases) has very little to do with reality - it holds *only* during the not very interesting low traffic periods which see mostly machine-generated router-to-router traffic. The fact that one can distinguish clearly - with respect to second-order statistical properties - between the existing models for Ethernet traffic and our measured data is surprising and clearly questions some of the modeling assumptions that have been made in the past.

While the distinction between currently proposed packet traffic models and actual packet traffic measurements from such applications as Ethernet LANs, VBR video codecs (see Beran et al. (1992)), and an ISDN office automation (see Meier-Hellstern et al. (1991)) is obvious from a statistical perspective, potential traffic engineering implications of this distinction are currently under intense scrutiny. In this section, we emphasize three direct implications of the self-similar nature of packet traffic for traffic engineering purposes: modeling individual sources, e.g., Ethernet users, inadequacy of conventional notions of "burstiness", and effects on congestion control/management for packet networks. We conclude with some guidelines for a first step toward modeling self-similar traffic and suggest possible refinements and some open problems.

5.1 ON THE NATURE OF TRAFFIC GENERATED BY AN INDIVIDUAL ETHERNET USER

In Section 4, we showed that irrespective of when and where the Ethernet measurements were collected, the traffic is self-similar, with different degrees of self-similarity depending on the load on the network. We did so without first studying and modeling the behavior of individual Ethernet users (sources). Although historically, accurate source modeling has been considered an absolute necessity to a successful modeling of *aggregate* traffic, we show here that in the case of self-similar packet traffic, knowledge of fundamental characteristics of the aggregate traffic can provide new insight into the nature of traffic generated by an individual user. Thus, this section

gives a physical explanation for the visually obvious (see Figure 3.1) and statistically significant (see Figure 4.2.4) self-similarity property of Ethernet LAN traffic in terms of the behavior of individual Ethernet users.

To this end, we recall Mandelbrot's construction of fractional Brownian motion (see Section 3.2.3) that uses renewal reward processes as simple building blocks, and interpret it in the context of possible Ethernet user behaviors. That is, the renewal reward process $W^{(m)} = (W^{(m)}(t) : t = 0, 1, 2, ...)$ given by equation (3.2.5) represents the amount of information (in bits, bytes, or packets) generated by user m at time t $(1 \le m \le M,\ t \ge 0)$. In fact, if bits or bytes are the preferred units, the renewal reward process source model resembles the popular class of fluid models (see Anick et al. (1982)). On the other hand, if we think of packets as the underlying unit of information, the renewal reward process is basically a packet train model in the sense of Jain and Routhier (1986). For ease of presentation, we can assume that the "rewards" $W_0, W_1, W_2, \cdots$ take only the values 1 and 0 (or, to keep $E[W] = 0$, +1 and -1 with equal probabilities), where the value 1/0 during a renewal interval indicates an "active"/"inactive" period during which the source sends 1/0 unit(s) of information every time unit. Although Ethernets are not slotted communications systems, there is a natural time unit (slot size) imposed by the speed of an Ethernet for transmitting a given unit of information (e.g., about 60μs for the smallest allowed Ethernet packet of size 64 bytes), and the definition (3.2.5) of a renewal reward process makes implicit use of this apparent slotted feature of Ethernets. At the same time, it is more intuitive to assume (and we will do so) that the inter-renewal times $U_0, U_1, U_2, \cdots$ are real-valued rather than discrete as is the case in Section 3.2.3; although the results in Section 3.2.3 are stated for discrete U's, this assumption is not crucial for our purposes.

The crucial property that distinguishes the renewal reward process source model from the above mentioned fluid model or packet train model is that the inter-renewal intervals (i.e., the lengths of the "active"/"inactive" periods) are *heavy-tailed* in the sense of equation (3.2.3) or, using Mandelbrot's terminology, exhibit the *infinite variance syndrome*. Intuitively, (3.2.3) states that with relatively high probability, the "active"/"inactive" periods are very long, i.e., each W_m can assume the same value for a long period of time. While this heavy-tailed property of the active/inactive periods seems rather plausible in light of the way a typical workstation user contributes to the overall traffic on the Ethernet, we have not yet analyzed the traffic generated by individual Ethernet users in order to validate the simple renewal reward source model assumption. However, evidence in support of the infinite variance syndrome in packet traffic measurements already exists. For example, in a recent study of traffic measurements from an ISDN office automation application, Meier-Hellstern et al. (1991) observed that the extreme variability in the data (e.g., interarrival times of packets, number of successive packet arrivals in certain states) cannot be adequately captured using traditional packet traffic models but, instead, seems to be best described with the help of heavy-tailed distributions of the form (3.2.3). They subsequently propose an elaborate and highly parameterized model for the measured terminal-generated packet traffic that practically rules out a simple approximate model for the aggregate or superposition traffic.

In contrast, the renewal reward source model for the traffic generated by an individual workstation user is extremely simple; moreover, we have seen in Section 3.2.3 that when aggregating the traffic of many such source models, the resulting superposition process is a fractional Brownian motion, i.e., a self-similar Gaussian process with self-similarity parameter $H = (3 - \alpha)/2$, where α is given in (3.2.3). In fact, we conjecture that more sophisticated source models (typically with some dependence among and between the "active" and "inactive" periods) will result in a similar superposition process (possibly with additional structure for the short-range dependence), as long as a typical distribution for the length of the "active"/"inactive" period follows (3.2.3). Clearly, since the superposition process of many (independent) renewal reward source models is fractional Brownian motion with $H = (3 - \alpha)/2$, the time series representing, for example, the total number of bytes or Ethernet packets every 10 millisecond, behaves like fractional Gaussian noise, i.e., an exactly self-similar discrete-time Gaussian process with the same self-similarity parameter $H = (3 - \alpha)/2$. In this sense, our analysis in Section 4 suggests that a simple renewal reward process is an adequate traffic source model for an individual Ethernet user and that often, a more detailed source modeling might not be needed since the convergence result in Section 3.2.3 shows that many of the details disappear during the process of aggregating the traffic of many sources and only property (3.2.3) is required for the fractional Brownian motion behavior of the superposition process to hold. Note that we have reached this conclusion by treating the Ethernet packets essentially as black boxes; with the exception of determining the number of hosts that spoke up during a given measurement period (see Section 2), we did not look into the packet header fields or distinguished packets based on their source or destination. Further work on extracting the relevant source-destination addresses from our measurements and on statistically validating the infinite variance property of the inter-renewal periods of a single source is currently in progress.

5.2 SELF-SIMILARITY AND SOME COMMONLY USED NOTIONS OF BURSTINESS

On an intuitive level, the results of our statistical analysis of the Ethernet traffic measurements in Section 4 can be summarized by saying that typically, the higher the load on the Ethernet the higher the estimated Hurst parameter H, i.e., the degree of self-similarity in the arrival rate process (in terms of packets or bytes). Visual comparisons between the different traces also suggest that the larger H, the "burstier" the corresponding trace appears. Trying to capture the intuitive notion of "burstiness" with the help of the Hurst parameter H becomes particularly appealing in light of the relation $H = (3 - \alpha)/2$ mentioned in the previous section between the self-similarity parameter H and the parameter α that characterizes the "thickness" (see (3.2.3)) of the tail of the inter-renewal time distribution (i.e., of the lengths of the "active"/"inactive" periods). Clearly, the heavier the tail in (3.2.3) (i.e., the closer α gets to 1), the greater the variability of the "active"/"inactive" periods and hence, the burstier the traffic generated by an individual source. Going from α to H relates burstiness of an individual source to burstiness of the aggregate traffic: the higher the H, the burstier the aggregate traffic. The fact that the Hurst parameter H captures the intuitive notion of burstiness in a mathematically rigorous manner through the concept of self-similarity and, at the same time, also seems to agree well with the visual assessment of bursty behavior challenges the feasibility of some of the most commonly used measures of "burstiness". The latter include the *index of dispersion (for counts)*, the *peak-to-mean ratio*, and the *coefficient of variation (of inter-renewal times)*.

A commonly used measure for capturing the variability of traffic over different time scales is provided by the *index of dispersion (for counts)* and has recently attracted considerable attention (e.g., Heffes and Lucantoni (1986), Sriram and Whitt (1986), Fendick et al. (1991)). For a given time interval of length L, the index of dispersion for counts (IDC) is given by the variance of the number of arrivals during the interval of length L divided by the expected value of that same quantity. Figure 5.2.1 depicts the IDC as a function of L in log-log coordinates; it shows the IDC for both internal (solid lines) and external (dashed lines) traffic from the high hour of the January '90 (a) and February '92 data (b). Note in particular that the IDC increases monotonically throughout a time span that covers 4-5 orders of magnitude. This behavior is in stark contrast to conventional traffic models such as Poisson or Poisson-like processes and the popular Markov-modulated Poisson processes (MMPPs), where the IDC is either constant or converges to a fixed value quite rapidly. On the other hand, self-similar traffic models are easily shown to produce a monotonically increasing IDC. In fact, assume for simplicity that the process X representing the total number of packets seen in every 10 millisecond interval, is fractional Gaussian noise (with positive drift) with self-similarity parameter H. Then we have

$$IDC(L) = var(\sum_{j=1}^{j=L} X_j) / E[\sum_{j=1}^{j=L} X_j] \approx c\, L^{2H-1}, \tag{5.2.1}$$

where c is a finite positive constant independent of L. Note that when plotting $log(IDC(L))$ against $log(L)$, property (5.2.1) results in an asymptotic straight line with slope $2H - 1$. The dotted lines in Figure 5.2.1 represent the IDC curves predicted by self-similar traffic models with $H \approx 0.94$ (JAN90.HP) and $H \approx 0.96$ (FEB92.HP), respectively. Similarly striking agreement between the empirical and theoretical IDC curves can be observed for the corresponding external traffic data sets; in agreement with our earlier observations (see Section 4.3.2), the slopes of the IDC curves for the external traffic data are smaller than those for the corresponding internal traffic data. The close agreement between the empirical and theoretical curves suggests that plotting IDC curves provides yet another fast graphical method (in addition to the earlier discussed variance-time plots and pox diagrams of R/S) for estimating the Hurst parameter of a given time series. Compared to the closely related variance-time plot method, plotting the IDC curve and estimating its slope provides a quick and simple engineering-based approach to testing for self-similarity of a set of traffic measurements.

Leland and Wilson (1991) have pointed out the problem with using the *peak-to-mean ratio* as a measure for "burstiness" in the presence of self-similar traffic. The observed ratio of peak bandwidth (i.e., peak arrival rate of, say, bytes) to mean bandwidth depends critically on the time interval over which the peak and mean bandwidth is determined, i.e., essentially any peak-to-mean ratio is possible, depending on the length of the measurement interval. For a two-week long trace of the October '89 measurements, they show that the peak rate in bytes for the external traffic observed in any 5 second interval is about 150 times the mean arrival rate, while the peak rate observed in any 5 millisecond interval is about 710 times the mean. The dependence of this burstiness measure on the choice of the time interval is clearly undesirable.

Finally, we remark that the use of the *coefficient of variation* (for interarrival times), i.e., the ratio of the standard deviation of the interarrival time to the expected number of the interarrival time, as a measure of

"burstiness" becomes questionable because of the potential "heavy-tailedness" (in the sense of (3.2.3)) of the interarrival times and the implied infinite variance property. Although the empirical standard deviation can always be calculated, it will depend crucially on the sample size and can attain practically any value as the sample size increases.

5.3 CONGESTION MANAGEMENT IN THE PRESENCE OF SELF-SIMILAR TRAFFIC

In order to illustrate the effect of self-similar traffic on basic architectural issues concerning high-speed, high-bandwidth communications systems of the future, we concentrate below on the complex problem of congestion management for B-ISDN. For our purposes, congestion is defined to be a user-visible degradation of network performance and hence, understanding congestion requires knowledge of the nature of traffic transported by the network. However, since (i) B-ISDN is expected to offer a broad variety of new services exhibiting a wide range of different traffic characteristics, (ii) it is uncertain at this time which services will be offered and which will contribute significantly to the overall traffic, and (iii) traffic characteristics of many of the new services are, in general, poorly understood, the nature of the congestion to be managed for B-ISDN remains almost completely unknown.

