

The Nature of the Beast

This research is supported by NWO-SWON Grant 613-03-039.

Rijksuniversiteit Groningen

The Nature of the Beast

*Analyzing and Modeling
Computer Network Traffic*

Proefschrift

ter verkrijging van het doctoraat in de
Wiskunde en Natuurwetenschappen
aan de Rijksuniversiteit Groningen
op gezag van de
Rector Magnificus, dr. D.F.J. Bosscher,
in het openbaar te verdedigen op
vrijdag 3 mei 2002
om 16.00 uur

door

Albertus Wilhelmus Stegeman

geboren op 25 juni 1975
te Deventer

Promotores: Prof. dr. H.G. Dehling
Prof. dr. T. Mikosch

Beoordelingscommissie: Prof. dr. O.J. Boxma
Prof. dr. W. Schaafsmma
Prof. dr. M.S. Taqqu

ISBN 90-367-1598-9



Every ending carries within itself a new beginning

Contents

1	Introduction	1
1.1	Computer networks	2
1.2	Measurements of computer network traffic	3
1.3	Statistical analysis of traffic measurements	5
1.3.1	Heavy tails	5
1.3.2	Long-range dependence	5
1.3.3	Non-stationarity or long-range dependence?	7
1.3.4	Self-similarity	7
1.3.5	Multifractals	9
1.3.6	Computer networks versus the telephone system	9
1.4	Modeling the workload in computer networks	10
1.4.1	Intuitive explanations	10
1.4.2	Mathematical explanations	10
1.5	Outline of the thesis	11
1.5.1	The Nature of the Beast	13
2	Computer Networks	15
2.1	Network hardware	16
2.1.1	Local Area Networks (LANs)	16
2.1.2	Wide Area Networks (WANs)	17
2.1.3	Inter Networks	18
2.2	Network software	18
2.3	Protocols in the Internet	21
2.3.1	Connection-oriented versus connectionless	21
2.3.2	The Internet network layer	22
2.3.3	The Internet transport layer	23
2.4	ATM networks	26
3	Mathematical Concepts	29
3.1	Heavy tails	29
3.1.1	Definition and examples	29
3.1.2	Detecting heavy tails	33
3.1.3	Stable Lévy motion	35
3.2	Long-range dependence	35

3.2.1	Definitions	35
3.2.2	Proofs	37
3.2.3	Detecting long-range dependence	39
3.2.4	Application to the Bellcore data	41
3.2.5	Long-range dependence or non-stationarity?	41
3.2.6	ARIMA models and long-range dependence	44
3.2.7	Fractional Brownian motion	45
3.3	Self-similarity	47
3.3.1	Definition	47
3.3.2	Self-similarity ‘by picture’	48
4	Traffic Data Analysis	51
4.1	Overview	52
4.2	LRD or ARIMA?	53
4.2.1	Two goodness-of-fit tests	56
4.2.2	Fitting BC-pAug to an ARIMA model	57
4.2.3	Analysis of other workload series	58
4.3	Visualizing non-stationarity	62
5	Modeling Computer Network Traffic	67
5.1	Heavy tails as the cause of LRD	68
5.1.1	LAN traffic: the ON/OFF model	68
5.1.2	The ON/OFF model and reality	71
5.1.3	WAN traffic: the infinite source Poisson model	72
5.1.4	The infinite source Poisson model and reality	73
5.2	Self-similar limits	74
5.2.1	Convergence to fractional Brownian motion	74
5.2.2	Convergence to stable Lévy motion	76
5.2.3	Simultaneous limit regimes	79
5.3	Network performance – a discussion	83
5.3.1	Queueing results	84
5.3.2	Simulation studies	86
5.A	Proof of Theorem 5.5: Slow Growth	88
5.A.1	The basic decomposition	88
5.A.2	Vanishing remainder terms	88
5.A.3	Convergence of the marginal distributions	94
5.A.4	Convergence of the finite-dimensional distributions	104
5.B	Proof of Theorem 5.5: Fast Growth	107
5.C	Large deviations of heavy-tailed sums	109
5.D	Bounds for regularly varying functions	111
6	Extremal Behavior of ON-Periods	113
6.1	Slow Growth	116
6.2	Fast Growth	120
6.3	A Central Limit Theorem	123
6.A	Proof of Theorem 6.1	125

6.B Proof of Theorem 6.3	129
6.C Proof of Theorem 6.5	134
6.D Convergence to a simple point process	137
6.E An identity in law for stopped random sums	138
Bibliography	141
Summary	151
Samenvatting	155
Dankwoord	159

One

Introduction

The key technological achievements of the last century were in the area of information and communication. Examples are the installation of worldwide telephone networks, the invention and widespread use of radio and television, the explosive developments in the computer industry and the launching of communication satellites. Due to fast technological progress the ability to gather, process and distribute information has grown rapidly. As a consequence, the demand for more sophisticated information processing systems has increased. This, in turn, has stimulated scientific research in this direction. The subject of this thesis is an example of such research: it considers the *statistical analysis and modeling of data traffic in computer networks*.

As surely anyone has experienced, the operation of transferring data between two computers in a network does not always end in success. Especially on the worldwide *Internet*, with hundreds of millions of users and many different kinds of applications, long delays and failing attempts to establish a connection are not uncommon. Among the research disciplines involved in examining and improving the performance of computer networks, statistics and mathematics play an important role. Since the beginning of the 1990s a large number of empirical studies of network traffic measurements has seen the light. The main conclusion of this body of work seems to be that the workload in computer networks is characterized by *long-range dependence* and *self-similarity*. Here, the workload is often defined as the number of bytes flowing through (a particular point in) the network per time unit. Also, it has been observed that transmission durations and file lengths seem to follow a *heavy-tailed* distribution. In general, these three canonical properties are believed to mark the difference between computer network traffic and voice traffic in the traditional telephone network. Attempts to incorporate these features in a mathematical framework have yielded two popular models for the workload in computer networks: the *ON/OFF model* and the *infinite source Poisson model*. Using these models, it has been established that the above characteristics of computer network traffic are responsible for lower network performance, as compared to the telephone system, in terms of queuing and delay statistics.

This introductory chapter provides the framework within which this thesis

has been written. In Section 1.1 we briefly introduce the technology and terminology of computer networks. Section 1.2 describes various data sets that can be obtained from measurements of computer network traffic. Section 1.3 introduces the concepts of *heavy tails*, *long-range dependence* and *self-similarity* and discusses their presence in traffic measurements on computer networks. Section 1.4 provides a glance at the *ON/OFF model* and the *infinite source Poisson model*, in which the characteristics of computer network traffic have been incorporated. Finally, Section 1.5 gives an outline of the rest of the thesis.

1.1 Computer networks

A *computer network* is a collection of computers that are able to exchange information. Computer networks are often classified by their scale. In general, three types are distinguished: *Local Area Networks* (LANs), *Wide Area Networks* (WANs) and *Inter Networks*. LANs are privately-owned networks within a single building or campus of up to a few kilometers in size. A WAN spans a large geographical area, like a country or a continent. It contains two types of machines: *hosts* and *routers*. Hosts are intended for running user programmes (i.e. applications), while *routers* are used as intermediate machines for data sent by one host to another host. An Inter Network is a collection of interconnected networks, e.g. a number of LANs connected by a WAN. The best known example of an Inter Network is the *Internet*, which spans the whole planet. The most popular application on the Internet is the *World Wide Web* (WWW), a system of interlinked sites on which information (in the form of text, pictures, sound and video) is displayed.

Data is transmitted between computers in a network in the following way. The sending computer is referred to as the *source* computer, while the receiver of the data is known as the *destination* computer. When a file is sent from a source to a destination in the network, the various layers in the network software decompose the file into small *packets*. These packets are sent through a *physical medium* (e.g. copper wire or fiber optics) to the destination computer. When arrived without transmission errors, the network software puts the packets together and thus reconstructs the original file. The duration of the whole transmission depends on the size of the file, the speed allowed by the physical medium (i.e. the *bitrate* or *bandwidth*), the occurrence of transmission errors (the packets that are damaged or lost during the transport have to be sent again) and the state of the network, e.g. congestion, queues, defects, etc. The network software includes algorithms that collect information on the amount of congestion and queuing in the network and accordingly adapt the rate at which data is sent from the source computer. In this way, an attempt is made to prevent packet losses due to, for example, buffer overflows. The most well-known of such algorithms is the *Transmission Control Protocol* (TCP), which is used in the Internet and most WANs. It sets up a *connection* between the source and destination hosts and, depending on states of the destination host and the intermediate routers, determines the sending rate. A disadvantage

of the reliable transmission service offered by TCP is that the transmission duration is increased. For some applications, however, speed is more important than maintaining high quality (e.g. digitized voice or video) and an ‘unreliable’ service, i.e. without setting up a connection as above, is more appropriate. In the Internet this service is offered by the *User Datagram Protocol* (UDP).

1.2 Measurements of computer network traffic

We distinguish two categories of measurements of computer network traffic: those taken at the *application level* and those taken at the *packet level*. At the application level the variables of interest are, for example, file sizes, connection durations, transmission durations, the number of files requested from a *file server* or the number of *hits* on a website. At the packet level, information is gathered on the packet stream flowing through a link or cable in the network. All information obtained during the measurement period (ranging from a few hours up to several weeks) comprises the so-called *trace*. For example, a trace can contain the size, timestamp, source host address and destination host address of each packet, as well as whether the packet is controlled by the TCP or UDP algorithm and other application-specific information. To give an idea of the huge quantities of data that are obtained in this way, consider the measurements on the LAN at the Bellcore company in August 1989: in less than 53 minutes already 1.000.000 packets were captured (packet sizes are between 64 and 1518 bytes). Still, with an operating speed of 10 Mbps (Mbps = Megabit per second, i.e. 1.000.000 bits per second) this network is among the slowest ones around today.

From a trace, one can obtain several different data sets. The *local behavior* of the traffic, by which we mean the traffic at small time scales, can be analyzed by considering packet inter-arrival times, idle times and packet sizes. An idea of the *global behavior* of the traffic at the measurement point can be obtained by counting the number of bytes or packets per time unit (e.g. 10 milliseconds or 1 second). The resulting time series is called the *workload*. Notice that ‘global’ refers to aggregation of the traffic over time. In Figure 1.1 a plot of the number of bytes per second passing the measurement point on the Bellcore LAN can be seen. If the source and destination of each packet are recorded in the trace, then one can consider the traffic between separate source-destination pairs. Also, traffic generated by different applications can be distinguished (provided this information is contained in the trace).

For analyzing local behavior it is important that the timestamps are sufficiently accurate. The inter-arrival and silent (or idle) times are computed from differences in the timestamps of consecutive packets. Consider for example the first 1 million packets of the Bellcore measurements (available in the Internet Traffic Archive [49]). Let the timestamp t_i mark the end of packet i , q_i denote the packet size in bytes and s_i the silence time between packets i and $i + 1$. Since 1 byte consists of 8 bits and the network has a speed of 10 Mbps, 1 byte

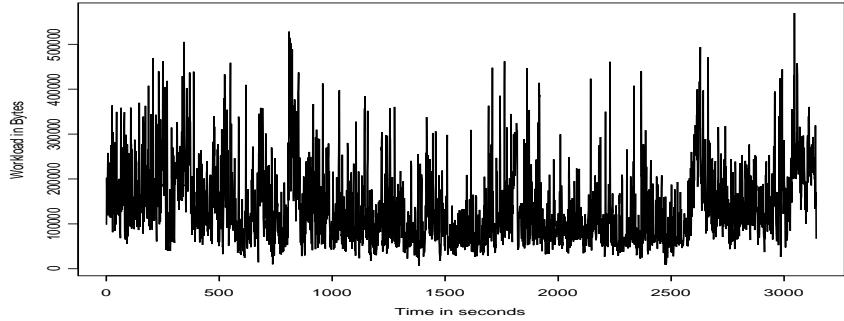


Figure 1.1: Graph of the number of bytes arriving per second on the Bellcore LAN.

corresponds to

$$\frac{8}{10.000.000} = 8 \cdot 10^{-7} \text{ seconds.}$$

The silence times s_i can now be obtained from

$$s_i = t_{i+1} - t_i - 8 q_{i+1} \cdot 10^{-7}. \quad (1.1)$$

The *Ethernet protocol* running on the Bellcore LAN imposes a minimum silence time between two packets of $52.6 \mu\text{s}$ (μs = microsecond, i.e. 10^{-6} seconds). Using (1.1) for calculating silence times from the Bellcore trace, however, results in 12 percent of the s_i being smaller than the minimal $52.6 \mu\text{s}$. Apparently, the accuracy of roughly $10 \mu\text{s}$ of the timestamps is not good enough.

The following schedule gives an overview of the various types of measurements we discussed above.

<ul style="list-style-type: none"> I Application level <ul style="list-style-type: none"> ★ file sizes, connection durations, transmission durations II Packet level (traces) <ul style="list-style-type: none"> 1 local data <ul style="list-style-type: none"> ★ inter-arrival times, silent times, packet sizes 2 global data <ul style="list-style-type: none"> ★ packet counts, byte counts

1.3 Statistical analysis of traffic measurements

Computer network traffic is believed to be characterized by heavy tails, long-range dependence and self-similarity. Here, we briefly introduce these notions and provide a glance at the statistical evidence of their presence in traffic measurements.