Here we discuss the congestion behavior that may be induced by a network that provides interconnection for LANs or VBR video service through connectionless data service such as Switched Multimegabit Data Service (SMDS) or Frame Relay, operating at 1.5 to 45 Mb/s. Our discussion below follows closely Fowler and Leland (1991), and Leland and Wilson (1991), who use long segments of actual Ethernet traffic measurements for their trace-driven simulations. Because of the statistical groundwork established in Section 4, their conclusions about the nature of congestion and the task of congestion management for B-ISDN provide convincing evidence for the significance of self-similar or fractal network traffic for engineering future integrated high-speed data networks. In light of the same fractal-like behavior of VBR video traffic (see Beran et al. (1992)), their conclusions are likely to apply also to a more heterogeneous B-ISDN traffic environment. While Leland and Wilson (1991) employ the Ethernet data in a trace-driven simulation of a LAN/B-ISDN interface proposed for SMDS and thus, shed light on the problem of designing and sizing of network access controllers, Fowler and Leland (1991) use the data for modeling the offered traffic of a connectionless service that provides LAN interconnections and show subsequently, what congestion might look like in a B-ISDN environment. Below, we briefly summarize their findings and interpret them in light of the statistical properties of the Ethernet data established in Section 4.

5.3.1 TRAFFIC SHAPING AT THE NETWORK ACCESS POINT

The underlying idea of an LAN/B-ISDN access proposal is to reduce the burstiness of the external LAN traffic submitted to a wide-area network by delaying the LAN packets arriving at the network access point for time intervals of the order of a few milliseconds. Leland and Wilson (1991) consider the access class scheme proposed for SMDS on public B-ISDN. SMDS is a connectionless data service with interfaces that split the variable length Ethernet packets into fixed-sized cells at the access point and reassemble the packets at the egress point. Packets arriving from a LAN are buffered at the interface and delivered to the cell-based network at some maximum rate, subject to traffic shaping that limits the burstiness and the sustained input rate. Packets that do not fit entirely within the available buffer space are dropped so that partial packets are not transmitted. Clearly, the amount of buffering at the interface strongly affects the packet delay distribution and packet loss probability observed by a LAN. At the lowest rate, DS-1 based SMDS service access path, the proposed quality of service requirement asks for 95% of the packets to be delivered with less than 140 milliseconds delay and for less than 0.01% (10^{-4}) of the packets to be lost.

Simple conventional models based on the external LAN traffic studied in Section 4.3 suggest that these quality of service requirements can readily be met. However, both packet loss and delay behaviors differ drastically between trace-driven simulations based on the actual traffic measurements and those based on these formal traffic models and indicate that the proposed quality of service requirements might be hard to meet for realistic traffic scenarios. In particular, Leland and Wilson (1991) observe that overall packet loss decreases very slowly with increasing buffer capacity, in sharp contrast to Poisson-based models where losses decrease exponentially fast with increasing buffer size. Moreover, packet delay (95th percentile) always increases with buffer capacity, again in contrast to the formal models where delay does not exceed a fixed limit regardless of buffer size. This distinctive loss and delay behavior seen with measured Ethernet traffic is typical for self-similar or fractal traffic and can be readily explained using either one of the fractal properties referred to earlier. For example, the frequency domain

manifestation of self-similarity (see (3.3.1))shows that both low and high frequencies are significant; heavy losses and long delays occur during long time-frame bursts (due to the presence of low frequencies) and can, therefore, not be dealt with effectively by larger buffers. Obviously, the presence of the low frequencies in the Ethernet data established in Section 4 distinguishes the measured traffic from the currently proposed formal models and results in the observed discrepancy between the trace-driven simulations and those using these standard models.

5.3.2 BROADBAND NETWORK CONGESTION MANAGEMENT

Once traffic has entered the network, congestion management decisions depend largely on the future behavior of the traffic stream, over timescales of milliseconds for immediate buffer allocations at individual switches within the network, through hundreds of milliseconds for congestion feedback policies, out through seconds and minutes for call admission policies and routing decisions. Thus, in contrast to voice networks that provide plain old telephone service (POTS), broadband congestion management requires models which accurately describe network traffic over such a wide range of time scales. Self-similar models were found in Section 4 to describe Ethernet LAN traffic well over time scales ranging from milliseconds to hours and can, therefore, be expected to provide a very accurate picture of the nature of congestion for broadband networks and, subsequently, for the feasibility/infeasibility of certain congestion management functions.

To this end, Fowler and Leland (1991) simulate a simple network that provides LAN interconnection. The switches within the network are assumed to be perfect output queueing switches and to provide connectionless transport for packets of varying length which, in turn, are composed of fixed-length, 53-octet cells. Each LAN is assumed to be a 10 Mb/s Ethernet and to have a separate access connection to its home switch; the switches are connected by one or more interoffice links. Access and interoffice links are operating at 45 Mb/s. There is a queue associated with each interoffice link that can receive packets from any LAN that homes on the switch. When a packet arrives at the link buffer, it either fits entirely into the buffer, or is immediately discarded without affecting the packets already in the buffer. The combined traffic offered to an interoffice link by multiple LANs was modeled using the measured Ethernet traces. Congestion due to traffic access contention for the interoffice links is the major focus of their simulation study.

There are three major components to congestion control in high-speed networks: *prevention*, *avoidance*, and *recovery*. Congestion prevention involves designing and building a network that minimizes the probability that congestion will occur, and it includes judiciously engineered components (based on the expected variation of traffic on time scales of hours, days, or months), well-designed routing algorithms, and policing features to ensure that a user's access line does not exceed its subscribed traffic rate. Congestion avoidance is basically a prediction problem and involves detecting when congestion is imminent, and taking actions designed to prevent it. Congestion recovery is action taken by the network after performance degradation is detected to limit the effects of congestion. In order to know what actions to take for congestion recovery, it is essential to know "How long does congestion typically persist?", and "What patterns of loss (or delay) occur during congestion?"

The main lessons learned from Fowler and Leland's (1991) simulation study are: (i) There exists large variation in network traffic on time scales of hours, days, or months; while this aggravates careful sizing of network components, small errors in engineering can incur drastic penalties in loss or delay. (ii) Although some of the standard traffic models suggest that congestion problems essentially disappear with sufficient buffer capacity, realistic network traffic shows that such behavior cannot be expected; large buffers will not prevent congestion from occurring but introduce instead undesirable delay characteristics. (iii) During congestion periods, congestion persists long enough for the effects of user and protocol responses to be felt. (iv) A detailed examination of congestion periods shows that when congestion occurs, losses are severely concentrated and are far greater than the background loss rate; losses may exceed the long-run loss probability by an order of magnitude during the first second following the onset of congestion, while the losses are elevated by over two orders of magnitude during the first 100 milliseconds. (v) Fortunately, many congestion episodes are preceded by signs of impending danger; whether detecting congestion or activating congestion avoidance responses can be done reliably far enough in advance of an actual congestion period requires further study.

In view of our analysis in Section 4 of Ethernet traces similar to the ones used by Fowler and Leland (1991), lessons (i)-(iv) are readily explained by the self-similar or fractal nature of the traffic used in the simulation (see Section 5.3.1 above). Moreover, the existence of long-term correlations in the data (see (3.1.1)) or, equivalently, the $1/f$-noise property of the power spectral density around the origin (see (3.3.1)), is good news for more accurate

predictions of congestion. For methods to determine the low frequency components, see Section 4. As in the previous section, it is the self-similar property of Ethernet LAN traffic that gives rise to the observed nature of congestion and, at the same time, provides an intuitive explanation for it. While many formal standard network traffic models provably show that congestion control "works" (e.g., large buffers provide protection against congestion, average loss rates are a sensible quality of service measure, and predicting bursts in the offered traffic is difficult), self-similar traffic models reveal a far more complex and challenging picture for B-ISDN congestion management.

5.4 SELF-SIMILAR TRAFFIC: PARSIMONIOUS MODELS AND SYNTHETIC TRACES

Self-similarity is often explained as being equivalent to the existence of a multilevel hierarchy of underlying mechanisms. While it is tempting to invoke such multilevel mechanisms to account for self-similarity in packet traffic (e.g., think of the 7 Open Systems Interconnection (OSI) levels), it is practically impossible to demonstrate why their contributions should result, for example, in an asymptotic power law for the autocorrelations of the form (3.1.1). Even if their physical reality could be established, the resulting models for packet traffic are likely to have a large number of parameters. Similarly, we have seen in Section 5.1 that conventional modeling approaches that stress the importance of source modeling are bound to produce highly overparameterized models for aggregate traffic. Below, we mention two very promising approaches in order to deal with *fractal properties* of packet traffic, i.e., with properties such as self-similarity (see Section 4), long-range-dependence (see Beran et al. (1992)), infinite variances (see Meier-Hellstern et al. (1991)), and fractal dimensions (see Erramilli and Singh (1992)) observed in actual packet traffic measurements. The two approaches are (i) stochastic models based on self-similar processes and (ii) deterministic models using nonlinear chaotic maps. Both methods emphasize the need for parsimonious modeling, i.e., for models with a small number of parameters, where every parameter can be given a physically meaningful interpretation, and the desire for generating synthetic traffic easily and quickly. For related work on "fractal" arrival processes, see the recent paper by Veitch (1992) whose approach is motivated by the seminal paper of Mandelbrot (1965).

5.4.1 STOCHASTIC MODELS: SELF-SIMILAR PROCESSES

As we have noted in Section 4, exactly self-similar models such as fractional Gaussian noise (or some nonlinear transformation of fractional Gaussian noise) or asymptotically self-similar models such as fractional ARIMA processes can be used to fit hour-long traces of Ethernet traffic very well. Note that fractional Gaussian noise is characterized by its mean μ, variance σ^2, and Hurst parameter H. Each of these 3 parameters has an obvious physical interpretation in light of our discussion in Section 5.2. At times, we have seen indications of a particular short-range dependence structure in the traffic measurements in which case asymptotically self-similar models such as the fractional ARIMA(p,d,q) seem to be more appropriate. That is, by adding a few parameters (typically 1 or 2), we are not only able to fit the low-frequency components in the data but also to capture the contributions of the high-frequency components.

Although our analysis in Section 4 has shown that the Hurst parameter can be expected to change during a measurement period of an hour or more, fractional Gaussian noise or fractional ARIMA models describe the statistical properties of Ethernet traffic more accurately than currently proposed models. Refinements such as modeling the change points of H may be needed in order to produce more realistic traffic models in the future, but for the time being, we suggest using either fractional Gaussian noise-based models or an appropriately chosen fractional ARIMA model to describe the fractal nature of (internal or external) Ethernet LAN traffic. In both cases, parameter estimation techniques are known but they often turn out to be computationally too intensive in order to work for large data sets. However, we have illustrated in Section 4 how to estimate the Hurst parameter H for large data sets, and methods to adapt the existing parameter estimation techniques and to apply them to long time series are currently being studied (see Beran and Terrin (1992)).

An important requirement of practical traffic modeling is to generate synthetic data sequences that exhibit similar features as the measured traffic. While exact methods for generating synthetic traces from fractional Gaussian noise and fractional ARIMA models exist (see McLeod and Hippel (1978), and Hoskings (1981), respectively), they are, in general, only appropriate for short traces (about 1000 observations). For longer time series, short memory approximations have been proposed such as the *fast fractional Gaussian noise* by Mandelbrot (1971). However, such approximations also become often inappropriate when the sample size becomes exceedingly large. Here, we briefly mention two methods for generating asymptotically self-similar observations, and provide

approximate running times for producing a long time series (for more details, see Garrett et al. (1993)). The first method simulates the buffer occupancy in an $M/G/\infty$ queue, where the service time distribution G satisfies the heavy-tail condition (3.2.3), i.e., G has infinite variance. Cox (1984) showed that an infinite variance service time distribution results in an asymptotically self-similar buffer occupancy process, and he relates the tail-behavior of the former to the degree of self-similarity of the latter. Generating a time series of length 100000 this way requires about 2 hours of CPU-time on a Sun SPARCstation 2. The second method exploits a convergence result obtained by Granger (1980) who showed that when aggregating many simple AR(1)-processes, where the AR(1) parameters are chosen from a beta-distribution on $[0, 1]$ with shape parameters p and q, then the superposition process is asymptotically self-similar; Granger also showed that the Hurst parameter H depends linearly on the shape parameter q of the beta-distribution. This method is obviously tailor-made for parallel computers, and producing a synthetic trace of length 100000 on a MasPar MP-1216, a massively parallel computer with 16384 processors, takes about 5-10 minutes. In contrast, Hosking's method to produce 100000 observations from a fractional ARIMA$(0, d, 0)$ model requires about 10 hours of CPU time on a Sun SPARCstation 2.