1.3.1 Heavy tails

Consider a positive random variable X with distribution function F and right tail

$$\bar{F} = 1 - F.$$

We say that X (or F) has a *heavy tail* if for $x \rightarrow \infty$,

$$P(X > x) = \bar{F}(x) \sim (\text{const}) x^{-\alpha},$$

where \sim means that the ratio of the two sides tends to 1 as $x \rightarrow \infty$ and $\alpha \in (0, 2)$ (a more general definition will be given in Section 3.1.1). If X has a heavy tail then $\text{Var}(X) = \infty$. Moreover, if $\alpha < 1$ then also $E(X) = \infty$. The tails of the distribution of a random variable contain information about the probability of observing values far from the median. Compared to the exponential distribution

$$P(\mathcal{E} > x) = e^{-\lambda x}, \quad \lambda > 0, \quad x > 0,$$

a heavy-tailed distribution has much more variability in the sense that values far from the median are more likely to occur.

In practice, exploratory graphical methods are used to detect the presence of heavy tails. We will discuss some of them in Section 3.1.2; see also Resnick [79]. In computer networks heavy tails have been observed at both the application level and the packet level. Crovella and Bestavros [20, 21] analyze file requests in the WWW and find evidence of heavy tails in the distributions of file lengths, connection durations and idle periods (i.e. times when a workstation is not receiving data). Paxson and Floyd [70] analyze Telnet connections and FTP sessions on WANs. They find that Telnet packet inter-arrival times within a Telnet connection and FTP data burst lengths within an FTP session are heavy-tailed. Finally, Willinger et al. [106] consider the traffic between individual source-destination pairs in the Bellcore LAN. For most source-destination pairs they observe heavy-tailed distributions of the activity periods (i.e. when the source is sending data to the destination) and idle periods.

1.3.2 Long-range dependence

For a weakly stationary, and, hence, finite-variance stochastic process $(X_t, t = 0, \pm 1, \pm 2, \dots)$ dependence between observations at times t and $t + k$ is usually measured by the *autocorrelation function* (ACF) ρ at lag k , i.e.

$$\rho(k) = \frac{\text{Cov}(X_t, X_{t+k})}{\text{Var}(X_t)} = \frac{\text{Cov}(X_0, X_k)}{\text{Var}(X_0)}, \quad k = 0, \pm 1, \pm 2, \dots$$

By plotting $\rho(k)$ against k one can gain an idea of the second order dependence structure of the process. Naturally, one expects that $|\rho(k)|$ decreases as k increases. The notion of *long-range dependence* (LRD) refers to the relative size of $\rho(k)$ at large lags k . If ρ is non-negligible at large lags, we have dependence over a *long range* of sequential observations. Here, we say that a stationary process has LRD if for $k \rightarrow \infty$,

$$\rho(k) \sim c_\rho k^{-\beta}, \quad (1.2)$$

where c_ρ is a positive constant and $\beta \in (0, 1)$. Other definitions will be discussed in Section 3.2.1; see also Beran [4]. In contrast to *short-range dependent* processes (e.g. ARMA) $\rho(k)$ does not decay at an exponential rate, but as a power function. In Figure 1.2 the sample ACF of the number of packets arriving per second on the Bellcore LAN is depicted. Clearly, the graph lies outside of the 95 percent confidence bands around zero for a large number of lags. For comparison we simulated a Poisson arrival process with the same mean inter-arrival time and plotted the sample ACF of the number of arrivals per second. In this case, the autocorrelations can hardly be distinguished from zero.

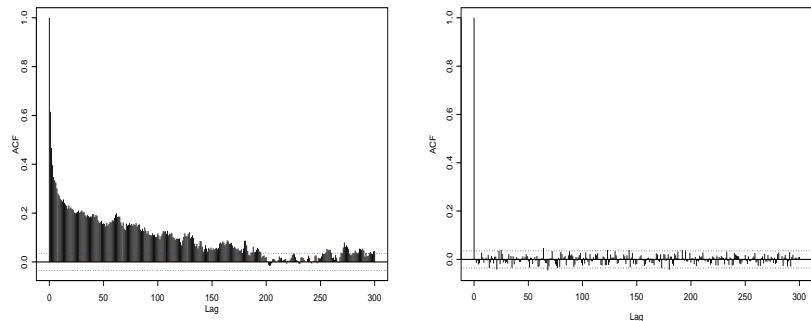


Figure 1.2: Sample autocorrelations of the number of packets arriving per second on the Bellcore LAN (left) compared to those of the number of arrivals in a realization of a mean-matched Poisson arrival process.

Several heuristic graphical tools are used to detect LRD in a time series. We will discuss some of them in Section 3.2.3 (see also Beran [4]). A statistical evaluation of these exploratory methods can be found in Taqqu and Teverovsky [99, 100]. For a wavelet based method, see Abry and Veitch [1]. In computer networks, LRD has been detected in various time series obtained from measurement traces. Abry and Veitch [1] find evidence of LRD in the sequences of inter-arrival times, silence times and packet sizes of the Bellcore measurements. Cao et al. [16] analyze several traces of Internet packets and conclude that LRD is present in the inter-arrival times and packet sizes. Also, LRD has been observed in packet and byte counts in the Bellcore LAN (Leland et al. [60], Abry

and Veitch [1]) and in various WANs (Paxson and Floyd [70]). Finally, Beran et al. [5] find LRD in byte counts obtained from traces of video traffic.

1.3.3 Non-stationarity or long-range dependence?

It is clear that the concept of LRD, as given by (1.2), only applies to weakly stationary stochastic processes. However, no general test for the stationarity of an observed time series is available. Also, the graphical methods that are often used to detect LRD in a time series are not very reliable. For example, it has been observed in the literature that non-stationarities like shifts in the mean or a slowly decaying trend can also be the cause of such slowly decaying autocorrelations as in Figure 1.2 (left).

When determining whether it is reasonable to consider a measured workload series as stationary or not, the level of aggregation, i.e. the time unit used to define the workload, and the length of the series of measurements play an important role. At different time scales different factors may induce non-stationarities in the measurement series. These factors include yearly, monthly, weekly and day-of-week effects and the diurnal cycle, describing fluctuations during 24 hours. Usually a portion of the measurements containing not more than one hour of data traffic is considered and the workload is defined per second. The implicit assumption is that “traffic is stationary over short periods”. Obviously, all effects mentioned above are eliminated in this way. However, also within one hour stationarity is not guaranteed. Apart from system crashes or other ‘shocks’, a changing number of active connections may be a reason to get worried. Since, in the latter case, the traffic is a superposition of a varying number of data streams, this may result in a shifting mean of the process and, hence, give rise to slowly decaying autocorrelations. A way around this would be to consider even shorter time intervals like 5 minutes and to make sure that the number of active connections is fairly stable during this period. But then one might argue that within each connection data is sent in bursts rather than at a fixed rate and, hence, the traffic still behaves in a non-stationary way.

Although the issue of non-stationarity is often settled with only a few words, it is worth a more detailed consideration. This will feature in Sections 3.2.5–3.2.6 and 4.2–4.3.

1.3.4 Self-similarity

A stochastic process $(X_t, t \geq 0)$ is said to be *self-similar* if the finite-dimensional distributions of (X_{at}) and $(a^H X_t)$ are identical for any $a > 0$ and some $H \in (0, 1)$, i.e. if

$$(X_{at}, t \geq 0) \stackrel{d}{=} (a^H X_t, t \geq 0).$$

Hence, the distribution of a self-similar process is invariant under a particular scaling of time and space. The parameter H is called the *index of self-similarity*. The concept of self-similarity became popular due to the work of Mandelbrot

and his co-workers, see e.g. [62, 64]. A thorough mathematical description of self-similarity is given in Samorodnitsky and Taqqu [89].

Workload measurements in computer networks (i.e. packet or byte counts) show a high level of variability on every time scale that is considered, from milliseconds to minutes. For the Bellcore measurements this conclusion is drawn by Leland and Wilson [59] and Fowler and Leland [34]. In Leland et al. [60] and Willinger et al. [106] the variability of the workload on the Bellcore LAN is shown to be roughly the same on five different time scales. This invariance under scaling in time and space is taken to be evidence of self-similarity in the workload measurements. In Figure 1.3 we show the number of packet arrivals in the Bellcore LAN per 10 seconds and per 0.1 second. As we see, the amount of variability does not decrease much at the aggregated level. This is in contrast to the number of arrivals governed by a simulated Poisson process. At the 10-second level the number of arrivals is approximately equal to its mean rate.

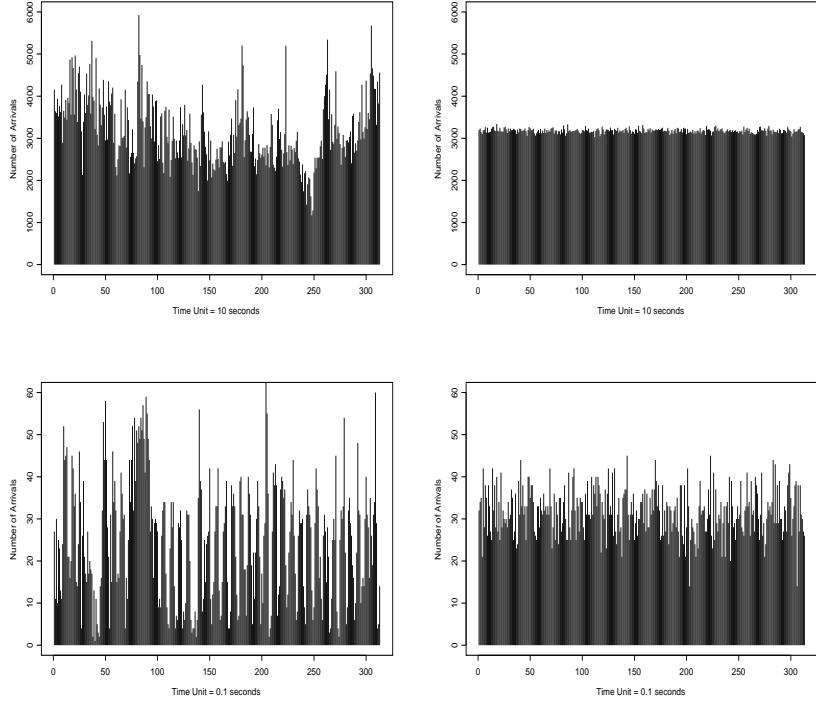


Figure 1.3: The number of (packet) arrivals per time unit on the Bellcore LAN (left) and in a realization of a mean-matched Poisson process (right).

1.3.5 Multifractals

It has been observed that the scaling properties of the workload in WANs at small time scales (roughly below 100 milliseconds) differ from those at larger time scales. While in the latter case the scaling can be described by self-similarity, in the former one it is more complex. The scaling parameter H cannot be taken constant throughout the whole measurement period. This conclusion is drawn by Feldmann et al. [30, 31], Riedi and Lévy-Véhel [84] and Gilbert et al. [36]. It has been suggested that the workload at small time scales can be described by the class of so-called *conservative cascades*. This class consists of multiplicatively generated *multifractal processes* that can incorporate more complex scaling behavior than self-similar (or *monofractal*) processes. Analysis of conservative cascades using wavelets is considered in Riedi [85] and Resnick et al. [81].

It is believed that the multifractal scaling property of the workload at small time scales is caused by the multiplicative nature of the TCP algorithm operating on the WANs under consideration. An effort to give a mathematical description of this multiplicative structure can be found in Resnick and Samorodnitsky [83]. At larger time scales the multiplicative effects are dominated by the additive properties of network traffic (i.e. aggregating the workload generated by different sources or connections) and self-similar scaling behavior is observed. For an overview of the scaling phenomena of the workload, both at small and large time scales, see Riedi and Willinger [86].

In this thesis, however, the multifractal scaling properties of the workload are beyond our consideration. This section is included for completeness only.

1.3.6 Computer networks versus the telephone system

The statistical properties of computer network traffic differ significantly from those of voice traffic in the telephone system (see e.g. Fowler and Leland [34] or Willinger and Paxson [107]). Telephone calls can be modeled by a Poisson process, i.e. their inter-arrival times are roughly exponentially distributed. The lengths of telephone calls have an exponentially bounded right tail. This implies that the autocorrelations of the network workload decrease exponentially in the time between observations. Moreover, on a sufficiently large time scale the workload smooths out, i.e. the number of call arrivals is approximately equal to the long-term arrival rate of the Poisson process.

As we have seen above, these properties are not observed in computer network traffic. On the contrary, file lengths, transmission durations and connection lengths are heavy-tailed, workload processes exhibit LRD and show ‘burstiness’ across an extremely wide range of time scales (i.e. traffic does not smooth out). Therefore, computer networks are a far greater challenge to an engineer than the telephone system.

1.4 Modeling the workload in computer networks

1.4.1 Intuitive explanations

At the application level, file sizes, connection lengths and transmission durations are found to be heavy-tailed. Given constant bitrate transmission the heavy-tailed connection lengths and transmission durations can be explained by the heavy-tailed file sizes. Also, the heavy-tailed file sizes are related to the LRD observed in packet inter-arrival times, silent times and packet sizes. For example, suppose a source is transmitting an extremely large file to a destination host. Due to the observed heavy-tails the probability of extremely large files is non-negligible. Before transmission, the file is decomposed into small packets, on which the bandwidth of the physical medium imposes a certain maximum packet size. Since the file is extremely long it is most efficient if it is decomposed into packets of this maximum size. Hence, a long stream of packets of the same size occurs. Moreover, if there is no interference from other transmissions, the inter-arrival times and silence times between the packets will also be the same. This explains how the transmission of extremely large files causes dependence over a long range of observations (i.e. LRD) in the sequences of packet inter-arrival times, silent times and packet sizes.

The reasoning above is not necessarily valid when the traffic between several independent source-destination pairs is considered. This procedure is known as *multiplexing*. Cao et al. [16, 17] argue that multiplexing weakens the dependence in the packet sequences, since data bursts of several sources interfere and disturb the regular pattern observed when transmitting a single extremely large file. Analyzing several traces from LANs and ATM networks, [16, 17] observe that the sequences of packet sizes inter-arrival times change from LRD to independent if the average number of connections (i.e. data streams) increases.

The regular patterns observed when transmitting extremely long files can also explain the LRD in sequences of packet and byte counts. Cao et al. [16, 17] find that at this aggregated level the effect of multiplexing on the dependence structure of the workload is less pronounced.

1.4.2 Mathematical explanations

The idea of heavy tails as the cause of LRD in workload measurements has been captured in two popular models. The first one is the *ON/OFF model* proposed by Willinger et al. [106]. Here, traffic is generated by M independent and identically distributed ON/OFF sources. If a source is ON it transmits data at unit rate (e.g. 1 byte per time unit). If it is OFF it remains silent. In this way, an individual ON/OFF source generates a binary ON/OFF process W_t , where $W_t = 1$ if at time t the source is ON and $W_t = 0$ if the source is OFF at time t . The lengths of periods in which the source is ON, the *ON-periods* X_i , are independently drawn from a heavy-tailed distribution. Analogously, the *OFF-periods* Y_i are also heavy-tailed. The X - and Y -sequences are assumed independent. It has been shown by Heath et al. [45] that the stationary ver-

sion of the ON/OFF process W_t exhibits LRD. Moreover, since the M sources are independent, the sum of their ON/OFF processes, i.e. the total workload generated by the M sources, also exhibits LRD.

In the second model, the *infinite source Poisson model*, the number of sources in the network is taken infinite. Traffic is generated by independent connections arriving according to a Poisson process, i.e. with exponential inter-arrival times. During a connection, traffic is generated at unit rate. The lengths of the connections are independent and taken from a heavy-tailed distribution. Also, the connection lengths are independent from the connection inter-arrival times. Cox [19] shows that the workload process generated by this model exhibits LRD.

So far we have not offered an explanation for the supposed self-similarity in network traffic. In probability theory a self-similar process can be understood as a weak limit of a sequence of scaled and time-dilated processes (see Lamperti [58]). Such limit theorems have been proved for the cumulative workload processes in the ON/OFF- and infinite source Poisson model and will be dealt with in Chapter 5.

1.5 Outline of the thesis

- Chapter 2 follows Tanenbaum [96] and is an introduction into the terminology, hardware and software of computer networks. In Section 2.1 we discuss network hardware and focus on the differences between LANs, WANs and Inter Networks. This is followed by a closer look at network software in Section 2.2. To facilitate network design the software is organized as a series of *layers* with different functions. The purpose of each layer is to offer a service to the one above it. We describe the process of data transmission as a journey through the layers of the network. In Section 2.3 we focus on the software in the Internet, in particular on the *Internet Protocol* (IP) and the TCP algorithm. Finally, Section 2.4 introduces a relatively new class of networks, the ATM networks.
- In Chapter 3 we discuss in detail the concepts of heavy tails, long-range dependence and self-similarity. In Section 3.1 we describe several classes of heavy-tailed distributions and discuss two graphical methods to detect heavy tails in a random sample: the log-log *complementary distribution plot* and the *Hill plot*. Section 3.2 provides various definitions of LRD and proves their equivalence. Four exploratory tools are discussed to detect LRD: the log-log *correlogram*, the log-log *periodogram*, the log-log *variance-time plot* and the *R/S* method. Also, we apply them to the Bellcore measurements. Next, we address the issue of non-stationarity as the cause of LRD-like phenomena. We show that applying the four exploratory methods above to a realization from a non-stationary ARIMA model is likely to yield the conclusion that LRD is present in the data. In Section 3.3 we define self-similarity and also discuss the concept of *second-order self-similarity* coined by Cox [19]. A common method to ob-

serve self-similarity is to plot the data on a wide range of time scales. If the relative variability remains roughly the same evidence is said to be found for self-similar scaling behavior. We show that this conclusion can also be drawn when ‘zooming in’ on a realization from an ARIMA model.

- Chapter 4 is devoted to the statistical analysis of computer network traffic. In Section 4.1 gives an overview of the relevant measurement studies in the literature. Section 4.2 is based on Stegeman [95]. Here, we provide our own analysis of sequences of workload measurements, including the Bellcore measurements, several traces WAN traffic available in the Internet Traffic Archive [49] and measurements on an ATM network. We show that most workload series can be modeled by a non-stationary ARIMA($p,1,q$) model, with small values of p and q . Hence, it is virtually impossible to distinguish between LRD and non-stationarity. Section 4.3 takes a closer look at the sequence of packet sizes in the Bellcore measurements of August 1989. It appears that several roughly uncorrelated groups of packet sizes can be distinguished, each group being generated by a different process or application running on the Bellcore LAN at the time of the measurements. Examining the arrival processes of packets of fixed sizes leads to the conclusion that they are not consistent with the assumption of stationarity.
- In Chapter 5 we consider the modeling of computer network traffic. In Section 5.1 we provide a detailed definition of the ON/OFF- and infinite source Poisson model, as well as a discussion of their relation to actual networks. In Section 5.2 we present limit theorems for the cumulative workload process in the ON/OFF model, with two parameters converging to infinity: the number of sources M and the time-dilation parameter T . Depending on the order in which the limits are taken, the limit process is either stable Lévy motion or fractional Brownian motion (see Willinger et al. [106] and Taqqu et al. [101]). Both of these processes are self-similar, but the former has independent increments while the latter has an LRD increment process. Next, we present a result in which simultaneous limit regimes of M and T are considered. The relative growth rate of M with respect to T determines which process is obtained in the limit. If M grows ‘fast’ the limit process is fractional Brownian motion, while under ‘slow growth’ convergence to stable Lévy motion can be shown. The proof of this theorem is presented in the Appendix of Chapter 5 and can also be found in Mikosch and Stegeman [65]. A preliminary analysis was done in Stegeman [92]. See Stegeman [93] for an overview of convergence results for the cumulative workload to self-similar limits. In Mikosch et al. [66] the special case when the ON-periods have a ‘heavier tail’ than the OFF-periods is considered. The paper [66] combines convergence results for the cumulative workload in both the ON/OFF- and infinite source Poisson models and grew out of two different projects, each focusing on one of the two models. Convergence results for the cumulative workload in the infinite source Poisson model are also included in Section 5.2. Finally,

in Section 5.3 we discuss some queuing results and simulation studies to illustrate the effect of the characteristics of computer network traffic on network performance.

- Chapter 6 is based on Stegeman [94]. We use the framework of the ON/OFF model to study the number of ON-periods up to time T exceeding a high threshold. Again, we consider simultaneous limits of M and T . Moreover, also the threshold depends on T . We distinguish between the ‘slow’ and ‘fast’ growth conditions on M . Although different approaches are needed, in both cases we are able to show that the number of exceedances converges to a Poisson random variable if the threshold satisfies a balancing condition (guaranteeing a constant number of exceedances in the limit). The ‘slow growth’ case is dealt with in Section 6.1, while Section 6.2 considers the ‘fast growth’ situation. In Section 6.3 we show that the number of exceedances satisfies the Central Limit Theorem. Here, we do not need to impose a condition on the growth rate of M with respect to T . The only requirement is that the number of exceedances is increasing as $T \rightarrow \infty$. The proofs of the above results are presented in the Appendix of Chapter 6.

1.5.1 The Nature of the Beast

With its wild behavior the workload in computer networks does not obey the standard assumptions of the Poisson process. It seems that a new class of models incorporating heavy tails and exhibiting LRD is needed to capture its nature. However, also here problems arise. For example, the issue of LRD versus non-stationarity, the influence of multiplexing on traffic characteristics and the question as to whether the scaling of network traffic is multifractal or self-similar. With respect to the last issue, it has been observed that multifractal scaling behavior occurs at small time scales while at larger time scales traffic seems self-similar. However, a ‘physical’ model that incorporates both these effects still has to be found.

Excursion: William Blake

The theme of reason trying to capture a wildly fluctuating and irregular phenomenon can also be found in the work of the English poet, thinker and artist William Blake (1757-1827). In Blake’s poetic theory several restraining forces are identified which function as the circumference of the natural state of being or reality. Being as such, they degrade the latter in quality to such a degree that it is only a shadow of its true potential. According to Blake the restraining forces must be constantly battled in order to avoid the horrid state of repression. This war takes place in the physical and spiritual world as well as at the core of the human psyche.

A few important restrainers in Blake’s work are the following. First, there is sensory perception. According to Blake the world perceived by the five senses

is regarded as pathetically limited. To overcome these limitations of sensory perception man must use his power of creative imagination. Only those who are artistically awakened, i.e. those who do not accept reality as given, are able to obtain a vision of the world of delight that lies beyond the doors of perception. A second restrainer is dogmatic rule or tyranny. Here, the repressed entity is called human desire and a revolution to overthrow the repressive regime can only be started if human desire is strong enough. The most important restrainer in Blake's work is *Reason*. In this case we may think of a rational model trying to explain nature. But Reason can also act internally, e.g. the repression of human desire by internalized rules. In general, Blake states that Reason functions as a restrainer of what lies at the root of reality: *Energy*. Since Blake strongly advocates the powers of creative imagination, desire and Energy, his view is directly opposed to the scientific rationalism and mechanistic empiricism of his contemporaries Newton, Locke and Rousseau.

In what is generally considered as his most inspired and original work *The Marriage of Heaven and Hell*, Blake states (see [10], Plate 3):

Without Contraries is no Progression. Attraction and Repulsion, Reason and Energy, Love and Hate, are necessary to Human existence.

From these contraries spring what the religious call Good & Evil. Good is the passive that obeys Reason. Evil is the active springing from Energy.

Good is Heaven. Evil is Hell.

This shows that Blake is not concerned with the static and inert state of victory over the restraining forces, but rather with the vitality of the struggle between Reason and Energy. In this sense, he opposes religion which focuses on the victory of Good over Evil. Moreover, Blake regards the concepts of Good & Evil merely as stale abstractions from the dynamic process of the interacting Contraries.

Back to computer network traffic

Using the terminology of William Blake, the efforts of trying to find a model incorporating all features of computer network traffic can be seen as a battle between Reason (mathematical modeling) and Energy (the highly bursty nature of traffic measurements). The fact that no suitable model has been found yet has two implications. First, Reason and Energy continue their strife and Progression is the result. Second, computer network traffic does not (yet) obey Reason and, hence, belongs to the realm of Evil. In the Book of Revelations 13:11-18 Evil is represented by the Beast. Hence, metaphorically, this thesis is an investigation into *The Nature of the Beast*.

Two Computer Networks

Nowadays computers play a predominant role in the gathering, processing and distribution of information. By linking computers together, i.e. building *computer networks*, information can travel long geographical distances at low costs. With the arrival of the Internet, information from all over the world is available at the touch of a button.

By a ‘computer network’ we mean an *interconnected* collection of *autonomous* computers. Two computers are interconnected if they are able to exchange information. With the autonomy requirement we exclude master/slave relationships, such as a large computer with remote terminals.

Computer networks have many benefits. For companies Tanenbaum [96] mentions the following points. The first issue is resource sharing: data is made available without regard to the physical location of the resource and the user. Also, by putting all files on two or three machines in the network, a high reliability of information access is obtained. These machines are called *file servers*. In this network, the users are called *clients*. They request the servers for information or to do some other job. The server performs these tasks and replies. An advantage of this so-called *client-server model* is that it is cheaper than having a large mainframe computer with remote terminals, since small computers have a better price/performance ratio than large ones. Finally, computer networks are a powerful communication medium among widely separated employees.

With the coming of the World Wide Web as one of the applications on the Internet, private individuals are also able to benefit from computer networks. They have access to remote information and are able to engage in person-to-person communication through *e-mail* and in *chatrooms*. There are worldwide newsgroups where people can exchange messages with like-minded individuals. Also, interactive entertainment is available.

Next we focus on the technical issues in computer network design. In Sections 2.1 and 2.2 we discuss network hardware and network software, respectively. We specifically deal with the Internet in Section 2.3. Finally, in Section 2.4 we consider ATM networks. For this chapter we closely follow Tanenbaum [96] both in wording and presentation.

2.1 Network hardware

Computer networks are often classified by their *transmission technology* and *scale*. There are two types of transmission technology. In *broadcast networks* there is a single communication channel that is shared by all the machines on the network. A message, called a *packet*, sent by any machine is received by all the others. An address field within the packet specifies the destination. When a machine receives a packet, it checks the address field. The packet is ignored if it is intended for another machine. Otherwise, it is processed. In broadcasting networks it is often possible to send a packet to all machines or to a specified subset of machines.

A different transmission technology is used in *point-to-point networks*. There are many pairwise connections between individual machines. If a packet is sent from a source machine it may have to visit one or more intermediate machines before arriving at the destination machine. Often multiple routes, of different lengths, exist between the source and destination machines, so routing algorithms play an important role in point-to-point networks.