5.4.2 DETERMINISTIC MODELS: NONLINEAR CHAOTIC MAPS

Another promising approach to modeling packet traffic was recently proposed by Erramilli and Singh (1990, 1992) and is based on deterministic *chaotic* maps. Chaos is a dynamical system phenomenon in which simple, low order, nonlinear deterministic equations can produce behavior that mimics random processes. To illustrate the underlying idea, consider a nonlinear map $f(\cdot)$ that describes the evolution of a state variable $x_n \in (0, 1)$ over discrete time as $x_{n+1} = f(x_n)$. The packet generation process for an individual source can now be modeled by stipulating that the source generates one or no packet at time n depending on whether x_n is above or below an appropriately chosen threshold. If f is a chaotic map, the resulting packet process can mimic complex packet traffic phenomena. In particular, Erramilli and Singh (1992) have shown that a simple, two parameter nonlinear chaotic map referred to as the *intermittency map* can capture many of the above-mentioned fractal properties in actual packet traffic measurements (see also Erramilli and Willinger (1992) for additional evidence for the fractal behavior of packet traffic measurements).

Clearly, the generation of synthetic traffic via nonlinear chaotic maps makes the dynamical system approach to packet traffic modeling particularly appealing. Once an appropriate chaotic map has been derived from a set of traffic measurements, generating a packet stream for an individual source is generally quick and easy. On the other hand, the inference part for dynamical systems, i.e., deriving an appropriate nonlinear chaotic map based on a set of actual traffic measurements, currently requires considerable guessing and experimenting; developing more rigorous statistical estimation methods for dynamical systems has recently attracted considerable attention in the statistics literature (e.g., see Casdagli (1991), Smith (1991), Berliner (1992), and Chatterjee and Yilmaz (1992), and references therein).

Note that the motivation for both approaches, self-similar stochastic modeling and deterministic non-linear modeling, is the desire for a relatively simple description of the complex packet traffic generation process. Moreover, both modeling approaches yield a single parameter that describes the fractal nature of traffic (i.e, the Hurst parameter for self-similar models, and the fractal dimension for chaotic map models) and seems to capture the intuitive notion of "burstiness" where conventional measures of burstiness no longer apply. However, the connection on a mathematical level between the Hurst parameter of a process and its corresponding fractal dimension is not always clear. Finally, for purposes of performance analysis, both approaches pose new and challenging problems. While traditional performance modeling favors the use of stochastic input models, studying arrival streams to queues that are generated by non-linear chaotic maps may well provide new insight into the performance of queueing systems where the arrival processes exhibit fractal properties.

6. CONCLUSIONS

Understanding the nature of traffic in high-speed, high-bandwidth communications systems such as B-ISDN is essential for engineering, operations, and performance evaluation of these networks. In a first step toward this goal, it is important to know the traffic behavior of some of the expected major contributors to future high-speed network traffic. In this paper, we analyze LAN traffic offered to a high-speed public network supporting LAN interconnection, an important and rapidly growing B-ISDN service. The main findings of our statistical analysis of a few hundred million high quality, high time-resolution Ethernet LAN traffic measurements are that (i) Ethernet

LAN traffic is statistically self-similar, irrespective of when during the 4-year data collection period 1989-1992 the data were collected and where they were collected in the network, (ii) the degree of self-similarity measured in terms of the Hurst parameter H is typically a function of the overall utilization of the Ethernet and can be used for measuring the "burstiness" of the traffic (namely, the burstier the traffic the higher H), and (iii) major components of Ethernet LAN traffic such as external LAN traffic or external TCP traffic share the same self-similar characteristics as the overall LAN traffic.

We also mention two novel methods for modeling this self-similar or fractal nature of LAN traffic. One method is based on stochastic self-similar processes, the other relies on a dynamical systems approach and uses deterministic nonlinear chaotic maps. Both methods provide accurate and parsimonious models that can be easily used for generating synthetic network traffic. Refinements of both methods are possible and result in traffic models that are flexible enough to display the same variety of features observed in measured LAN data. None of the currently considered formal models for LAN traffic or, more generally, for packet traffic, is able to capture the fractal nature of real traffic. An important implication of the self-similarity characteristic of LAN traffic is that aggregating streams of such traffic does typically not produce a smooth (i.e., "Poisson-like") superposition process but instead, intensifies the burstiness (i.e., the degree of self-similarity) of the aggregation process.

Implications of the self-similar nature of packet traffic for engineering, operations, and performance evaluation of high-speed networks are ample: (i) source models for individual Ethernet users are expected to show extreme variability in terms of interarrival times of packets (i.e., the infinite variance syndrome), (ii) commonly used measures for "burstiness" such as the index of dispersion (for counts), the peak-to-mean-ratio, or the coefficient of variation (for interarrival times) are no longer meaningful for self-similar traffic but can be replaced by the Hurst parameter, and (iii) the nature of congestion produced by self-similar network traffic models differs drastically from that predicted by standard formal models and displays a far more complicated picture than has been typically assumed in the past. Finally, in light of the same fractal-like behavior recently observed in VBR video traffic - another major contributor to future high-speed network traffic - the more complicated nature of congestion due to the self-similar traffic behavior can be expected to persist even when we move toward a more heterogeneous B-ISDN environment. Thus, we believe based on our measured traffic data that the success or failure of, for example, a proposed congestion control scheme for B-ISDN will depend on how well it performs under a self-similar rather than under one of the standard formal traffic scenarios.

ACKNOWLEDGEMENTS

This work could not have been done without the help of J. Beran and R. Sherman who provided the S-functions that made the statistical analysis of an abundance of data possible. We also acknowledge many helpful discussions with A. Erramilli about his dynamical systems approach to packet traffic modeling.

REFERENCES

1. D. Anick, D. Mitra, M.M. Sondhi, "Stochastic Theory od a Data-Handling System with Multiple Sources", *Bell System Technical Journal 61*, 1871-1894, 1982.
2. J. Beran, "Estimation, Testing and Prediction for Self-Similar and Related Processes", PhD Thesis, ETH Zurich, Switzerland, 1986.
3. J. Beran, "Statistical Methods for Data with Long-Range Dependence", *Statistical Science 7*, No. 4, 1992.
4. J. Beran, R. Sherman, M. S. Taqqu, W. Willinger, "Variable-Bit-Rate Video Traffic and Long-Range Dependence", accepted for publication in *IEEE Trans. on Networking*, subject to revisions, 1992.
5. J. Beran, N. Terrin, "A Multivariate Central limit Theorem for Long-Memory Processes with Statistical Applications", preprint, 1992.
6. L. M. Berliner, "Statistics, Probability and Chaos", *Statistical Science 7*, 69-90, 1992.

7. D. R. Boggs, J. C. Mogul, C. A. Kent, "Measured Capacity of an Ethernet: Myths and Reality", *Proceedings of SIGCOMM'88, Computer Communications Review 18*, 222-234, 1988.
8. G. E. P. Box, G. M. Jenkins, *Time Series Analysis: Forecasting and Control*, 2nd Ed., Holden Day, San Francisco, 1976.
9. M. Casdagli, "Chaos and Deterministic versus Stochastic Non-Linear Modelling", *J. Roy. Statist. Soc. 54, Series B*, 303-328, 1991.
10. M. Cassandro, G. Jona-Lasinio, "Critical Behavior and Probability Theory", *Advances in Physics 27*, 913-941, 1978.
11. S. Chatterjee, M. R. Yilmaz, "Chaos, Fractals and Statistic", *Statistical Science 7*, 49-68, 1992.
12. D. R. Cox, "Long-Range Dependence: A Review", in: *Statistics: An Appraisal*, H. A. David and H. T. David (Eds.), The Iowa State University Press, Ames, Iowa, 55-74, 1984.
13. D. R. Cox, M. W. H. Townsend, "The Use of the Correlogram in Measuring Yarn Irregularity", *Proceedings of the International Wool Textile Organization*, Tech. Committee 2, 28-34, 1948.
14. R. Dahlhaus, "Efficient Parameter Estimation for Self-Similar Processes", *Ann. Statist. 17*, 1749-1766, 1989.
15. A. Erramilli, R. P. Singh, "Application of Deterministic Chaotic Maps to Model Packet Traffic in Broadband Networks", *Proc. 7th ITC Specialists Seminar*, Morristown, NJ, 8.1.1-8.1.3, 1990.
16. A. Erramilli, R. P. Singh, "An Application of Deterministic Chaotic Maps to Characterize Packet Traffic", preprint, 1992.
17. A. Erramilli, W. Willinger, "Fractal Properties of Packet Traffic Measurements", preprint, 1993.
18. D. Feldmeier, "Traffic Measurements of a Token ring Network", *Proceedings of the IEEE Computer Network Symposium*, 236-243, 1986.
19. K. W. Fendick, V. R. Saksena, W. Whitt, "Investigating Dependence in Packet Queues with the Index of Dispersion for Work", *IEEE Trans. Communications 39*, 1231-1244, 1991.
20. H. J. Fowler, W. E. Leland, "Local Area Network Traffic Characteristics, with Implications for Broadband Network Congestion Management", *IEEE Journal on Selected Areas in Communications 9*, 1139-1149, 1991.
21. R. Fox, M. S. Taqqu, "Large-Sample Properties of Parameter Estimates for Strongly Dependent Stationary Gaussian Time Series", *Ann. Statist. 14*, 517-532, 1986.
22. M. W. Garrett, G. Latouche, R. Sherman, W. Willinger, "Fast Methods for Generating Long Sequences of Data with Long-Range Dependence" (in preparation, 1993).
23. M. W. Garrett, W. Willinger, "Statistical Analysis of a Long Sample of Variable Bit Rate Video Traffic", preprint, 1993.
24. L. Giraitis, D. Surgailis, "On Shot Noise Processes Attracted to Fractional Levy Motion", in: *Stable Processes and Related Topics*, S. Cambanis, G. Samorodnitsky, and M. S. Taqqu (Eds.), Birkhauser, Boston, 261-273, 1991.
25. C. W. J. Granger, "Long Memory Relationships and the Aggregation of Dynamic Models", *J. Econometr. 14*, 227-238, 1980.
26. C. W. J. Granger, R. Joyeux, "An Introduction to Long-Memory Time Series Models and Fractional Differencing", *J. Time Series Anal. 1*, 15-29, 1980.
27. R. Gusella, "A Measurement Study of Diskless Workstation Traffic on an Ethernet", *IEEE Transactions on Communications 38*, 1557-1568, 1990.
28. F. R. Hampel, "Data Analysis and Self-Similar Processes", *Proc. 46th Session ISI*, Book 4, 235-254, 1987.
29. F. R. Hampel, E. M. Ronchetti, P. J. Rousseeuw, W. A. Stahel, *Robust Statistics. The Approach Based on Influence Functions*, Wiley, New York, 1986.

30. H. Heffes, D. M. Lucantoni, "A Markov Modulated Characterization of Packetized Voice and Data Traffic and Related Statistical Multiplexer Performance", *IEEE Journal on Selected Areas in Communications 4*, 856-868, 1986.

31. J. R. M. Hosking, "Fractional Differencing", *Biometrika 68*, 165-176, 1981.

32. H. E. Hurst, "Long-Term Storage Capacity of Reservoirs", *Trans. Amer. Soc. Civil Engineers 116*, 770-799, 1951.

33. H. E. Hurst, "Methods of Using Long-Term Storage in Reservoirs", *Proc. Institution Civil Engineers*, Part I, 519-577, 1955.

34. R. Jain, S. A. Routhier, "Packet Trains: Measurements and a New Model for Computer Network Traffic", *IEEE Journal on Selected Areas in Communications 4*, 986-995, 1986.

35. V. Klemes, "The Hurst Phenomenon: A Puzzle?", *Water Resources Research 10*, 675-688, 1974.

36. H. R. Kuensch, "Statistical Aspects of Self-Similar Processes", *Proc. 1st World Congress of the Bernoulli Society*, Tashkent, Vol. 1, 67-74, 1986.

37. A. J. Lawrance, N. T. Kotegoda, "Stochastic Modelling of Riverflow Time Series (with Discussions)", *J. Roy. Statist. Soc. 140, Series A*, 1-47, 1977.