An alternative criterion for classifying computer networks is their scale. We will distinguish three categories: *Local Area Networks*, *Wide Area Networks* and *Inter Networks*, which are discussed in Sections 2.1.1–2.1.3.

2.1.1 Local Area Networks (LANs)

LANs are privately-owned networks within a single building or campus of up to a few kilometers in size. Since LANs are restricted in size, the worst-case transmission time of a packet is known in advance. This simplifies network management. LANs are often broadcast networks, where the computers are attached to a single communication channel. There are two well-known topologies for broadcast LANs: the *bus* and the *ring*. They are depicted in Figure 2.1.

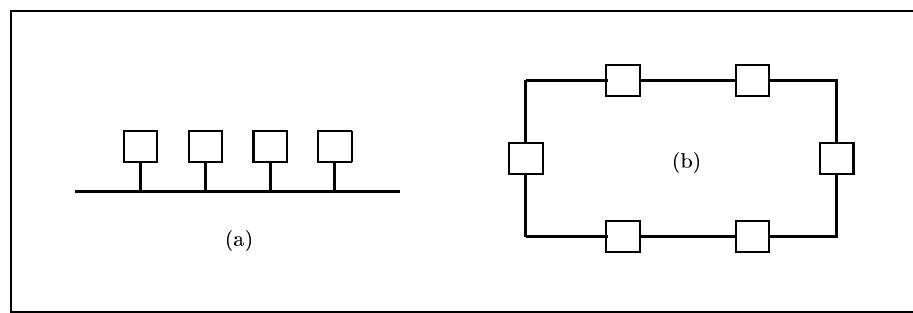


Figure 2.1: Two topologies for broadcast networks: (a) Bus and (b) Ring.

In a bus network the computers are connected through a linear cable. Since it is a broadcast network only one machine is allowed to send a message at any given time. If two machines start transmitting simultaneously, a collision occurs

and an arbitration mechanism is needed. In the popular *Ethernet* LAN, which is a broadcast bus network, this is dealt with as follows. If a machine wants to transmit, it listens to the cable. If the cable is busy, the machine waits until the cable is idle; otherwise it transmits immediately. If two or more machines begin transmitting simultaneously and a collision occurs, they terminate their transmission, wait a random time, and try again.

In a ring network the computers are situated in a circle, each connected to two neighbors. As in the bus network an arbitration mechanism is used to prevent collisions. A popular one is the *Token Ring*. In this network a special packet, called the *token*, circulates around the ring whenever the communication cable is idle. If a machine wants to transmit, it waits for the token to come, removes it from the ring, transmits a packet and finally puts the token back on the ring. Without the token a machine is not allowed to transmit. In this way no collisions can occur. Besides the token ring, also token bus networks exist.

2.1.2 Wide Area Networks (WANs)

A WAN spans a large geographical area, like a country or a continent. It contains two kinds of machines: *hosts* and *routers*. Hosts are intended for running user programmes (i.e. applications), while routers are used as intermediate machines for packets sent by a source host to a destination host. The routers are connected by communication channels, which together comprise the *subnet*. The subnet is usually a point-to-point network. Figure 2.2 contains an example of a WAN.

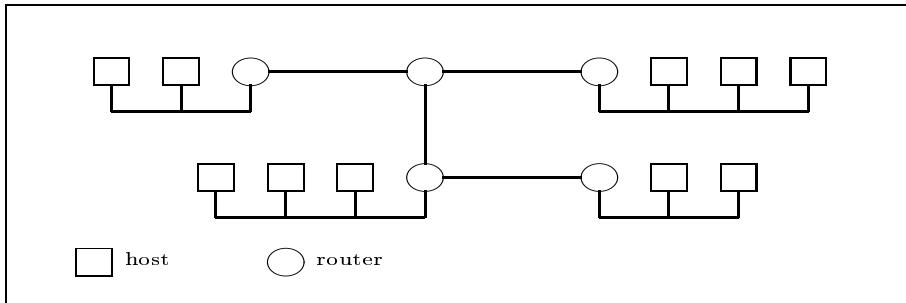


Figure 2.2: Hosts and routers in a WAN.

The strict separation of the communication aspects (i.e. packets traveling between routers in the subnet) from the application aspects (i.e. the hosts) simplifies network design. If host *A* sends a packet to host *B*, the packet is first sent to the closest router, which stores it until the required output line is free, and then forwards it. This principle is called *store-and-forward* or *packet-switching*. It is possible that the packet travels via one or more intermediate routers in the subnet. A routing algorithm is used to determine the route a packet has to travel through the subnet.

2.1.3 Inter Networks

An Inter Network is a collection of interconnected networks, e.g. a number of LANs connected by a WAN. Often the networks to be connected are incompatible in terms of their hardware or software. In these cases, machines called *gateways* are used to make the connection and provide the necessary translation, both in hardware and software.

The best known example of an Inter Network is the *Internet*, which spans the whole planet. The Internet is the successor of the ARPANET, which is basically the grandparent of all computer networks. The ARPANET was a research network sponsored by the U.S. Department of Defense (DoD). It was developed in the late 1960s by network researchers hired by the Advanced Research Projects Agency (which abbreviates as ARPA). During the Cold War the DoD wanted to have a computer network that would still work if some hosts, routers or communication lines would fail to function. This was one of the reasons for designing the ARPANET. The ARPANET was a packet-switched network, consisting of a subnet and host computers. It connected hundreds of universities and government installations. In the 1980s the NSFNET, the network of the National Science Foundation, was connected to the ARPANET. The resulting network became known as the Internet and many existing networks were connected to it. Right now the Internet literally spans the whole planet and has several hundreds of millions of users.

Traditionally, the Internet had only four applications: *e-mail* (sending and receiving electronic mail), *newsgroups* (users with a common interest exchanging messages), *remote login* (using Telnet or Rlogin users are able to log into a remote machine on which they have an account) and *file transfer* (with the FTP program files can be copied from one machine to another through the Internet). Up until the early 1990s, the Internet was largely used by academic, government and industrial researchers. When the *World Wide Web* (WWW) was introduced as a new application, however, also millions of non-academic users emerged on the Internet. The WWW also attracted the attention of companies who engaged in *e-commerce*.

2.2 Network software

In order to make computer networks easier to design, they are organized as a series of *layers*, each one built upon the one below it. The number of layers and the name, contents and function of each layer differ from network to network. However, in all networks the purpose of each layer is to offer certain services to the layer above it. In this way, the higher layers are shielded from the details of how the offered services are actually implemented.

In a computer network, layer n on one machine carries on a conversation with layer n on another machine. The conversation obeys the layer n *protocol*. The protocol is a set of rules and conventions used between the communicating parties. In reality, no data are directly transferred from layer n on one machine

to layer n on another machine. Instead, each layer passes data and control information to the layer below it, until the lowest layer in the network is reached. Below layer 1 is the *physical medium* through which the data is actually transferred to another machine. Between layer n and layer $n+1$ there is an *interface*, which defines the operations and services layer n offers to layer $n+1$. The set of layers and protocols is called the *network architecture*. As an example of a network architecture, we will discuss the reference model in Figure 2.3. Here virtual communication lines are dotted, while physical communication lines are solid. Virtual communication occurs between layers on different machines, using a protocol. Physical communication is carried out between adjacent layers on the same machine, using the interface, and through the physical medium.

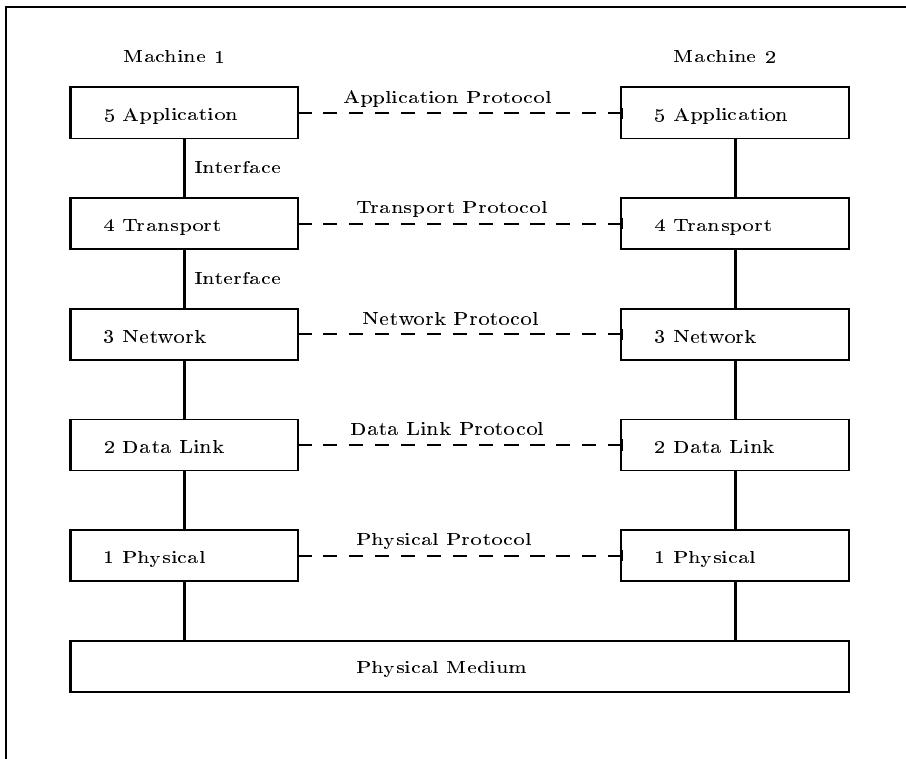


Figure 2.3: A reference model of network architecture.

Suppose a user on Machine 1 wants to send data to Machine 2. The data is first given to the *application layer*. Here, some control information, contained in the so-called *application header*, is put in front of the data. The application header makes sure that the application protocol is satisfied. For example, it specifies whether the data is an ASCII-file or a postscript-file, if some form of data-compression is used, if the data is an e-mail or what file-naming convention

is used.

The data and the application header are passed on to the *transport layer*. Here they are split up into smaller units (if needed). Each unit is provided with a transport header, which contains information on the size and order of the units. The transport layer also establishes the connection between Machine 1 and Machine 2 and decides whether one or multiple network connections are used. Finally, it is made sure that Machine 1 will not send the data faster than Machine 2 can accept them. This mechanism is called *flow control*.

The data then arrive at the *network layer*. Here it is determined how the packets are routed through the network from Machine 1 to Machine 2. Remember that in a WAN or an Inter Network there are often multiple routes via intermediate machines called routers. If too many packets are present in the subnet at the same time, they will get in each other's way and form bottlenecks. Another task of the network layer is to prevent this congestion; this is called *congestion control*. The routing information and the addresses of the source and destination host (Machines 1 and 2, respectively) are added to the data in the form of a network header.

The layers discussed so far are all concerned with *end-to-end transmission*, i.e. the source and destination machines are taken into account explicitly. The following two layers are only concerned with sending and receiving data through the physical medium. To them the path from the source machine to the destination machine is not visible.

In the *data link layer* the data arriving from the network layer are broken up into *data frames* of a few hundred or thousand bytes. Each data frame has a frame header, containing information on its size and reference number. The data frames are transmitted sequentially. Since no communication channel is completely error-free, the data received may deviate from the data sent. Therefore, the data link layer also has the jobs of *error-detection* and *error-correction*. Usually, the procedure is as follows. A common method is used to compute a so-called *checksum* number for each dataframe. The checksum is added to the frame header. After transmission the checksum is again computed. If it differs from the number in the data link header, a transmission error has occurred. In broadcasting networks the data link layer also takes care of obtaining access to the shared communication channel.

The data frames are passed on to the *physical layer*. Here they are transmitted as a sequence of bits over the physical medium. The physical layer has to take into account the properties of the physical medium, such as how many bits can be transmitted per second (the so-called *bitrate* or *bandwidth*; for example, a typical Ethernet LAN operates at 10 Mbps, which is 10.000.000 bits per second), the time it takes for 1 bit to travel from Machine 1 to Machine 2 (this is called the *delay* of the connection) and whether data can be sent in both directions simultaneously (a *full-duplex* connection) or only sequentially (a *half-duplex* connection).

The bits arrive through the physical medium (usually copper wire or fiber optics, but also wireless networks exist) at the physical layer on Machine 2. The data link layer then recognizes the data frame boundaries in the incoming

stream of bits and checks the data frames for transmission errors. For each correct frame, it sends back an *acknowledgment frame*. If, some time after transmitting a frame, the data link layer on Machine 1 has not received an acknowledgment frame, it sends the frame again. If all frames have arrived on Machine 2 without errors, the frame headers are removed, the frames are put together into a packet and passed on to the network layer. Here the incoming packets are put into the right order (it might happen that the packets do not arrive in the order they were sent due to using different routes) and the network headers are removed. The data is passed on to the transport layer, where it is checked if the right connection is used. The transport header is removed and the data arrive at the application layer. Here the application header is removed and the data are put into a file of the desired format, which can then be accessed by the user on Machine 2.

2.3 Protocols in the Internet

Here we will discuss the protocols that are used in the network and transport layers in the Internet. First we explain the difference between *connection-oriented* and *connectionless* networks and services.

2.3.1 Connection-oriented versus connectionless

Consider a point-to-point network containing hosts and routers, the latter connected by a subnet. We say that the subnet is *connection-oriented* if it is modeled after the telephone system. Before data is sent, a connection, called a *virtual circuit*, is set up and the sender and receiver negotiate about the rate at which data can be sent or the required buffer space. This process is called *option negotiation*. If these parameters are set, data can be sent in both directions, following the same route through the virtual circuit. In this way, packets arrive in the same order as they were sent. Due to option negotiation, flow control is provided automatically. Also, congestion control is easy, since all packets travel the same route through the subnet.