38. W. E. Leland, "LAN Traffic Behavior from Milliseconds to Days", *Proceedings of the ITC 7th Specialist Seminar*, Morristown, NJ, 6.1.1-6.1.6, 1990.

39. W. E. Leland, D. V. Wilson, "High Time-Resolution Measurement and Analysis of LAN Traffic: Implications for LAN interconnection", *Proceedings of the IEEE INFOCOM'91*, Bal Harbour, FL, 1360-1366, 1991.

40. J. B. Levy, M. S. Taqqu, "On Renewal Processes Having Stable Inter-Renewal Intervals and Stable Rewards", *Ann. Sc. Math. Quebec 11*, 95-110, 1987.

41. A. W. Lo, "Long-Term Memory in Stock Market Prices", *Econometrica 59*, 1279-1313, 1991.

42. B. B. Mandelbrot, "Self-Similar Error Clusters in Communication Systems and the Concept of Conditional Stationarity", *IEEE Trans. Communications Technology COM-13*, 71-90, 1965.

43. B. B. Mandelbrot, "Long-Run Linearity, Locally Gaussian Processes, H-Spectra and Infinite Variances", *Intern. Econom. Rev.* 10, 82-113, 1969.

44. B. B. Mandelbrot, "A Fast Fractional Gaussian Noise Generator", *Water Resources Research 7*, 543-553, 1971.

45. B. B. Mandelbrot, *The Fractal Geometry of Nature*, Freeman, New York, 1983.

46. B. B. Mandelbrot, M. S. Taqqu, "Robust R/S Analysis of Long Run Serial Correlation", *Proc. 42nd Session ISI*, vol. XLVIII, Book 2, 69-99, 1979.

47. B. B. Mandelbrot, J. W. Van Ness, "Fractional Brownian Motions, Fractional Noises and Applications", *SIAM Review 10*, 422-437, 1968.

48. B. B. Mandelbrot, J. R . Wallis, "Computer Experiments with Fractional Gaussian Noises", *Water Resources Research 5*, 228-267, 1969.

49. B. B. Mandelbrot, J. R. Wallis, "Some Long-Run Properties of Geophysical Records", *Water Resources Research 5*, 321-340, 1969.

50. A. I. McLeod, K. W. Hippel, "Preservation of the Rescaled Adjusted Range 1: A Reassessment of the Hurst Phenomenon", *Water Resources Research 14*, 491-508, 1978.

51. K. Meier-Hellstern, P. E. Wirth. Y-L Yan, D. A. Hoeflin, "Traffic Models for ISDN Data Users: Office Automation Application", in: *Teletraffic and Datatraffic in a Period of Change* (Proc. 13th ITC, Copenhagen, 1991), A. Jensen, V. B. Iversen (Eds.), North Holland, 167-172, 1991.

52. R. M. Metcalfe. D. R. Boggs, "Ethernet: Distributed Packet Switching for Local Computer Networks", *Communications of the ACM 19*, 395-404, 1976.

53. V. Ramaswami, "Traffic Performance Modeling for Packet Communication: Whence, Where and Whither", *Proc. 3rd Australian Teletraffic Seminar*, Adelaide, Australia, 1988.

54. J. F. Shoch, J. A. Hupp, "Measured Performance of an Ethernet Local Network", *Communications of the ACM 23*, 711-721, 1980.

55. W. D. Sincoskie, "Transparent Interconnection of Broadcast Networks", *Proceedings 1986 International Zurich Seminar on Digital Communications*, 131-134, March 1986.

56. R. L. Smith, "Estimating Dimensions in Noisy Chaotic Time Series", *J. Roy. Statist. Soc. 54, Series B*, 329-351, 1991.

57. F. Sowell, "Modeling Long Run Behavior with the Fractional ARIMA Model", manuscript, Graduate School of Industrial Administration, Carnegie Mellon University, 1990.

58. K. Sriram, W. Whitt, "Characterizing Superposition Arrival Processes in Packet Multiplexers for Voice and Data", *IEEE Journal of Selected Areas in Communications 4*, 833-846, 1989.

59. E. H. Spafford, "The Internet Worm Incident", *Proceedings of the 1989 European Software Engineering Conference (ESEC 89)*; see also: *Lecture Notes in Computer Science*, vol. 87, Springer-Verlag, New York, 1989.

60. M. S. Taqqu, "A Bibliographical Guide to Self-Similar Processes and Long-Range Dependence", in: *Dependence in Probability and Statistic*, E. Eberlein and M. S. Taqqu (Eds.), Birkhauser, Basel, 137-165, 1985.

61. M. S. Taqqu, "Self-Similar Processes", in: *Encyclopedia of Statistical Sciences*, Vol. 8, S. Kotz and N. Johnson (Eds.), Wiley, New York, 1987.

62. M. S. Taqqu, J. B. Levy, "Using Renewal Processes to Generate Long-Range Dependence and High Variability", in: *Dependence in Probability and Statistics*, E. Eberlein and M. S. Taqqu (Eds.), Progress in Prob. and Stat. Vol. 11, Birkhauser, Boston, 73-89, 1986.

63. D. Veitch, "Novel Models of Broadband Traffic", preprint, 1992.

64. P. Whittle, "Estimation and Information in Stationary Time Series", *Ark. Mat. 2*, 423-434, 1953.

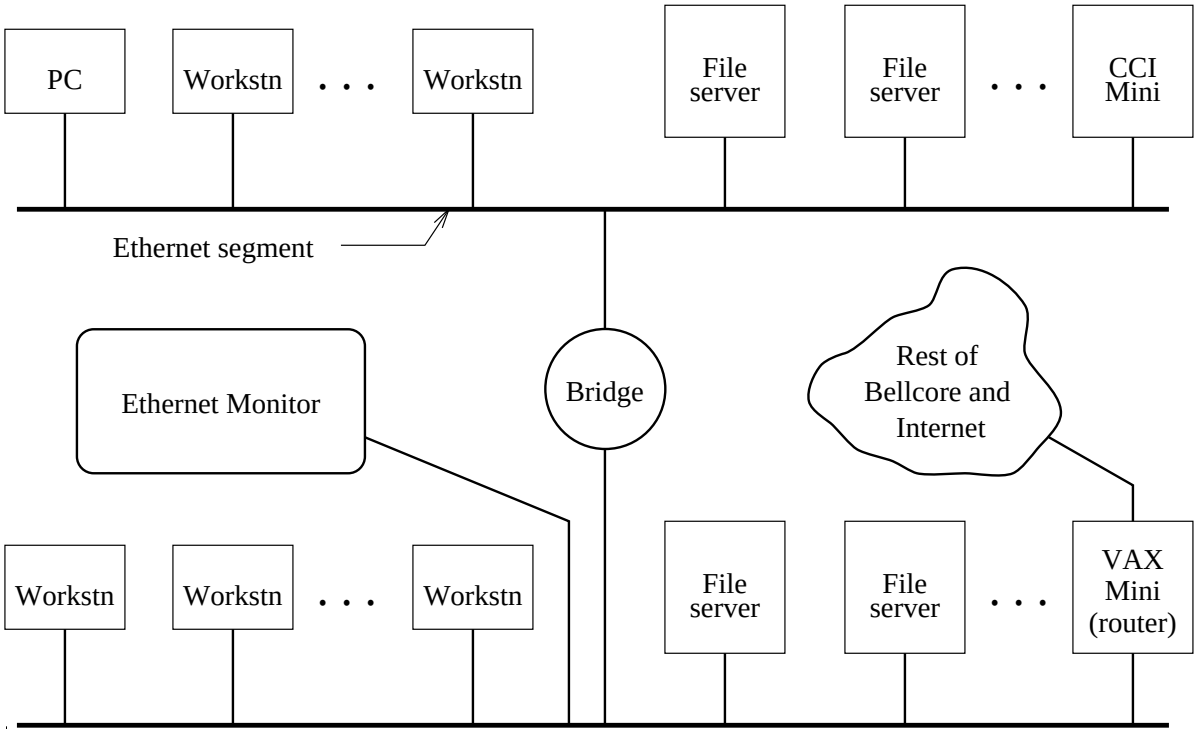

Figure 2.2.1. Network from which the August and October 1989 measurements were taken.

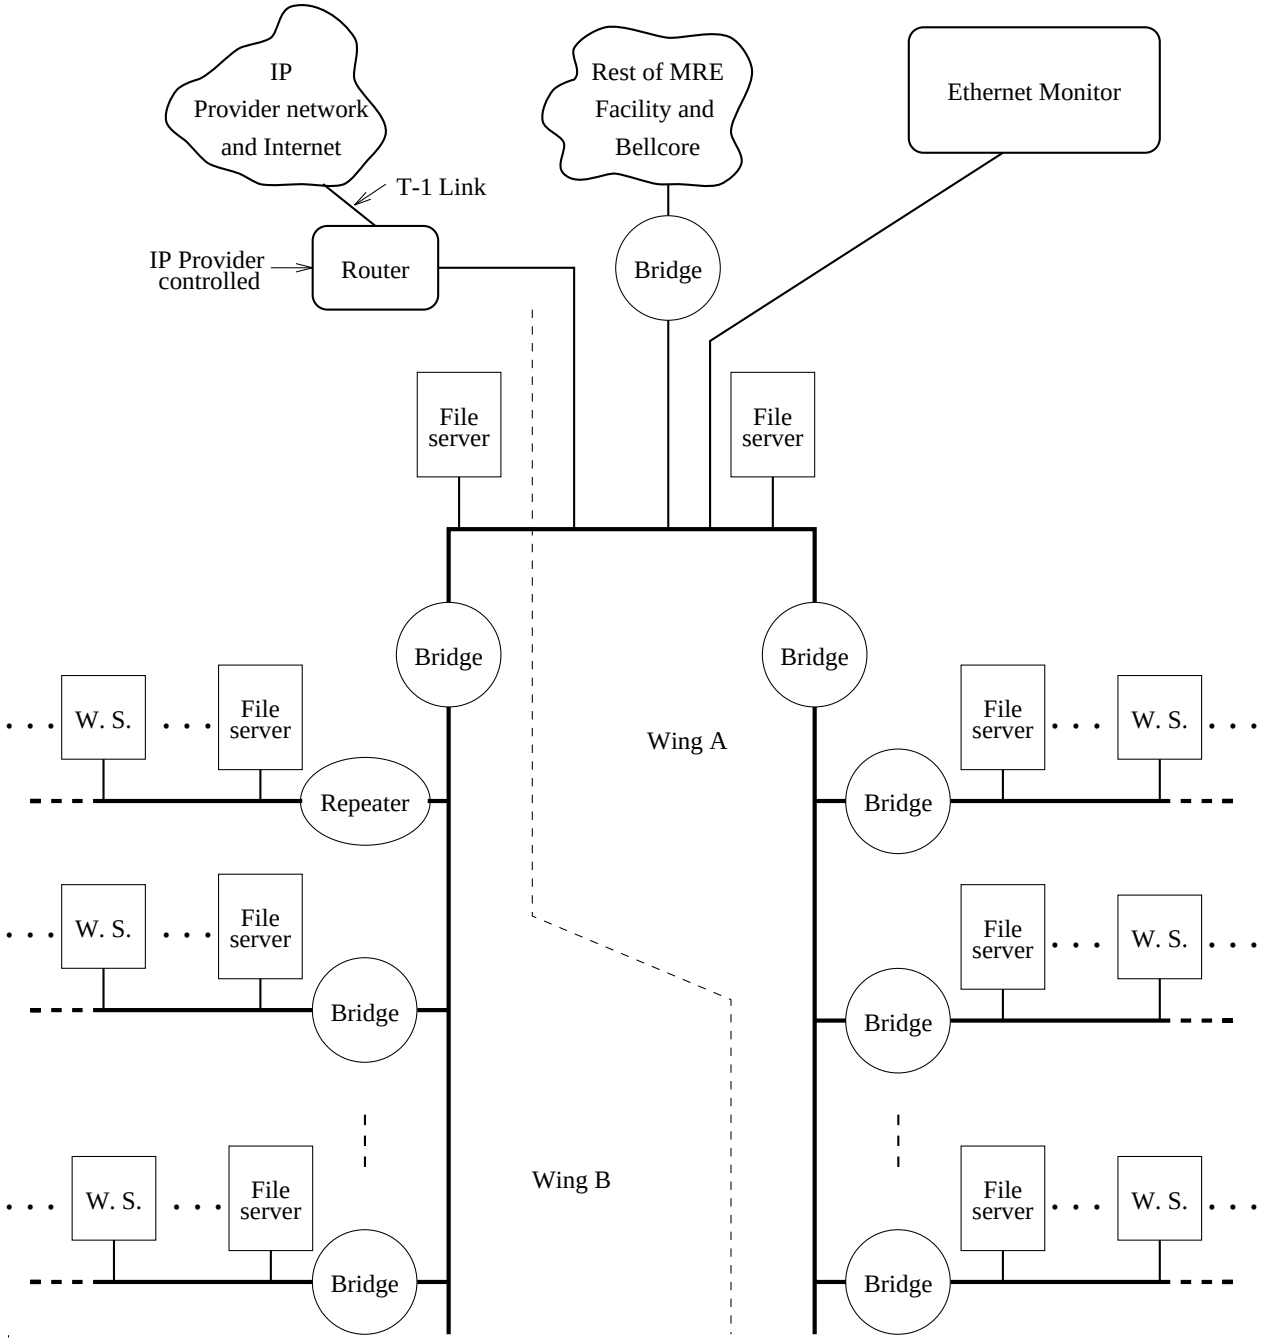

Figure 2.2.2. Network for second laboratory from which the January 1990 measurements were taken.

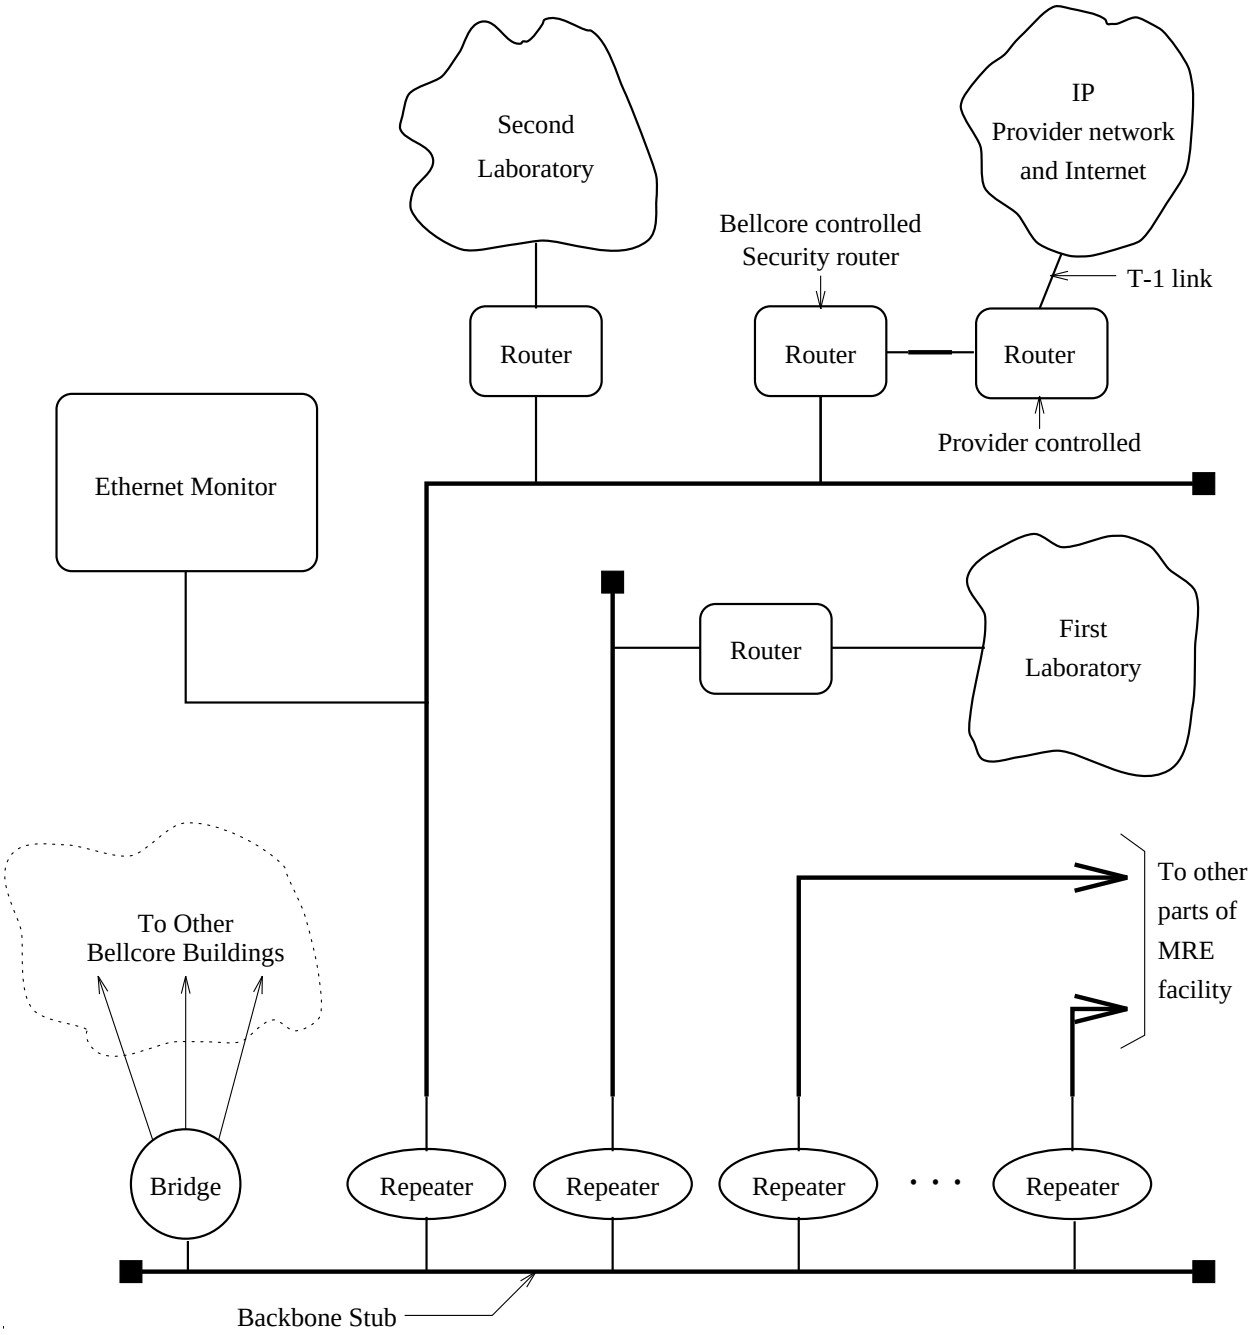

Figure 2.2.3. Backbone network for MRE facility from which the February 1992 measurements were taken.

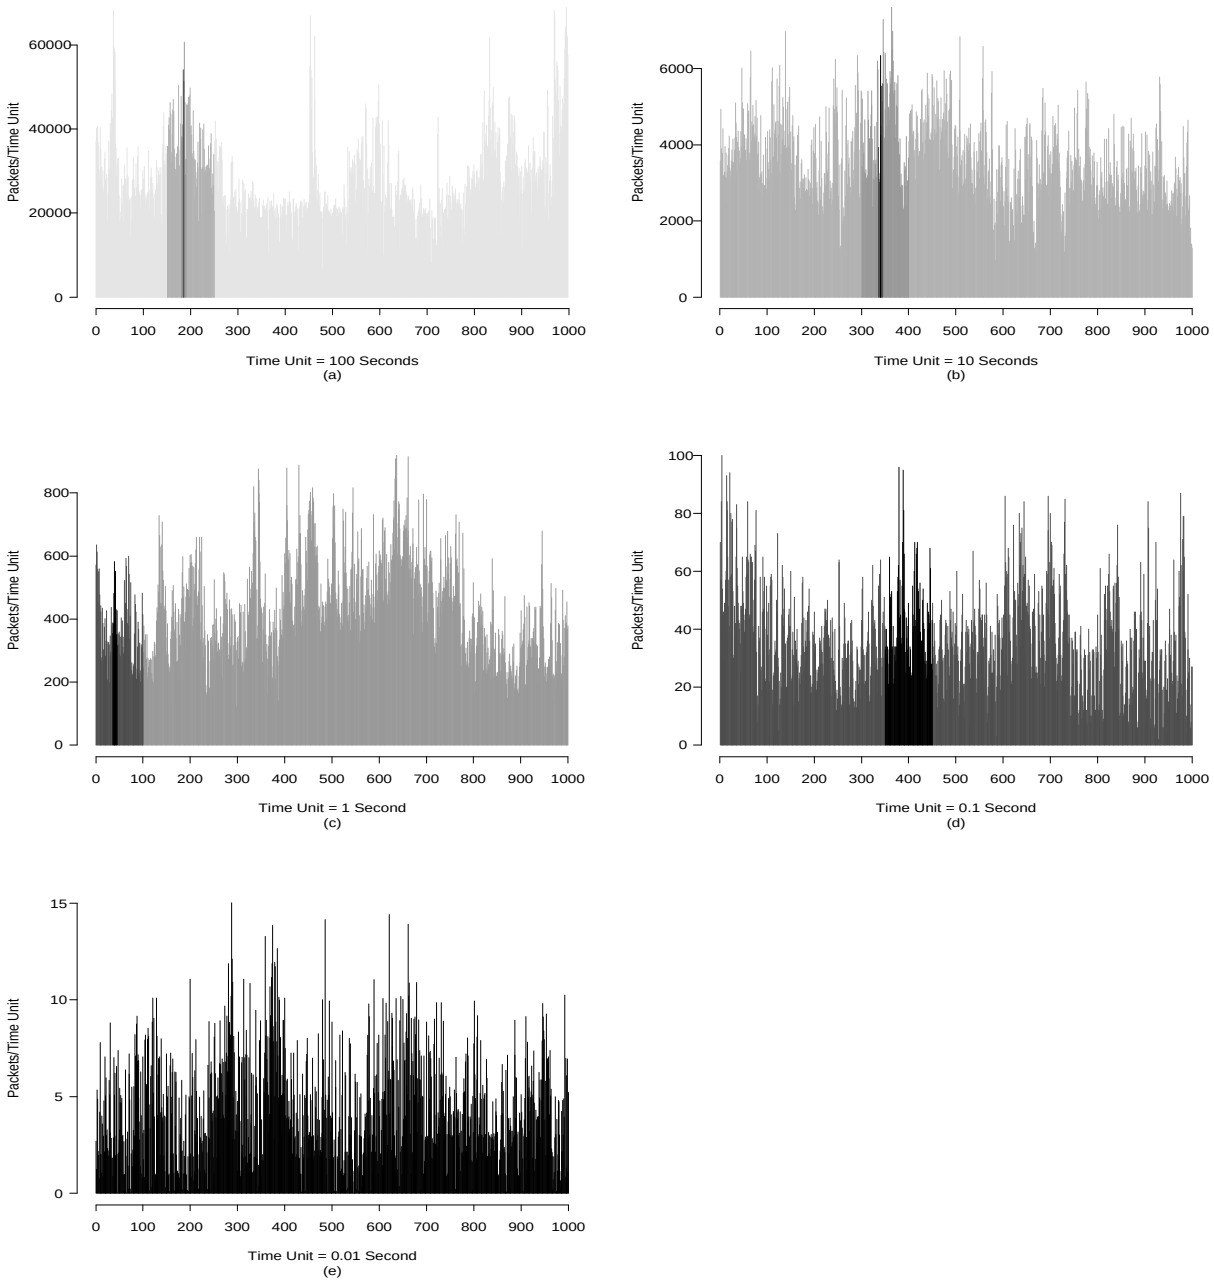

Figure 3.1 (a) - (e). "Pictorial" proof of self-similarity: Ethernet traffic (packets per time unit) on 5 different time scales. (Different gray levels are used to identify the same segments of traffic on the different time scales.)