Alternatively, a subnet can be *connectionless*. In this case, it resembles the postal system. All packets, also called *datagrams*, contain a full source and destination address and are routed independently through the subnet. Hence, they may not arrive in the order they were sent. In such a subnet congestion control is difficult, since each packet may travel a different route. However, it is less vulnerable to router failure. The packets are simply rerouted along another path. In a connection-oriented subnet all virtual circuits passing through a failed router are terminated.

The distinction between connection-oriented and connectionless does not only apply to actual networks but also to services offered by the transport layer in a network. In this sense, the user of a connection-oriented service establishes a connection, uses the connection and then releases the connection. The data arrive in the same order as they were sent. It is important to notice that a

connection-oriented service can also be provided when the physical network is connectionless.

An important aspect of a provided service is *reliability*, i.e. whether data is not lost during transmission. Usually, a reliable service is implemented by having the receiver acknowledge (to the sender) the receipt of each message. It depends on the application whether a reliable or unreliable, a connection-oriented or connectionless service is used. A reliable connection-oriented service is often used for file transfer. In this way all the bits arrive correctly and in the same order as they were sent. However, the acknowledgment process introduces delays. For some applications, like digitized voice or video, speed is more important than maintaining high quality and an unreliable connection-oriented service is more appropriate. Other applications, like e-mail, do not require a connection and can use a connectionless service.

2.3.2 The Internet network layer

The Internet is a connectionless point-to-point Inter Network consisting of hosts and routers. It can be viewed as a collection of so-called *Autonomous Systems* (ASes), which are regional networks. The ASes are connected by a *backbone* of high-bandwidth lines and fast routers. The network layer in the Internet handles the routing of the datagrams. These datagrams are usually around 1500 bytes and obey the *Internet Protocol* (IP). The network layer puts an IP header in front of each datagram. The IP header contains the source and destination IP addresses, a checksum, service options (reliability versus speed) and routing options. Each host has a 32 bit IP address, specifying the network it is in and its address within this network.

Since the Internet is connectionless, each IP datagram may follow a different route to its destination host. There are two different routing protocols: one for routing within an AS (the *interior gateway protocol*) and one for routing between ASes (the *exterior gateway protocol*). We will first describe the interior gateway protocol, which is called OSPF (Open Shortest Path First). OSPF works by abstracting the collection of actual networks, routers and lines into a directed graph in which each arc is assigned a cost (e.g. delay time, distance, number of hops or delay time and queuing time together). It then computes the shortest paths between all pairs of routers using Dijkstra's shortest-path algorithm (see Dijkstra [23]). If there are m nodes in the graph, it can be shown that the computational complexity of Dijkstra's algorithm is $O(m^2)$. These shortest paths must be known by all routers. Upon receiving a datagram the router checks the destination IP address, determines the shortest path from itself to the destination host and forwards the datagram to the next router on this path. Therefore, each router must know the network topology inside the AS and the 'cost' of each line.

The algorithm used to get this information to all routers is called *link state routing*. First, a router sends a special HELLO packet on each line to its neighboring routers. These routers then respond by sending back a reply telling who they are. In this way each router gets to know its neighbors. Next, the router

sends an ECHO packet to its neighbors, containing a timestamp. This packet is immediately returned by the neighbors. By comparing the timestamp with the time at which the ECHO packet returns, the *round-trip time* can be computed. If the round-trip time is divided by two, the delay of the connection is known. In this way, each router determines the ‘cost’ of the lines to its neighbors. This information is put into a so-called *link state packet*. Each router then sends copies of its link state packet to all its neighbors and forwards incoming link state packets on every outgoing line except the one it arrived on. This method of distributing the packets is called *flooding*. Since flooding generates a lot of duplicate packets, some method must be used to damp the process. One way is, for each router, to keep track of the link state packets it has forwarded to avoid sending them out a second time. For other methods see Tanenbaum [96], Section 5.2.3. The process of flooding the link state packets ensures that all routers in the AS know the network topology and the ‘cost’ of each line. Hence, each router can calculate the shortest path to each other router, and determine on which line an incoming IP datagram must be forwarded by simply checking the destination IP address. Link state routing is a *dynamic routing algorithm*. Usually, the procedure of sending HELLO and ECHO packets is repeated at regular time intervals. In this way, congestion or router failure can be taken into account when determining the shortest paths. For more details see Tanenbaum [96], Sections 5.2.6 and 5.5.5.

Next we discuss the exterior gateway protocol, which is used for routing IP datagrams between ASes. This protocol is called BGP (Border Gateway Protocol). While the interior gateway protocol only has to take care of moving packets as efficiently as possible from source to destination, BGP also has to deal with politics. For example, a corporate AS might be unwilling to carry packets originating in a foreign AS and ending in a different foreign AS, even if its own AS was on the shortest path between the two foreign ASes. However, it might be willing to carry transit traffic for ASes that paid it for this service. In practice, each BGP router determines the shortest paths to the other BGP routers and discards any route violating a policy constraint. BGP uses *distance vector routing* as a routing algorithm. Each BGP router maintains the distance to each destination and also keeps track of the exact path used. Periodically, this information is transmitted to each neighboring BGP router, which uses it to update its own paths and distances. Updates are necessary if some line or BGP router fails, for example. As link state routing, also distance vector routing is a dynamic routing algorithm. Details can be found in Tanenbaum [96], Sections 5.2.5 and 5.5.6.

2.3.3 The Internet transport layer

The task of the transport layer in a computer network is to provide reliable, efficient data transport from the source machine to the destination machine, independent of the physical network. Although the Internet itself is a connectionless network, the Internet transport layer offers both (reliable) connection-oriented and (unreliable but fast) connectionless service to the application layer.

The protocol used for reliable connection-oriented service is called *Transmission Control Protocol* (TCP). Unreliable connectionless service is offered by the *User Datagram Protocol* (UDP).

Transport using TCP works as follows. First, a connection is set up between the source and destination hosts. Then the receiver lets the sender know how much bufferspace is available for incoming packets. The sender takes this information into account when sending packets. In this way flow control is incorporated in TCP. Also during the process of sending packets, bufferspace information is still sent to the sending host (this is called *dynamic buffer allocation*). For each undamaged packet the destination host receives, an acknowledgment is sent back to the source host. When the sender has sent all packets and received all acknowledgments, the connection is released.

The communicating parties on the sending and receiving hosts are called TCP *entities*. Communication between the TCP entities is carried out by sending TCP headers, possibly followed by some additional data. A TCP header has a standard format and includes source and destination addresses, a checksum and some options relating to establishing and releasing connections and acknowledgments of received messages. The procedure of establishing a connection is called the *three-way handshake*. First, the sender sends a CONNECTION REQUEST. The receiver sends back a CONNECTION ACCEPT. Finally, the sender sends an acknowledgment of having received the CONNECTION ACCEPT. All TCP connections are full-duplex, which means that data can flow in both directions simultaneously. Each of these directions has to be released independently of the other. When the sender has no more data to send, it sends a CONNECTION CLOSE. This is acknowledged by the receiver and one direction is released. The receiver then also sends a CONNECTION CLOSE to the sender. If this is also acknowledged, the connection between the two hosts is released.

In the Internet, congestion control can be found in both the network layer (if routing algorithms use delay and queuing time as a measure of 'cost' on a line) and the transport layer (TCP, not UDP). TCP congestion control is closely linked with flow control. The data stream passed on by the application layer is broken into *segments* by TCP. These segments (each with a TCP header) are handed to the network layer for transmission. The size of the segments may not exceed the *maximum segment size* associated with the connection. Congestion control and flow control are concerned with the number of bytes sent in one transmission by the source host (possibly consisting of several segments) and the timing of these transmissions. The number of bytes sent is constrained by the bufferspace of the receiver and the internal carrying capacity of the network. During the process of data transmission the receiver keeps the sender informed on the amount of available bufferspace. This results in an upper bound for the number of bytes sent in one transmission. This upper bound is called the *receiver's window*. The upper bound imposed on the number of bytes sent by network capacity is called the *congestion window*. The total number of bytes in one transmission is taken as the minimum of the receiver's window and the

congestion window. If a segment arrives at the TCP entity on the receiving host, an acknowledgment is sent back. If the time between sending a segment and receiving the acknowledgments exceeds the so-called *timeout interval*, the segment is retransmitted. Only when an acknowledgment is received in time, another segment is transmitted.

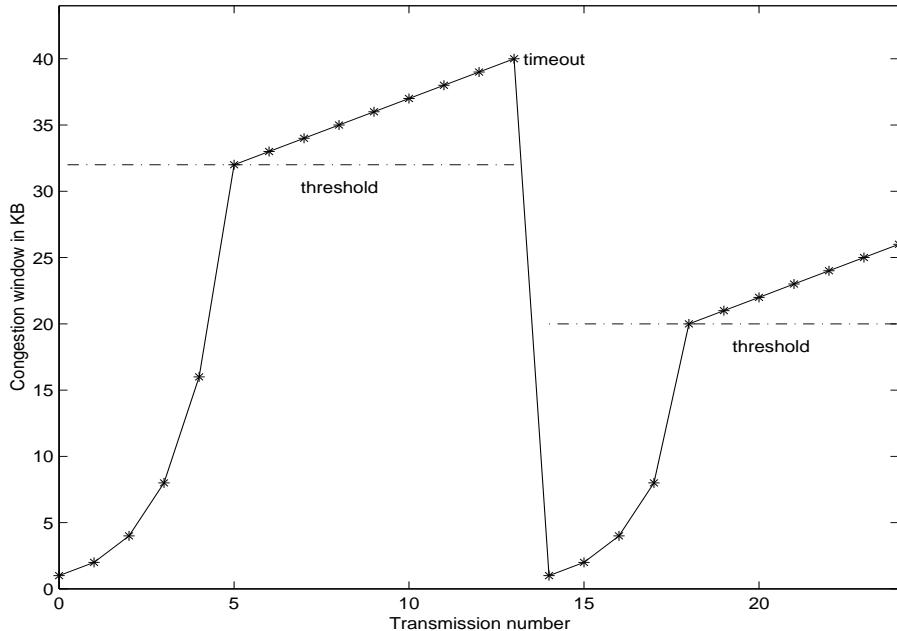


Figure 2.4: An example of the Internet congestion control algorithm.

During the process of transmission the congestion window is increased with each received acknowledgment. This algorithm is called *slow start* (see Jacobson [50]). It uses a *threshold*, which is initially set to 64 KB (1 KB is 1024 bytes). Also, the maximum segment size plays a role. If the congestion window is smaller than the threshold and all acknowledgments of the segments sent in the same transmission are received in time, the congestion window is doubled for the next transmission. If the congestion window is larger than the threshold, it is increased by the maximum segment size. If a timeout occurs (i.e. an acknowledgment takes too long to arrive or a segment is damaged or lost), the threshold is set to half of the current congestion window. The congestion window itself is reset to the maximum segment size.

In Figure 2.4 the congestion control algorithm is illustrated. Here the maximum segment size is 1 KB. The threshold is 64 KB, which is also the initial value of the congestion window. We assume that the receiver's window is large such that the size of the segments sent is equal to the congestion window. Initially, 64 KB is transmitted (i.e. 64 segments of 1 KB in this case), but a timeout

occurs. The threshold is set to 32 KB and the congestion window to 1 KB. As no timeout occurs the congestion window grows exponentially until it hits the threshold (32 KB). Then it grows linearly, increasing with 1 KB for each successful transmission. On transmission 13 a timeout occurs. The threshold is set to half of the current congestion window (which is now 40 KB). The congestion window itself is reset to 1 KB. The slow start procedure is resumed with a lower threshold value.

An implicit assumption in the Internet congestion control algorithm is that a timeout is usually caused by congestion. However, a timeout can also have been caused by noise on the transmission line, resulting in damaged packets. Nowadays, usually fiber optics are used, which have few transmission errors. The number of timeouts depends strongly on the length of the timeout interval. Timer management is of key importance to TCP congestion control. Work on improvements of TCP is still continuing. For example, Brakmo et al. [14] have reported improving TCP throughput by 40 percent to 70 percent by more accurate timer management and predicting congestion before timeouts occur. This variant is called TCP Vegas.

2.4 ATM networks

Over the last few years a new wide area service, called *Broadband Integrated Services Digital Network* (B-ISDN), is available to consumers and companies. It offers video on demand, live television from many sources, Internet connections, a phone combined with full motion video, CD-quality music, LAN interconnection and high-speed data transport. In the future, B-ISDN is supposed to replace the telephone system and combine the mentioned services in a single new network on which every home is connected. The technology used by B-ISDN is called *Asynchronous Transfer Mode* (ATM), because it is not synchronous (tied to a master clock) like most long distance telephone lines. The basic idea behind ATM is to transmit all information in small, fixed-size packets called *cells*. Each cell is 53 bytes long and consists of a header of 5 bytes and a *payload* of 48 bytes. ATM networks are connection-oriented, i.e. making a call requires setting up a connection first. After that, all cells follow the same path from the source to the destination. Cell delivery is not guaranteed but their order is (this is especially important for audio and video streams). ATM networks are organized like WANs, with communication lines and switches (routers). Each line is unidirectional. For full-duplex operation, two parallel links are needed, one for traffic each way. Typical bitrates for ATM networks are 155 and 622 Mbps. The technique of cell-switching in ATM networks (using fiber optics) contrasts the old tradition of circuit-switching (using copper wire) within the telephone system. Cell-switching was chosen since it is more flexible (both constant rate traffic (uncompressed audio and video) and variable rate traffic (data) can be handled easily) and faster than circuit-switching. Also, cell-switching can provide broadcasting (necessary for television distribution) and circuit-switching

cannot.