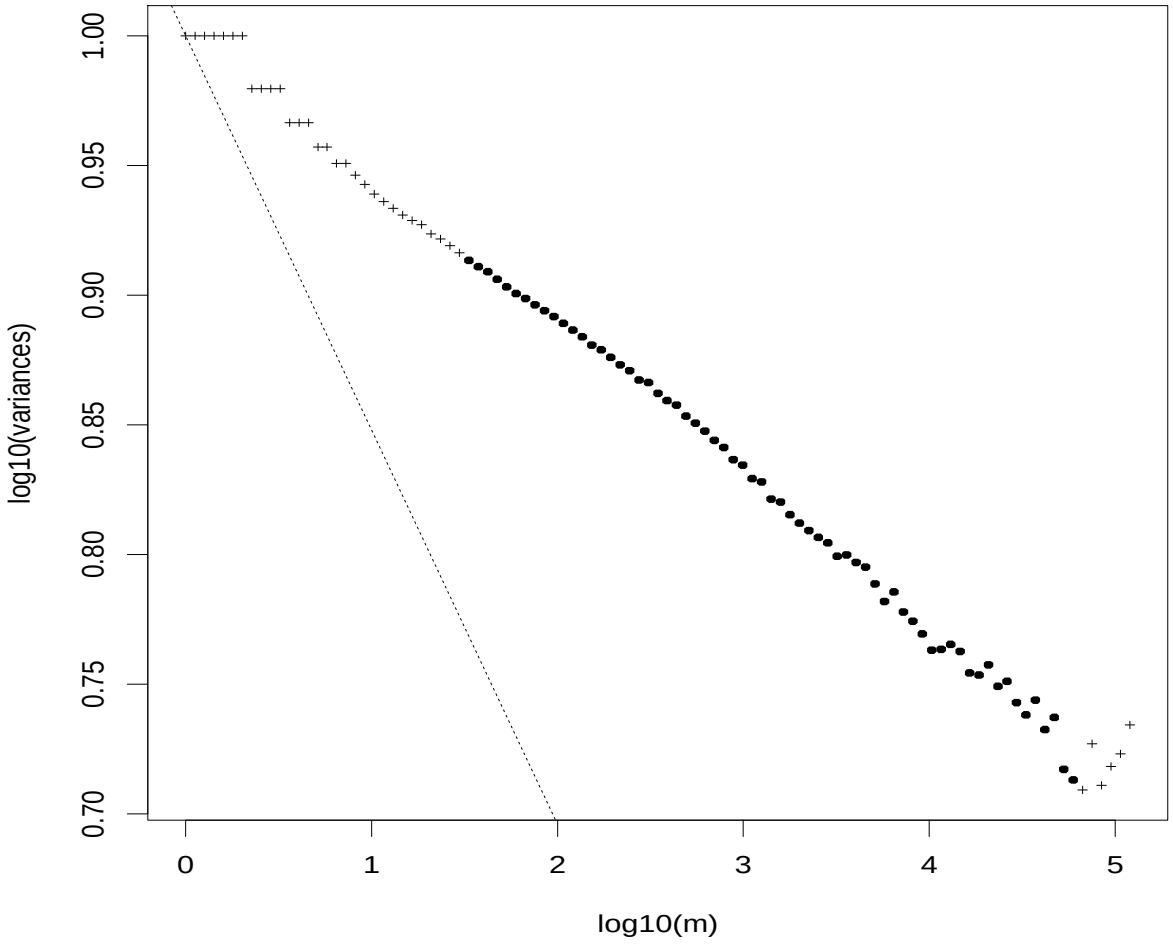

Figure 4.2.1 (a). Variance-time plot of sequence AUG89.MB. The asymptotic slope (determined using the "brushed" points) is clearly larger than the slope -1.0 of the dotted reference line and is readily estimated to be about -0.40.

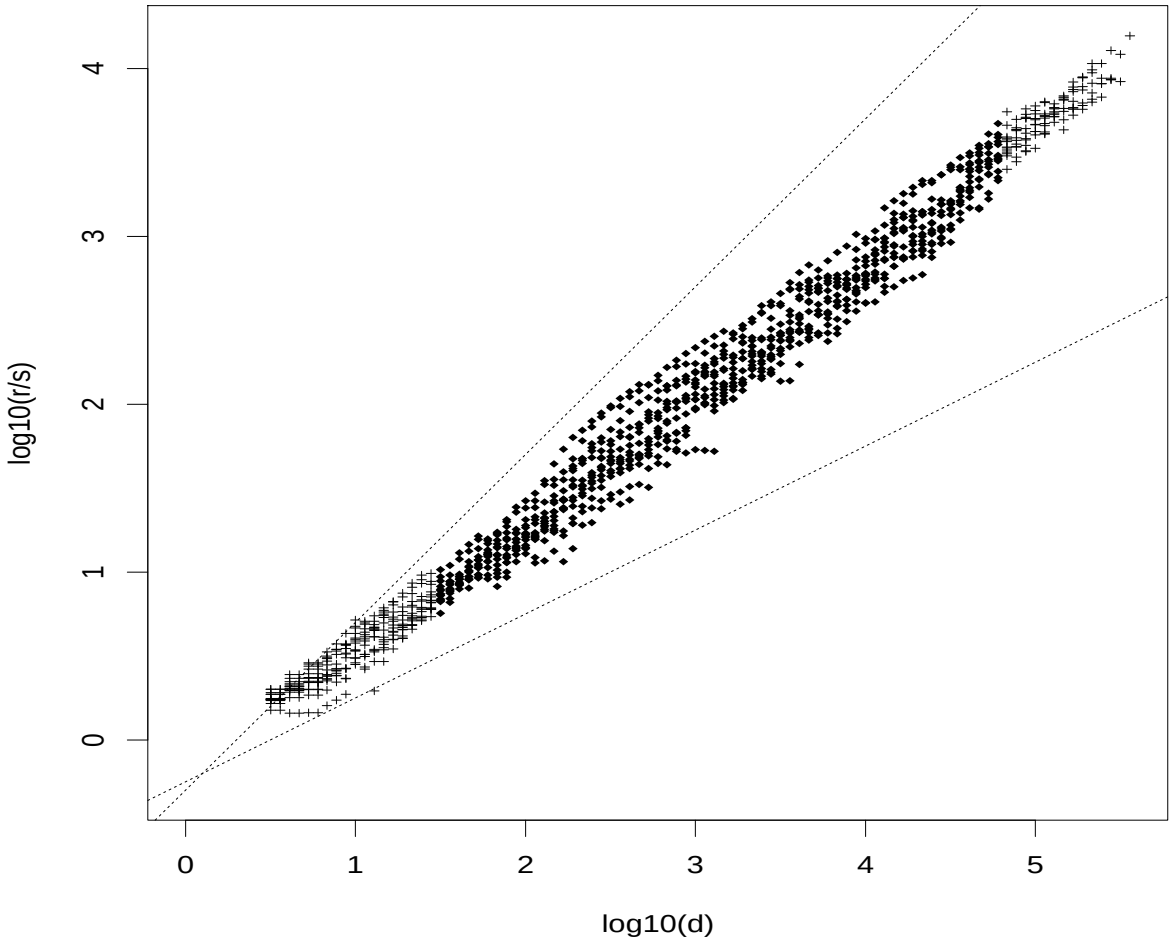

Figure 4.2.1 (b). Pox plot of R/S for sequence AUG89.MB. The plot tightly clusters around a straight line whose asymptotic slope clearly lies between the slopes 0.5 (lower dotted line) and 1.0 (upper dotted line) and is readily estimated (using the "brushed" points) to be about 0.79.

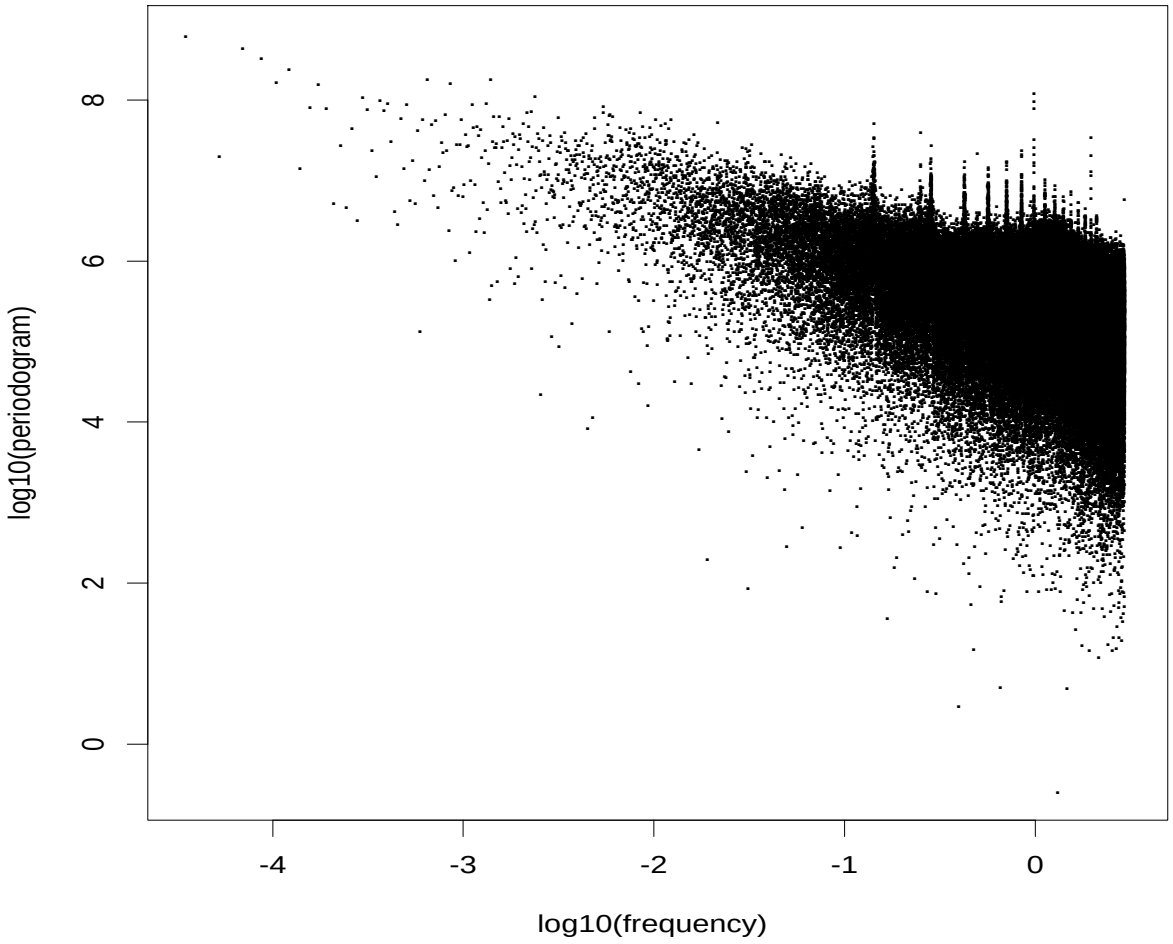

Figure 4.2.1 (c). Periodogram plot for sequence AUG89.MB. The slope of the low-frequency part of the plot is clearly different from 0. Using the lowest 10% of all frequencies, we obtain a slope estimate of about -0.64.

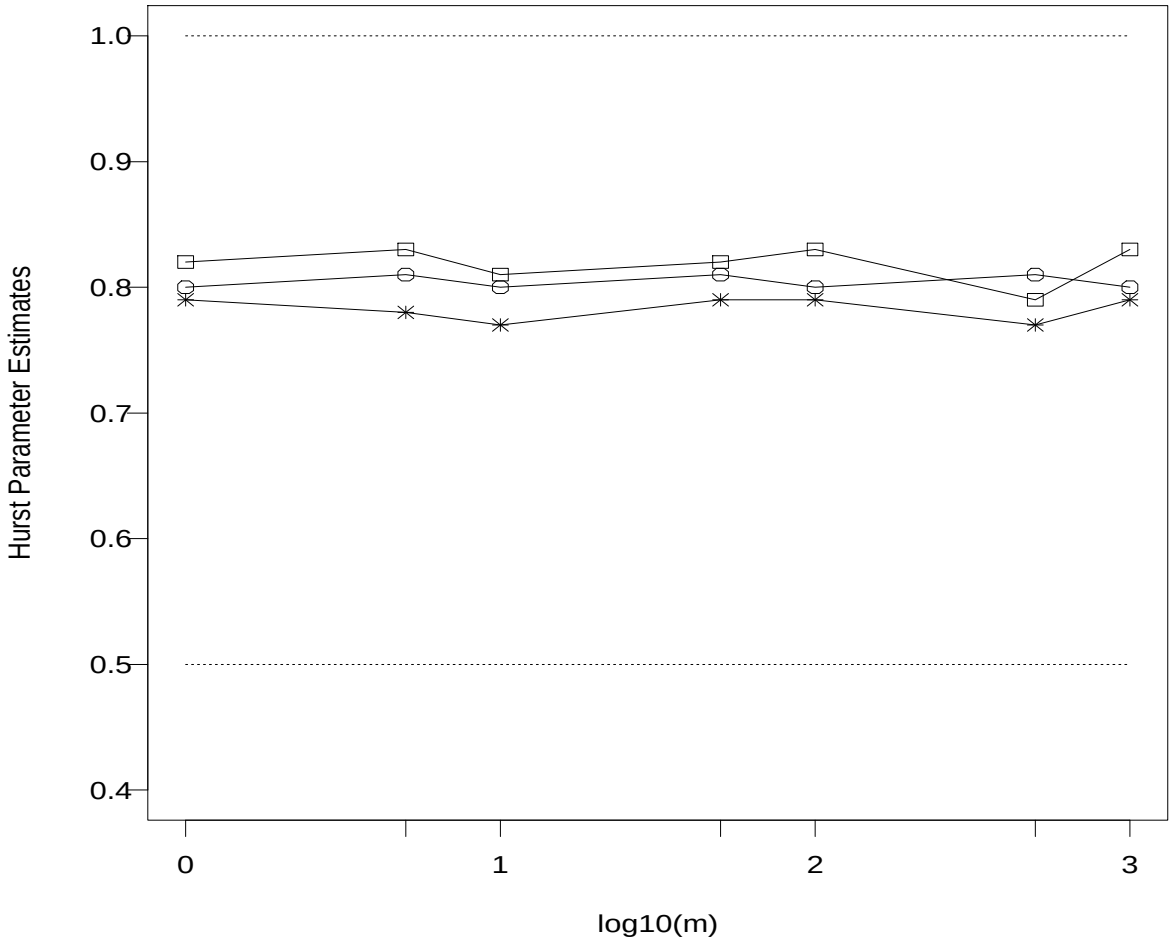

Figure 4.2.2. Estimates of the Hurst parameter H for sequence AUG89.MB, as a function of the aggregation level m. (○ estimate based on the asymptotic slope of the variance-time plot, ∗ estimate based on the the asymptotic slope of the pox plot of R/S, □ estimate based on the slope of the periodogram plot (using 10% of the lowest frequencies).)

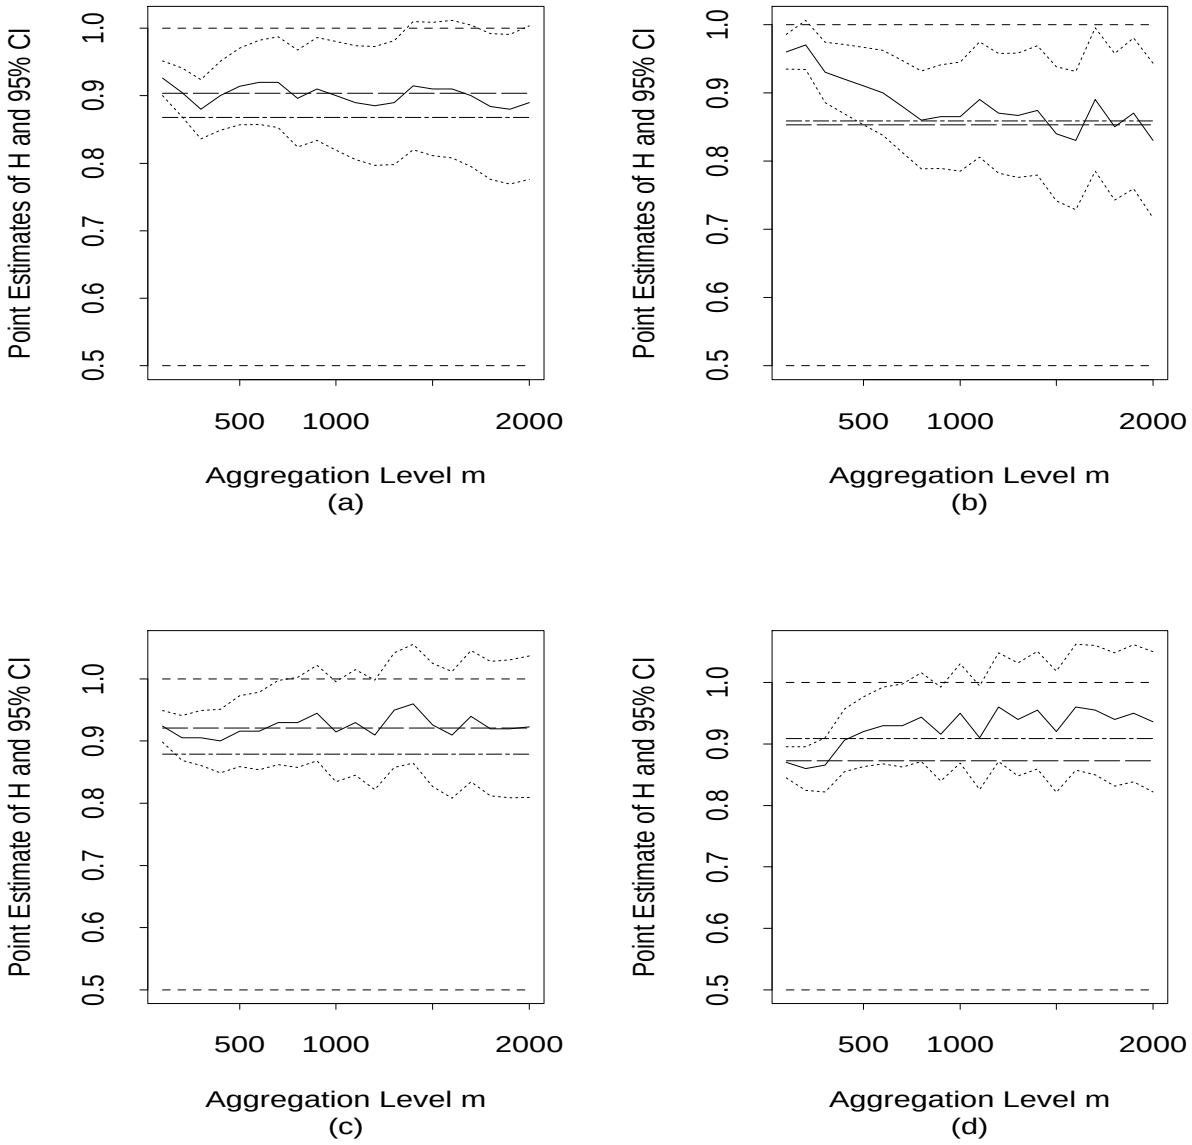

Figure 4.2.3 (a) - (d). Periodogram-based MLE estimate $\hat{H}^{(m)}$ of H (solid line) and 95%-confidence intervals (dotted lines), as a function of the aggregation level m, for sequences AUG89.MP (a), OCT89.MP (b), JAN90.MP (c), and FEB92.MP (d). For example, plot (a) shows that $m \approx 300$ is an appropriate aggregation level for sequence AUG89.MP, yielding a point estimate $H = \hat{H}^{(300)} = 0.90$ and a 95%-confidence interval [0.85,0.95]. For comparison, we also added to each plot the estimate of H based on the variance-time plot (–·–·–) and the R/S-based estimate of H (– – –).

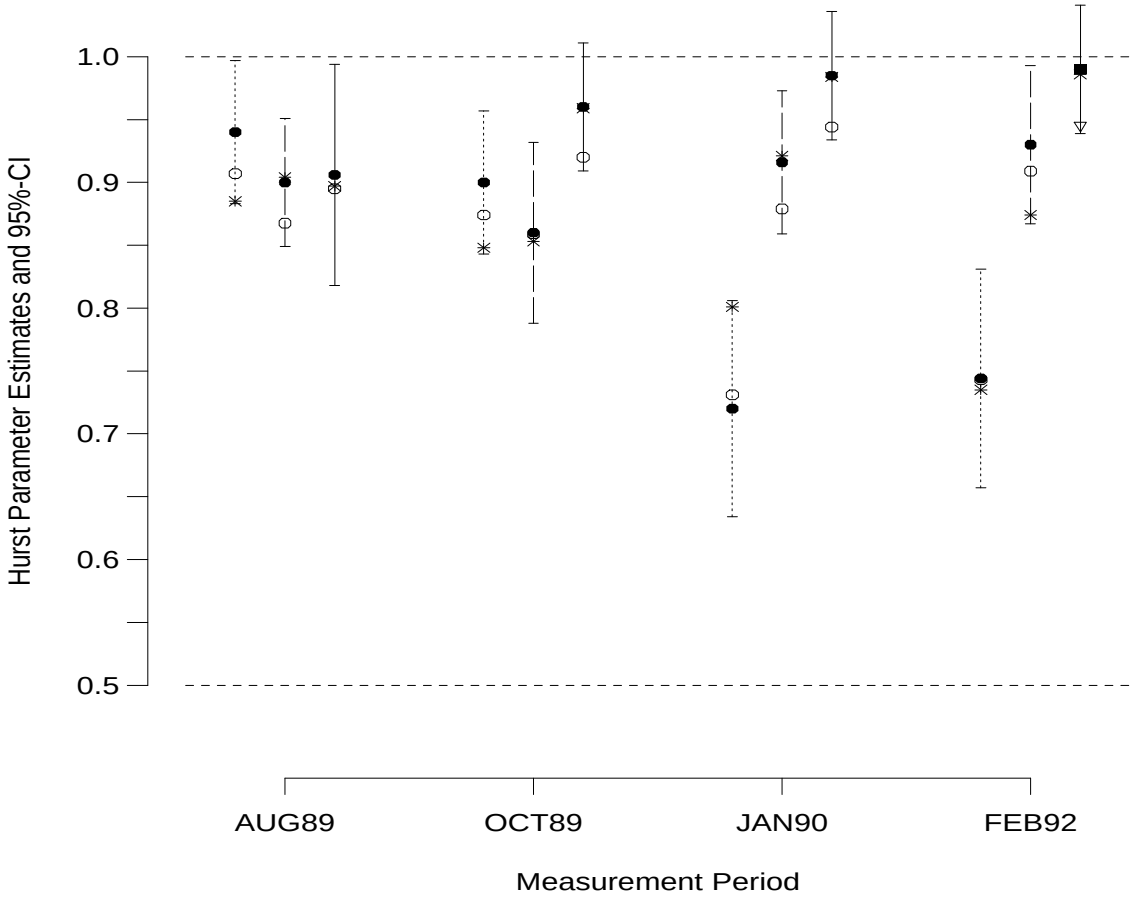

Figure 4.2.4 (a). Summary plot of estimates of the Hurst parameter H for all data sets (representing the number of *packets*) in Table 1. (● periodogram-based MLE estimate and corresponding 95%-confidence interval, ○ point estimate of H based on the asymptotic slope of the variance-time plot, ∗ point estimate of H based on the asymptotic slope of the pox plot of R/S.)