The layers in the ATM network architecture differ from the standard reference model in Figure 2.3. In Figure 2.5 the ATM reference model is depicted. It consists of three layers, the physical, ATM, and ATM adaption layers. On top of the ATM adaption layer the user can put several additional layers, for example the network and transport layers of the Internet.

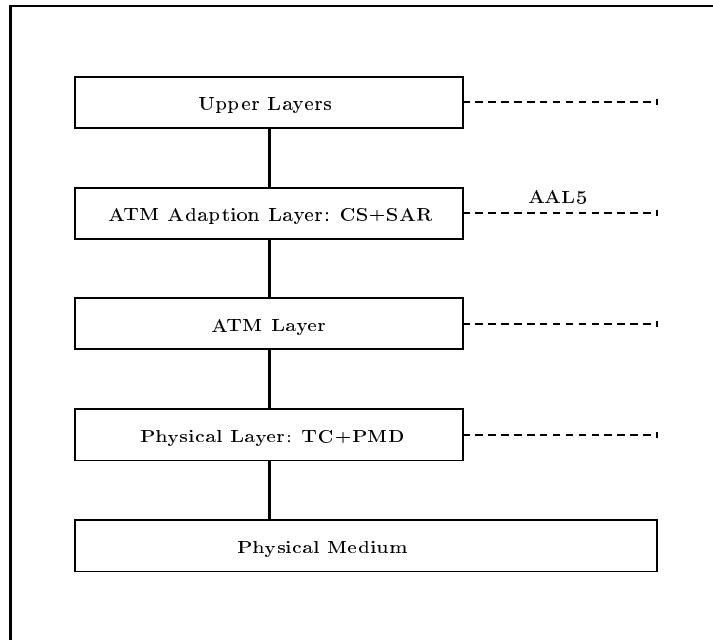


Figure 2.5: The ATM reference model.

The physical layer consists of two sublayers: the *Physical Medium Dependent* (PMD) sublayer and the *Transmission Convergence* (TC) sublayer. The PMD sublayer is analogous to the physical layer in the standard reference model in Figure 2.3. The TC sublayer is comparable to the data link layer in Figure 2.3. One of its tasks is error-detection. Of each cell the checksum of the header is computed to reduce incorrect delivery (error-detection in the payload takes more time and is left to higher layers). ATM is designed to be independent of the transmission medium. If an asynchronous system is used cells can be transmitted at any time with arbitrary idle periods. However, if the transmission is synchronous, it expects a cell at fixed intervals. If no cell is available when one is needed the TC sublayer has to send an idle cell. Another important task of the TC sublayer is to match the ATM output rate to the rate of the underlying transmission medium. For example, if frames are used to transmit data the cells must be fitted into these frames when sent and extracted when received. This involves recognizing cell boundaries in an incoming bitstream.

The ATM layer is concerned with moving cells from the source to the destination host. This involves setting up a connection (a virtual circuit), with possibly several intermediate switches. Each virtual circuit has a number which is added to the cell header. By checking the cell header intermediate switches know along which virtual circuit the cell is travelling and, hence, on which line the cell needs to be forwarded. Since fiber optics (for which ATM is designed) are reliable and delays in audio and video streams are unacceptable, no acknowledgement process is used to guarantee cell delivery. It is, however, guaranteed that cells never arrive out of order. Another task of the ATM layer is to enforce the quality of service demanded by the user. Several parameters (e.g. minimum and peak cell rate, cell loss ratio and cell transfer delay) may be specified by the user. Usually, the values depend on the application (e.g. videoconferencing or file transfer) and the service offered (e.g. constant bit rate or variable bit rate, real time or non-real time). The required quality of service is enforced using, for example, a *leaky bucket* algorithm to adjust the cell rate. Another factor influencing the quality of service is congestion at intermediate switches. Congestion control in ATM networks is organized by reserving the resources a virtual circuit needs before it is set up, i.e. the emphasis is on prevention rather than on actual control.

The ATM adaption layer (AAL) was designed to provide useful services to application programs and to shield them from the mechanics of chopping up data into cells before a transmission and the process of reassembling after it. The AAL consists of two sublayers: the Convergence sublayer (CS) and the Segmentation and Reassambly (SAR) sublayer. The CS sublayer puts a message between a header and trailer and hands it to the SAR sublayer. There it is segmented into cells. The AAL can provide serveral kinds of services: real time or non-real time, constant bit rate or variable bit rate and connection-oriented or connectionless. At its beginning, four protocols existed, each suitable for a specific class of services. Now, the most used protocol is AAL5 which is functionally similar to UDP in the Internet. AAL5 offers a choice between reliable (guaranteed delivery with flow control) and unreliable service. It can operate in two modes: stream or message. In message mode, each call from the application to AAL5 injects one message into the network. The message is delivered as such, i.e. message boundaries are preserved. In stream mode the boundaries are not preserved, which is more useful for audio and video streams. Message mode can be used for data transfer or small messages.

Three **Mathematical Concepts**

Empirical studies of traffic in computer networks suggest that three properties are invariantly present in the data: heavy tails, long-range dependence and self-similarity. In this chapter we introduce these mathematical notions and briefly discuss some methods to detect them in real-life data sets. In Section 3.1 we introduce the concept of heavy-tailed distributions, including Pareto and stable distributions. We discuss two common methods for detecting heavy tails, the log-log complementary distribution plot and the Hill estimator. Also, infinite variance stable Lévy motion is introduced. In Section 3.2 we discuss long-range dependence as a property of stationary stochastic processes, and describe four exploratory methods to detect long-range dependence in a real-life time series. We apply these methods to a series of workload measurements on the Ethernet LAN at the Bellcore company in August 1989. Next, we address the issue of non-stationarity versus long-range dependence and show that also a realization from a non-stationary ARIMA model can exhibit features resembling those of the long-range dependence situation. Finally, we introduce fractional Gaussian noise, a stochastic process with long-range dependence. In Section 3.3 we give the definition of a self-similar stochastic process. A method used to detect self-similarity is by plotting the data on different scales. If the relative variability remains roughly the same on the various scales it is believed that evidence has been found of the self-similar nature of the data. We indicate the doubtfulness of this method by showing that the same effect can be obtained with a realization from an ARIMA model, which is not self-similar.

3.1 Heavy tails

3.1.1 Definition and examples

The tails of the distribution of a real-valued random variable contain information about the probability of observing values far from the median. The amount of probability mass contained in the tails determines whether the distribution is considered ‘heavy-’ (also ‘fat-’) or ‘light-tailed’. A typical example of the light-tailed case is the exponential distribution $P(X > x) = e^{-\lambda x}$, $\lambda > 0$, $x > 0$. In

general, we consider a positive random variable X with distribution function F and right tail

$$\bar{F} = 1 - F.$$

We say that X (or F) has a *heavy tail* if for $x > 0$

$$P(X > x) = \bar{F}(x) = x^{-\alpha} L(x), \quad (3.1)$$

where $\alpha \in (0, 2)$ and L is slowly varying (at infinity), i.e. for all $t > 0$,

$$\lim_{x \rightarrow \infty} \frac{L(tx)}{L(x)} = 1.$$

Notice that if X has a heavy tail then $\text{Var}(X) = \infty$. Moreover, if $\alpha < 1$ then also $E(X) = \infty$. For a random variable on the whole real line (3.1) defines a heavy right tail. A heavy left tail can be defined analogously. A positive, measurable function $f(x) = x^{-\alpha} L(x)$, $x > 0$, with $\alpha \geq 0$ and L slowly varying is called *regularly varying (at infinity)*. The parameter $-\alpha$ is called the *index of regular variation*. Hence, (3.1) states that \bar{F} is regularly varying with index $-\alpha$. A slowly varying function is regularly varying with index 0.

In the literature, heavy tails are sometimes defined by the special case of (3.1) when $L(x) \rightarrow c$ as $x \rightarrow \infty$, for some positive constant c . The advantage is that mathematical calculations are simplified. Examples are *Pareto* and *stable* distributions. A possible parametrization of the *Pareto distribution* is given by

$$\bar{F}(x) = \left(\frac{\kappa}{\kappa + x} \right)^\alpha, \quad \alpha, \kappa > 0, \quad x \geq 0.$$

If $\alpha < 2$ the Pareto distribution is heavy-tailed in the sense of (3.1). Stable distributions are characterized by four parameters: the *index of stability* $\alpha \in (0, 2]$, a scale parameter $\sigma_0 > 0$, a skewness parameter $\beta \in [-1, 1]$ and a location parameter $\mu \in \mathbb{R}$. The characteristic function $E(\exp\{itX\})$ of a stable random variable X is given by

$$\begin{aligned} & \exp\{-\sigma_0^\alpha |t|^\alpha (1 - i\beta \text{sign}(t) \tan(\pi\alpha/2)) + i\mu t\} & \text{if } \alpha \neq 1, \\ & \exp\{-\sigma_0 |t| (1 + 2i\beta\pi^{-1}) \text{sign}(t) \log|t|) + i\mu t\} & \text{if } \alpha = 1. \end{aligned}$$

If X has a stable distribution this is denoted by $X \sim S_\alpha(\sigma_0, \beta, \mu)$. Although any stable distribution has a density, in general an explicit expression of the density in terms of elementary functions is unknown. Exceptions (excluding the degenerate case) are the Lévy distribution $S_{1/2}(\sigma_0, 1, \mu)$, the Cauchy distribution $S_1(\sigma_0, 0, \mu)$ and the Gaussian distribution $S_2(\sigma_0, 0, \mu) = N(\mu, 2\sigma_0^2)$ (β is irrelevant when $\alpha = 2$). The parameter β determines the skewness of the distribution to the left ($\beta < 0$) or to the right ($\beta > 0$). If $\beta = 0$ the distribution is symmetric about μ . The mean of a stable random variable is μ when $\alpha \in (1, 2]$; for $\alpha \leq 1$, $E|X| = \infty$. For $\alpha \in (0, 2)$ the variance is infinite. Using a Tauberian

theorem it can be shown (see e.g. Samorodnitsky and Taqqu [89], Property 1.2.15) that for $\alpha \in (0, 2)$, a stable random variable X has tails, as $x \rightarrow \infty$,

$$P(X > x) \sim C_\alpha \frac{1+\beta}{2} \sigma_0^\alpha x^{-\alpha}, \quad \text{and} \quad P(X \leq -x) \sim C_\alpha \frac{1-\beta}{2} \sigma_0^\alpha x^{-\alpha},$$

where the constant C_α is given by

$$C_\alpha = \begin{cases} \frac{1-\alpha}{\Gamma(2-\alpha) \cos(\pi\alpha/2)} & \text{if } \alpha \neq 1, \\ \frac{2}{\pi} & \text{if } \alpha = 1. \end{cases} \quad (3.2)$$

Hence, if $\alpha < 2$ a stable distribution has heavy tails in the sense of (3.1). For more properties of stable distributions and other ways to define them we refer to Samorodnitsky and Taqqu [89], Chapter 1. Some classical references are Gnedenko and Kolmogorov [37], Chapter 7, Feller [33], Chapter IV and Ibragimov and Linnik [48], Chapter 2.

There also exist broader classes of heavy tailed distributions (possibly with finite variance) encompassing the regularly varying class. One example is the class of *subexponential* distributions. A distribution F with support $(0, \infty)$ is called subexponential if for all (some) $n \geq 2$,

$$\lim_{x \rightarrow \infty} \frac{\overline{F^{n*}}(x)}{\overline{F}(x)} = n, \quad (3.3)$$

where F^{n*} denotes the n -fold convolution of F . Alternatively, (3.3) states that for iid positive random variables X_i , $i \geq 1$, with distribution F , for all $n \geq 2$,

$$P\left(\sum_{i=1}^n X_i > x\right) \sim P\left(\max_{1 \leq i \leq n} X_i > x\right), \quad x \rightarrow \infty. \quad (3.4)$$

For regularly varying tails (3.4) can be checked by observing that the class of distributions satisfying (3.1) is closed under convolution (see Embrechts et al. [28], Lemma 1.3.1). Intuitively, (3.4) says that, for extremely large values in a sample, the tail of the sum is approximately equal to the tail of the largest summand. This is in agreement with the idea of heavy tails: values extremely larger than the median are attained with non-negligible probability. For relations between different classes of heavy-tailed distributions we refer to Embrechts et al. [28], Section 1.4. An encyclopedic treatment of regular variation can be found in Bingham et al. [9].