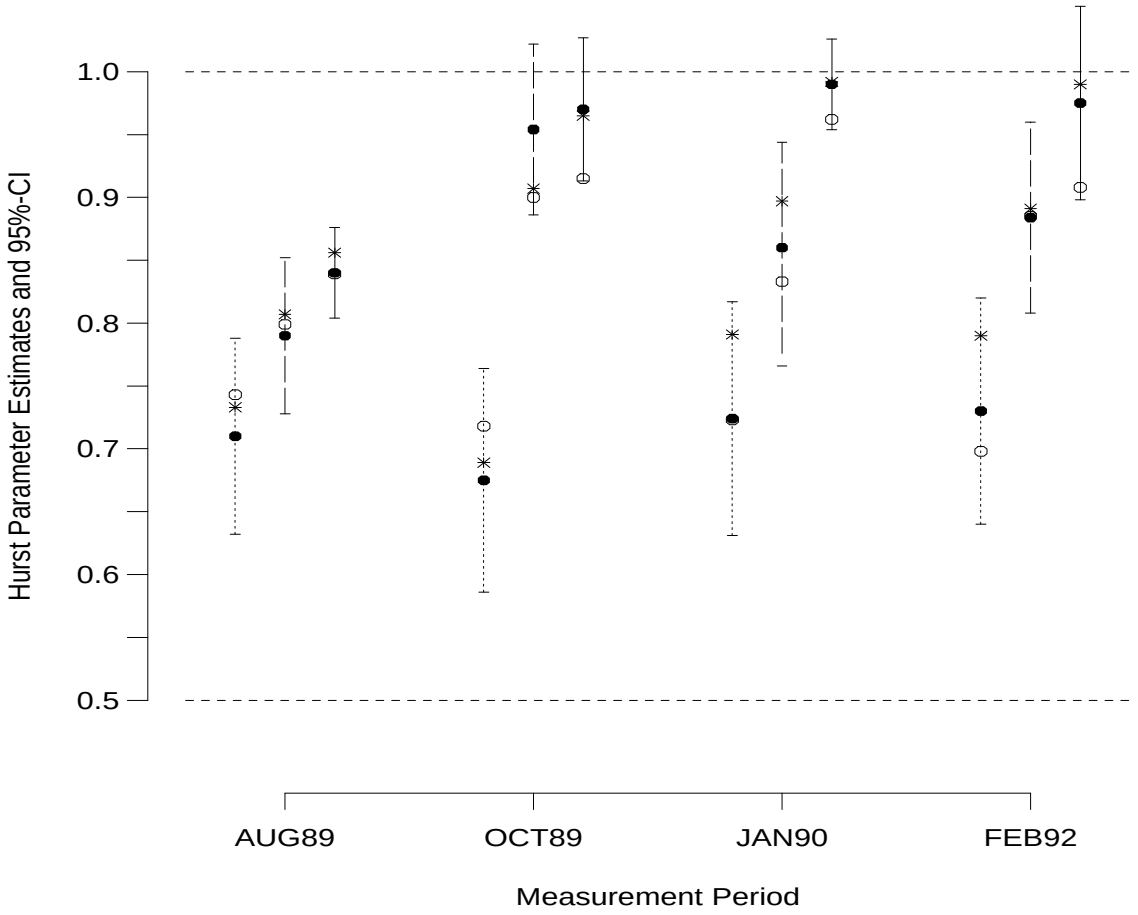

Figure 4.2.4 (b). Summary plot of estimates of the Hurst parameter H for all data sets (representing the number of *bytes*) in Table 1. (● periodogram-based MLE estimate and corresponding 95%-confidence interval, ○ point estimate of H based on the asymptotic slope of the variance-time plot, ∗ point estimate of H based on the asymptotic slope of the pox plot of R/S.)

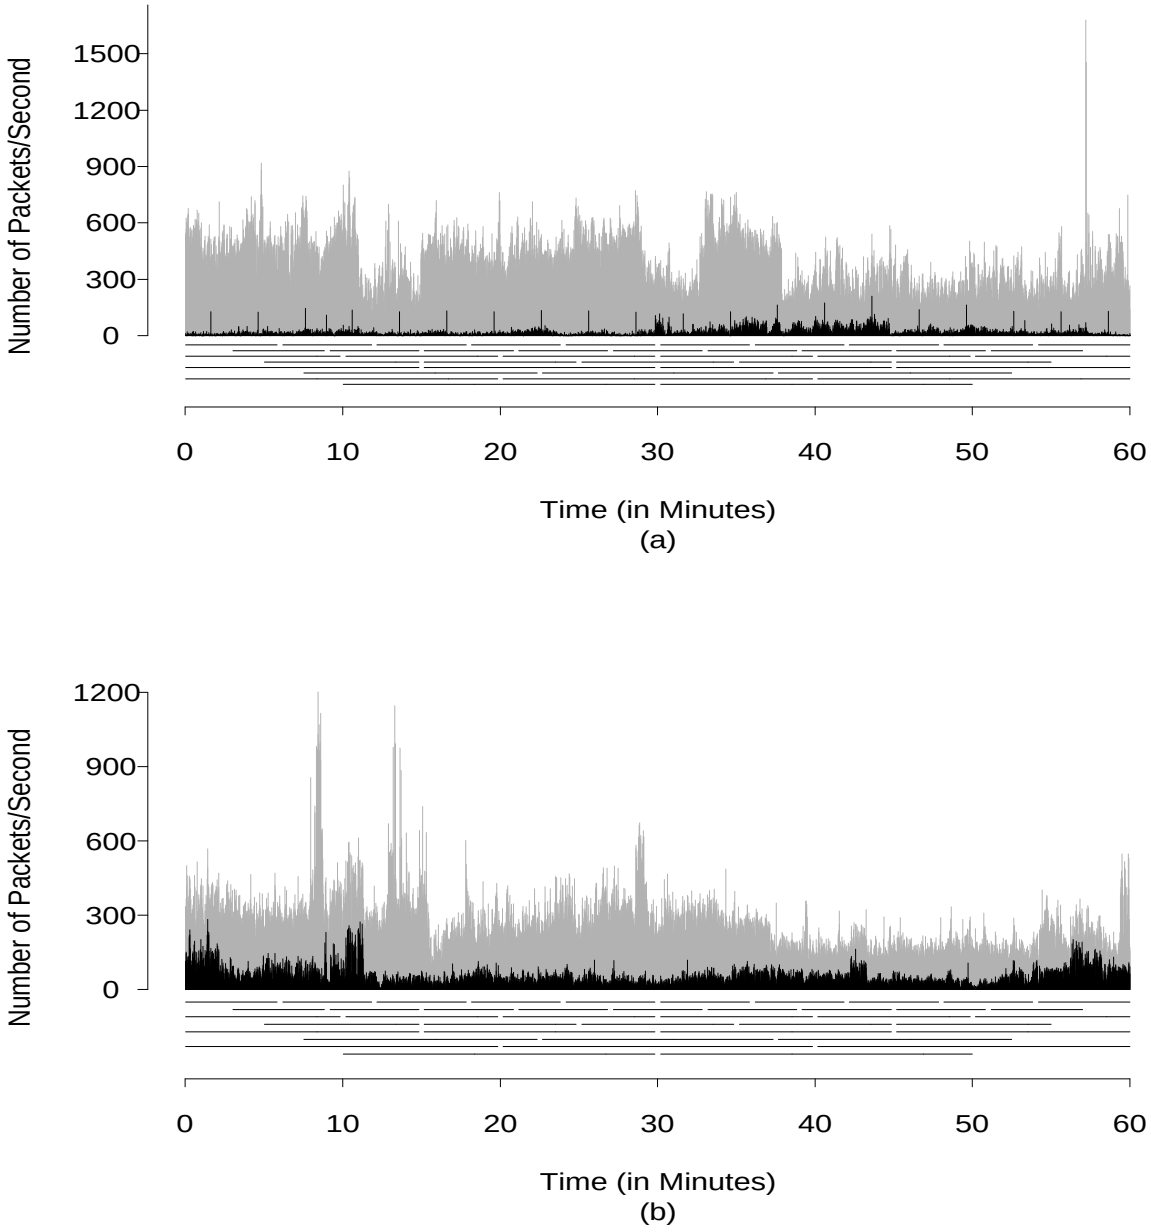

Figure 4.2.5 (a) - (b). Plots of the traffic (packets per second) during the busy hours of January ’90 (a) and February ’92 (b). The gray time series represents all Ethernet packets (i.e., internal traffic), and the portion in black depicts the amount of external traffic during these hours. Also indicated are the different ways of partitioning the data in disjoint blocks for testing the hypothesis that H is constant.

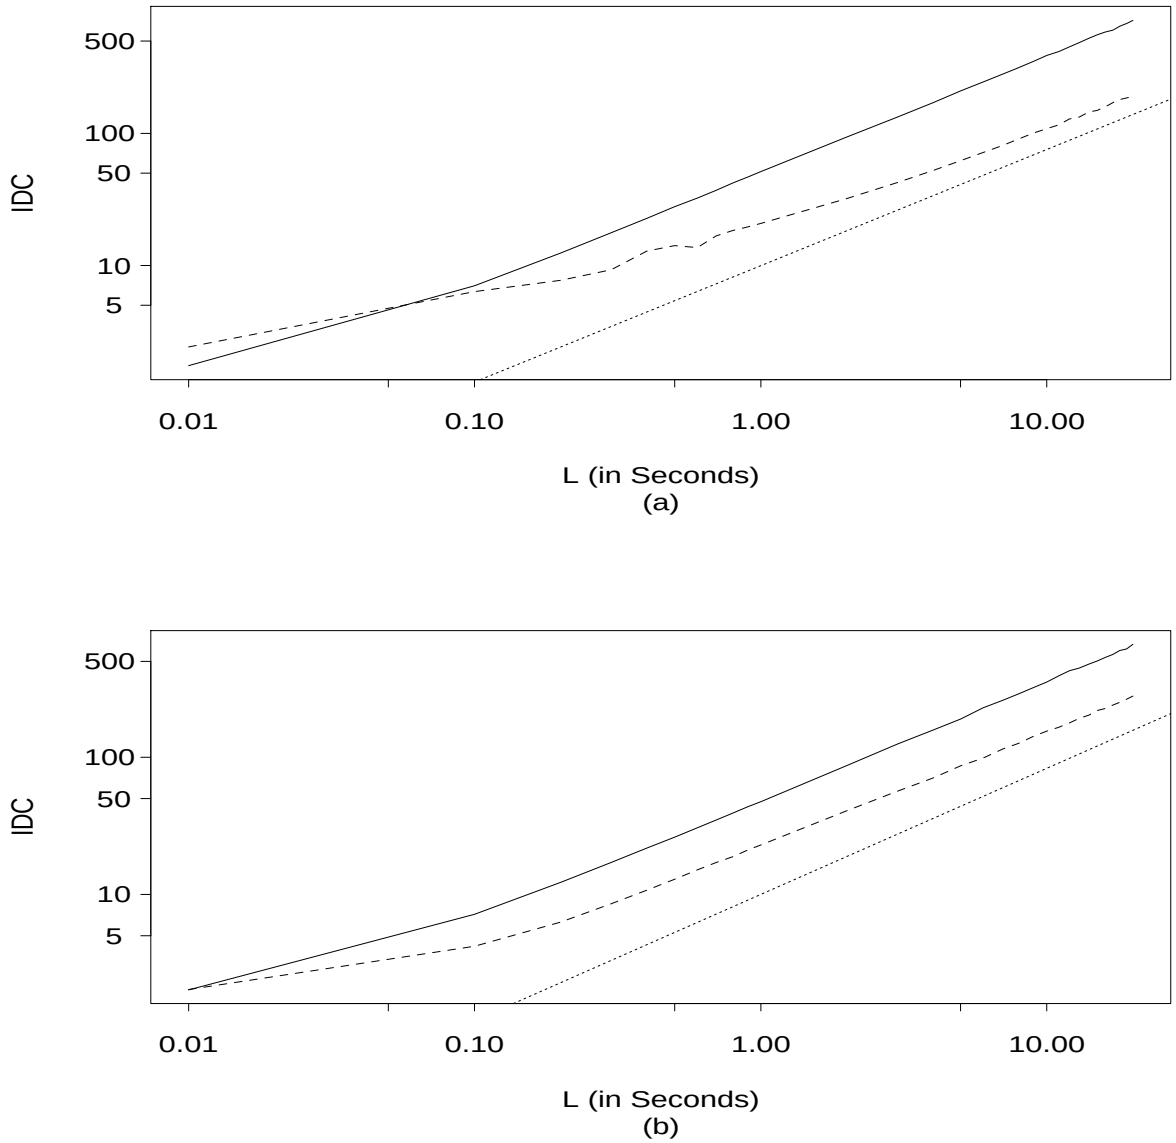

Figure 5.2.1 (a) - (b). Index of dispersion for counts (IDC) as a function of the length L of the time interval over which IDC is calculated, for the January '90 busy hour (a) and the February '92 busy hour (b). The solid lines are the IDC curves for the sequences JAN90.HP and FEB92.HP (internal traffic), the dashed lines depict the IDC curves for the sequences JAN90E.HP and FEB92E.HP (external traffic). The dotted lines are the IDC curves predicted by a self-similar model fitted to the time series JAN90.HP and FEB92.HP, respectively. Note that the asymptotic slopes of the solid lines agree with the slopes of the dotted lines.