Let Φ be the distribution function of $N(0, 1)$ with right tail $\overline{\Phi} = 1 - \Phi$. As $x \rightarrow \infty$,

$$\Phi(-x) = \overline{\Phi}(x) \sim \frac{1}{\sqrt{2\pi} x} e^{-x^2/2}. \quad (3.5)$$

Hence, the tails do not satisfy (3.1). Other examples of so-called *light-tailed* distributions are the exponential and gamma distributions. A graphical illustration can be found in Figure 3.1, where the tails of an exponential and a

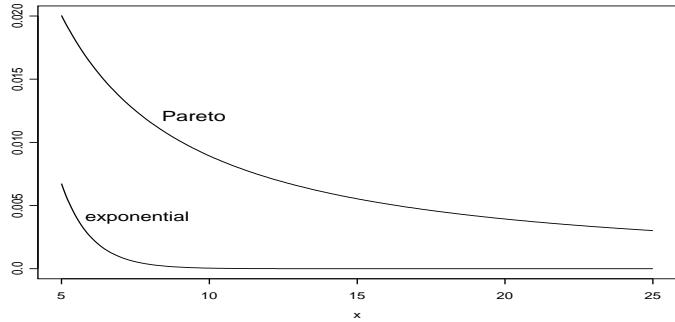


Figure 3.1: Tails of an exponential and a Pareto distribution, both with mean 1. For the Pareto distribution $\alpha = 1.2$.

Pareto distribution are depicted. In the light-tailed cases mentioned above the probability of a large sum of n iid random variables is asymptotically larger than the probability of a large maximum, i.e. the limit in (3.3) is infinite. Intuitively, this can be explained as follows. For $\sum_{i=1}^n X_i$ to exceed a high threshold x all n random variables X_i need to make a contribution, since for the individual X_i probabilities of extremely large values are negligible. The probability of the maximum being larger than x is asymptotically equivalent to $n P(X_1 > x)$. Since less probability mass is located in the right tails the probability of all X_i being large, but not as large as x , asymptotically outweighs the probability of one of them being larger than x .

For illustration, we give the following examples. Let the X_i be iid with common distribution F . Denote the maximum and the sum of X_1, \dots, X_n by M_n and S_n , respectively. Suppose F is exponentially distributed with mean $1/\lambda$ ($Exp(\lambda)$), then S_n has gamma distribution $\Gamma(n, \lambda)$. Moreover, as $x \rightarrow \infty$,

$$P(M_n > x) \sim n e^{-\lambda x}, \quad \text{and} \quad P(S_n > x) = e^{-\lambda x} \sum_{i=1}^n \frac{(\lambda x)^{i-1}}{i!}.$$

Next, let $F = \Phi$. We have

$$P(M_n > x) \sim n \overline{\Phi}(x),$$

and

$$P(S_n > x) = P(N(0, n) > x) = \overline{\Phi}\left(\frac{x}{\sqrt{n}}\right).$$

In both examples (using (3.5) in the second one) it is clear that $P(M_n > x) = o(P(S_n > x))$ as $x \rightarrow \infty$.

3.1.2 Detecting heavy tails

Here we discuss two methods to determine as to whether the distribution of an iid sample X_1, \dots, X_n has a heavy tail. Notice that by (3.1) a heavy tail is an asymptotic notion: the tail behavior is described for large x only. Hence, in determining whether a heavy tail is present only the largest values in the random sample play a role. Since it is virtually impossible to distinguish between the tails $cx^{-\alpha}$, $c > 0$, and $x^{-\alpha}L(x)$ for a more general slowly varying function L , the emphasis is on estimating the tail parameter α .

The first method uses the *empirical complementary distribution function* $1 - \hat{F}_n$, where \hat{F}_n is the empirical distribution function, i.e.

$$\hat{F}_n(x) = \frac{1}{n} \sum_{i=1}^n I_{(-\infty, x]}(X_i), \quad x \in \mathbb{R}.$$

The idea is to plot $1 - \hat{F}_n$ on log-log scales; this is called the log-log *complementary distribution*, or LLCD, plot. If the theoretical distribution function has a heavy right tail then, in agreement with the Glivenko-Cantelli theorem (see e.g. Billingsley [8], Theorem 20.6), for large n and moderately large values x , the LLCD plot should consist of points randomly scattered around a straight line with slope $-\alpha$. An estimate for α can then be obtained by least squares regression. In Figure 3.2 LLCD plots are shown for 100.000 realizations from a standard exponential and a Pareto distribution with $\alpha = 1.2$. In the exponential case the slope of the curve becomes more and more negative, while in the Pareto case the slope remains roughly the same. However, the estimate for α strongly depends on the number of points used for the least squares regression.

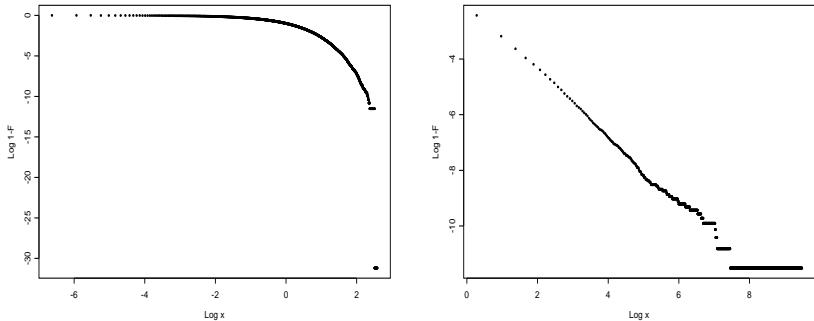


Figure 3.2: LLCD plots of 100.000 realizations from an exponential (left) and a Pareto (right) distribution, both with mean 1. For the Pareto distribution $\alpha = 1.2$.

The perhaps best known and most frequently used method for estimating the tail parameter α is the *Hill estimator*. It has the following form:

$$\hat{\alpha}_{k,n} = \left(\left(\frac{1}{k} \sum_{j=1}^k \log(X_{j,n}) \right) - \log(X_{k,n}) \right)^{-1},$$

where $X_{n,n} \leq \dots \leq X_{1,n}$ are the order statistics of the sample. For estimating α only the k largest order statistics are used. Under suitable conditions on $k = k(n) \rightarrow \infty$ asymptotic normality and strong consistency of the Hill estimator can be shown (see Embrechts et al. [28], Theorem 6.4.6). To obtain an estimate of α one plots $\hat{\alpha}_{k,n}$ against k for a variety of values k which are small compared to n ($k/n \rightarrow 0$ is required for consistency). This is called the *Hill plot*. In the heavy-tailed case the estimator stabilizes at a value $\hat{\alpha}$ for appropriate values of k . This value $\hat{\alpha}$ is taken as an estimate for the tail parameter α . In Figure 3.3 Hill plots are shown for 100.000 realizations from a standard exponential and a Pareto distribution with $\alpha = 1.2$. In the exponential case the Hill plot decreases, while in the Pareto case it stabilizes a little below the true value of 1.2. Notice that the range of values k for which the estimator is considered plays a crucial role. In Figure 3.3, $k/n \leq 0.1$.

Detailed information on various methods used to estimate the tail parameter α and possible pitfalls therein can be found in Resnick [79], Drees et al. [24] and Embrechts et al. [28], Section 6.4.

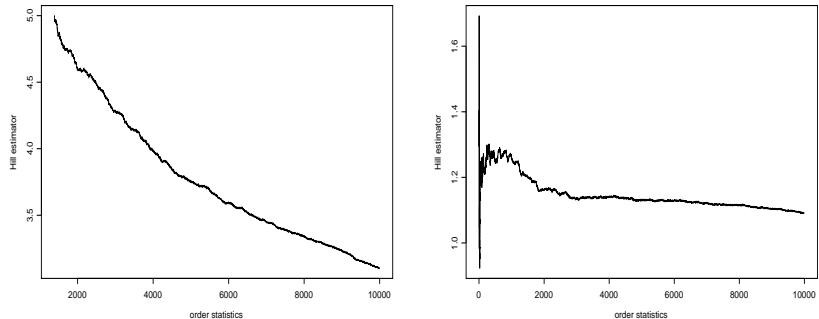


Figure 3.3: Hill plots for 100.000 realizations from an exponential (left) and a Pareto (right) distribution, both with mean 1. For the Pareto distribution $\alpha = 1.2$.

3.1.3 Stable Lévy motion

For later use, we define a stochastic process with stable marginal distributions. We say that a process $(\Lambda_{\alpha,\sigma,\beta}(t), t \geq 0)$ is α -stable *Lévy motion* if

- (1) $\Lambda_{\alpha,\sigma,\beta}$ has independent increments
- (2) $\Lambda_{\alpha,\sigma,\beta}$ has stationary increments
- (3) $\Lambda_{\alpha,\sigma,\beta}(t) \sim S_\alpha(\sigma t^{1/\alpha}, \beta, 0)$ for some $\alpha \in (0, 2]$, $\beta \in [-1, 1]$ and $\sigma > 0$
- (4) $\Lambda_{\alpha,\sigma,\beta}$ has right-continuous sample paths a.s.

Notice that if $\alpha = 2$ and $\sigma = 1/\sqrt{2}$, Λ is standard Brownian motion. In Figure 3.4 sample paths of stable Lévy motion and Brownian motion are compared. Stable Lévy motion has discontinuous sample paths with large jumps due to the heavy tailed marginal distributions, while Brownian motion has continuous sample paths and Gaussian marginals. For more information on stable processes and a constructive definition we refer to Samorodnitsky and Taqqu [89], Chapter 3.

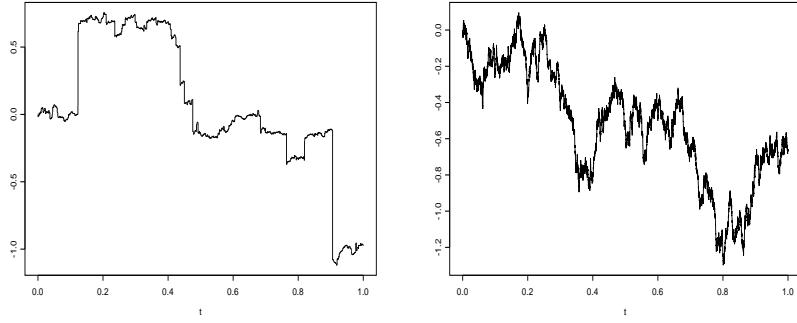


Figure 3.4: Sample paths of stable Lévy motion with $\alpha = 1.2$ (left) and Brownian motion (right).

3.2 Long-range dependence

3.2.1 Definitions

For a weakly stationary, and, hence, finite-variance stochastic process $(X_t, t = 0, \pm 1, \pm 2, \dots)$ dependence between observations at times t and $t + k$ is usually measured by the *autocorrelation function* (ACF) ρ at lag k , i.e.

$$\rho(k) = \frac{\text{Cov}(X_t, X_{t+k})}{\text{Var}(X_t)} = \frac{\text{Cov}(X_0, X_k)}{\text{Var}(X_0)}, \quad k = 0, \pm 1, \pm 2, \dots$$

By plotting $\rho(k)$ against k (called ACF plot) one can gain an idea of the second order dependence structure of the process. Naturally, one expects that $|\rho(k)|$ decreases as k increases. The notion of *long-range dependence* refers to the relative size of $\rho(k)$ at large lags k . If ρ is non-negligible at large lags, we have dependence over a *long range* of sequential observations. Formally, we say that a process has *long-range dependence* (LRD) if the sum of the absolute values of the autocorrelations is infinite, i.e.

$$\sum_{k=0}^{\infty} |\rho(k)| = \infty. \quad (3.6)$$

Alternatively, the term *long memory* is used for LRD. If the autocorrelations are absolutely summable the process is said to have *short-range dependence* or *short memory*. This is true for some important classes of Markov processes and ARMA processes, which are frequently applied in time series analysis (see Brockwell and Davis [15], Chapter 3). In the literature (see e.g. Beran [4]), LRD is also defined by describing the rate at which $\rho(k)$ decreases to zero as $k \rightarrow \infty$. Following this approach, a stationary process has LRD if

$$\rho(k) \sim c_{\rho} k^{-\beta}, \quad k \rightarrow \infty, \quad (3.7)$$

where c_{ρ} is a positive constant and $\beta \in (0, 1)$. Clearly, (3.7) implies (3.6). From an empirical perspective (3.7) is more suitable than (3.6) when checking for LRD in a supposedly stationary time series. Another reason why (3.7) is used to define LRD is that two very popular processes, fractional ARIMA and fractional Gaussian noise, both satisfy (3.7) with positive values $\rho(k)$. For fractional ARIMA, we refer to Brockwell and Davis [15], Section 13.2. Fractional Gaussian noise is introduced in Section 3.2.7.

An analogous way to define LRD is through the spectral density. The *spectral density* of a weakly stationary process $(X_t, t = 0, \pm 1, \pm 2, \dots)$ with $\sum_k [\rho(k)]^2 < \infty$ is defined (see Brockwell and Davis [15], Section 4.3) as

$$f(\lambda) = \frac{\sigma^2}{2\pi} \sum_{k=-\infty}^{\infty} \rho(k) e^{-ik\lambda}, \quad \lambda \in [-\pi, \pi],$$

where $\sigma^2 = \text{Var}(X_t)$. Since the autocorrelations are not absolutely summable under LRD, the spectral density can have a singularity at $\lambda = 0$. One can show that (3.7) is equivalent to

$$f(\lambda) \sim c_f \lambda^{\beta-1}, \quad \lambda \rightarrow 0, \quad (3.8)$$

where c_f is a positive constant depending on β and c_{ρ} . A proof will be provided in Section 3.2.2.

A third feature of LRD processes concerns the asymptotic behavior of the variance of the mean $\bar{X}_n = n^{-1} \sum_{t=1}^n X_t$. It can be shown that if (3.7) or, equivalently, (3.8) holds, then

$$\text{Var}(\bar{X}_n) \sim c_v \sigma^2 n^{-\beta}, \quad n \rightarrow \infty, \quad (3.9)$$

where c_v is a positive constant depending on c_ρ and β . The proof of (3.9) is contained in Section 3.2.2. For ARMA processes, which have short memory, $\text{Var}(\bar{X}_n)$ decreases to zero proportionally to n^{-1} .

3.2.2 Proofs

First, we show, under some technical conditions, the equivalence of (3.7) and (3.8). It is followed by a proof of the implication (3.7) \Rightarrow (3.9).

Proof of (3.7) \Rightarrow (3.8): Suppose (3.7) holds and write

$$\rho(k) = c_\rho |k|^{-\beta} L_1(|k|), \quad k = \pm 1, \pm 2, \dots, \quad (3.10)$$

where $L_1(x) \sim 1$ as $x \rightarrow \infty$. Hence, L_1 is slowly varying at infinity. Additionally, we assume that for any $\delta > 0$ and x large enough the functions $x^\delta L_1(x)$ and $x^{-\delta} L_1(x)$ are, respectively, increasing and decreasing in x . Using (3.10), the spectral density equals

$$\begin{aligned} f(\lambda) &= \frac{\sigma^2}{2\pi} \sum_{k=-\infty}^{\infty} \rho(k) \cos(k\lambda) \\ &= \frac{\sigma^2}{2\pi} + \frac{\sigma^2}{\pi} \sum_{k=1}^{\infty} \rho(k) \cos(k\lambda) \\ &= \frac{\sigma^2}{2\pi} + \frac{\sigma^2 c_\rho}{\pi} \sum_{k=1}^{\infty} k^{-\beta} L_1(k) \cos(k\lambda). \end{aligned}$$

An application of Zygmund [111], Chapter V.2, Theorem 2.6, yields

$$\frac{\sigma^2 c_\rho}{\pi} \sum_{k=1}^{\infty} k^{-\beta} L_1(k) \cos(k\lambda) \sim c_f \lambda^{\beta-1} L_1(\lambda^{-1}), \quad \lambda \downarrow 0,$$

where $c_f = \sigma^2 \pi^{-1} c_\rho \Gamma(1 - \beta) \sin(\pi\beta/2)$. Since the function $\lambda^{\beta-1} L_1(\lambda^{-1})$ is decreasing in λ the spectral density has a singularity at zero. Moreover, the behavior of $f(\lambda)$ as $\lambda \downarrow 0$ is determined by $\lambda^{\beta-1}$. This completes the proof. \heartsuit

Proof of (3.8) \Rightarrow (3.7): Suppose (3.8) holds and write

$$f(\lambda) = c_f |\lambda|^{\beta-1} L_2(|\lambda|), \quad \lambda \in [-\pi, \pi], \quad (3.11)$$

where $L_2(x) \sim 1$ as $x \downarrow 0$. Hence, L_2 is slowly varying at zero. Assume that for any $\delta > 0$ and x small enough the functions $x^\delta L_2(x)$ and $x^{-\delta} L_2(x)$ are, respectively, increasing and decreasing in x . Additionally, L_2 is assumed to be of bounded variation in (ϵ, π) for any $\epsilon > 0$. This last assumption rules out any

other singularities of $f(\lambda)$ with $\lambda \neq 0$. Using Herglotz's theorem (Brockwell and Davis [15], Section 4.3) and (3.11), the ACF equals

$$\begin{aligned}\rho(k) &= \frac{1}{\sigma^2} \int_{-\pi}^{\pi} f(\lambda) e^{ik\lambda} d\lambda \\ &= \frac{2}{\sigma^2} \int_0^{\pi} f(\lambda) \cos(k\lambda) d\lambda \\ &= \frac{2c_f}{\sigma^2} \int_0^{\pi} |\lambda|^{\beta-1} L_2(\lambda) \cos(k\lambda) d\lambda.\end{aligned}$$

An application of Zygmund [111], Chapter V.2, Theorem 2.24, yields

$$\frac{2c_f}{\sigma^2} \int_0^{\pi} |\lambda|^{\beta-1} L_2(\lambda) \cos(k\lambda) d\lambda \sim c_{\rho} k^{-\beta} L_2(k^{-1}), \quad k \rightarrow \infty,$$

where $c_{\rho} = 2\sigma^{-2}c_f\Gamma(\beta)\sin(\pi(1-\beta)/2)$. By the nature of the function L_2 , the asymptotic behavior of $\rho(k)$ is given by $k^{-\beta}$. This completes the proof. \heartsuit

Proof of (3.7) \Rightarrow (3.9): We write

$$\text{Var}(\overline{X}_n) = \frac{\sigma^2}{n} (1 + \delta_n(\rho)),$$

where

$$\delta_n(\rho) = 2 \sum_{k=1}^{n-1} \left(1 - \frac{k}{n}\right) \rho(k).$$

Suppose (3.10) holds with the same assumptions on L_1 as above. Additionally, assume that L_1 is locally bounded on $[1, \infty)$. Write

$$\begin{aligned}\delta_n(\rho) &= 2 c_{\rho} \sum_{k=1}^{n-1} \left(1 - \frac{k}{n}\right) k^{-\beta} L_1(k) \\ &= 2 c_{\rho} \left[\sum_{k=1}^{n-1} k^{-\beta} L_1(k) - \frac{1}{n} \sum_{k=1}^{n-1} k^{1-\beta} L_1(k) \right].\end{aligned}$$

Since the functions $k^{-\beta} L_1(k)$ and $k^{1-\beta} L_1(k)$ are, respectively, decreasing and increasing in k , we have the inequalities

$$\int_1^n x^{-\beta} L_1(x) dx \leq \sum_{k=1}^{n-1} k^{-\beta} L_1(k) \leq 1 + \int_1^{n-1} x^{-\beta} L_1(x) dx,$$

and

$$1 + \int_1^{n-1} x^{1-\beta} L_1(x) dx \leq \sum_{k=1}^{n-1} k^{1-\beta} L_1(k) \leq \int_1^n x^{1-\beta} L_1(x) dx.$$

Applying Karamata's theorem (see Embrechts et al. [28], Theorem A3.6) on both left- and right-hand sides yields, as $n \rightarrow \infty$

$$\sum_{k=1}^{n-1} k^{-\beta} L_1(k) \sim \frac{1}{1-\beta} n^{1-\beta} \quad \text{and} \quad \sum_{k=1}^{n-1} k^{1-\beta} L_1(k) \sim \frac{1}{2-\beta} n^{2-\beta}.$$

Hence, as $n \rightarrow \infty$

$$\delta_n(\rho) \sim \frac{2 c_\rho}{(1-\beta)(2-\beta)} n^{1-\beta},$$

which implies that

$$\text{Var}(\bar{X}_n) \sim c_v \sigma^2 n^{-\beta},$$

with $c_v = 2 c_\rho [(1-\beta)(2-\beta)]^{-1}$. This completes the proof. \heartsuit

Remark. Slowly varying functions L with $x^\delta L(x)$ increasing and $x^{-\delta} L(x)$ decreasing for x large enough and any $\delta > 0$ belong to the so-called *Zygmund class* of slowly varying functions.

3.2.3 Detecting long-range dependence

As mentioned above, in practice one usually tries to verify (3.7) or (3.8) rather than (3.6). There are several exploratory methods to detect the presence of LRD in a time series and more sophisticated statistical methods to estimate the parameter β in (3.7) and (3.8). A detailed description can be found in Beran [4]. See also Abry and Veitch [1], who describe a wavelet based method for detecting LRD. For a performance analysis of various estimators we refer to Taqqu and Teverovsky [99, 100]. Here, we describe four methods which are frequently used as exploratory tools.

Suppose we observed a time series $(X_t, t = 1, \dots, n)$ and we believe it is stationary and ergodic. One way to detect LRD makes use of the sample autocorrelations $\hat{\rho}(k)$. Here, $\hat{\rho}(k) = \hat{\gamma}(k)/\hat{\gamma}(0)$, where

$$\hat{\gamma}(k) = \frac{1}{n} \sum_{t=1}^{n-k} (X_t - \bar{X}_n)(X_{t+k} - \bar{X}_n), \quad 0 \leq k \leq n-1,$$

is the sample autocovariance function. The division by n , and not $n-k$, is needed to ensure that the function ρ is non-negative definite (see Brockwell and Davis [15], Section 1.5). To check whether the sample autocorrelations $\hat{\rho}(k)$ satisfy (3.7) for large k , a plot of $\log(\hat{\rho}(k))$ against $\log(k)$, also called the log-log *correlogram*, is made. If LRD in the sense of (3.7) is present, the points in this plot should be randomly scattered around a straight line with slope $-\beta$ for appropriate values of k and large n . A disadvantage of this method is that the estimate $\hat{\rho}(k)$ is unreliable for large k with respect to n (see Brockwell and Davis [15], Section 7.2).

An alternative method to detect LRD is estimation of the spectrum of X with the aim to verify (3.8). A natural estimator of the spectral density is the

periodogram (see Brockwell and Davis [15], Section 10.3), which we define by

$$I(\lambda) = \frac{1}{2\pi n} \left| \sum_{t=1}^n X_t e^{-it\lambda} \right|^2, \quad \lambda \in [-\pi, \pi]. \quad (3.12)$$

It is common use to evaluate the periodogram at the *Fourier frequencies* $\lambda_j = 2\pi j/n$, $j = -[(n-1)/2], \dots, [n/2]$. If LRD in the sense of (3.8) is present, a plot of $\log(I(\lambda_j))$ against $\log(\lambda_j)$ would result in an approximately straight line with slope $\beta - 1$ for small frequencies λ_j . By least squares regression a naive estimator of β can be obtained (for a discussion see Geweke and Porter-Hudak [35] and Robinson [87]). One has to take into account, however, that if the autocorrelations are not summable the periodogram fluctuates much more for frequencies tending to zero, than for short memory processes (see Beran [4], Theorem 3.8).

The third method focuses on the variance of the sample mean \bar{X}_n . Given a supposedly stationary time series X we can check if (3.9) holds in the following way. The data are split into blocks of size m (with m/n small). For the j th block we compute the mean value $\bar{X}_j^{(m)}$, for $j = 1, \dots, [n/m]$. The sample variance S_m^2 of $\{\bar{X}_1^{(m)}, \dots, \bar{X}_{[n/m]}^{(m)}\}$ is an estimator of $\text{Var}(\bar{X}_m)$. By computing S_m^2 for successive values of m and plotting $\log(S_m^2)$ against $\log(m)$ LRD in the sense of (3.9) can be detected. This is also called a *log-log variance-time* plot. The resulting points should be scattered around a straight line with slope $-\beta$. If the process has short memory the slope should be -1 .

The oldest and perhaps best known method to detect LRD is the *R/S* method. It has its origins in the field of hydrology and was first used by Hurst [47] when he discovered LRD-like features in the yearly minimal water levels of the Nile River. For a stationary process $(X_t, t = 1, 2, \dots)$ with partial sums $Y_k = \sum_{t=1}^k X_t$ and sample variance $S_k^2 = k^{-1} \sum_{t=1}^k X_t^2 - k^{-2} Y_k^2$, the *R/S* statistic, or *rescaled adjusted range*, is given by

$$\frac{R(k)}{S(k)} = \frac{1}{S(k)} \left[\max_{0 \leq t \leq k} \left(Y_t - \frac{t}{k} Y_k \right) - \min_{0 \leq t \leq k} \left(Y_t - \frac{t}{k} Y_k \right) \right].$$

It can be shown that if X_t is Gaussian, stationary, ergodic and (3.7) holds, then

$$k^{-H} \frac{R(k)}{S(k)} \xrightarrow{d} \xi, \quad k \rightarrow \infty, \quad (3.13)$$

where ξ is a non-degenerate random variable (see Mandelbrot [63], Theorems 5 and 11, combined with Taqqu [97]) and $H = 1 - \beta/2$. The parameter H is the so-called *Hurst parameter* and is frequently used as a measure of the strength of LRD in the data (H close to 1 corresponds to a strong presence of LRD). For various short memory processes (3.13) holds with $H = 1/2$ (see Mandelbrot [63] and Feller [32]). For fractional Gaussian noise and fractional ARIMA (with Gaussian innovations), we have

$$E(R(k)/S(k)) \sim c_r k^H, \quad k \rightarrow \infty, \quad (3.14)$$

where c_r is a positive constant. For an observed time series $(X_t, t = 1, \dots, n)$ we can try to verify (3.14) as follows. Partition the series in $[n/m]$ blocks of size m . Then, for each k , compute $R(m_i, k)/S(m_i, k)$, starting at points $m_i = i m + 1$, $i = 0, 1, \dots$, such that $m_i + k \leq n$. For values of k smaller than m , we get $[n/m]$ different estimates of $R(k)/S(k)$. For values of k approaching n , we get fewer values, as few as 1 when $k \geq n - m$. Next, plot $\log(R(m_i, k)/S(m_i, k))$ against $\log(k)$ and get, for each k several values on the plot. For large k , the points should lie around a straight line with slope H .

3.2.4 Application to the Bellcore data

We apply the four methods described above to a well-known time series in the field of teletraffic, **BC-pAug**, representing the number of packet arrivals per second on the Ethernet LAN at the Bellcore company. **BC-pAug** has length $n = 3142$, dates from August 1989 and, for now, we assume it to be stationary. More details of **BC-pAug** can be found in Section 4.2.3. In Figure 3.5 the series, its sample ACF and its log-periodogram are depicted. The slowly decaying ACF and the peak of the periodogram at zero suggest the presence of LRD. In Figure 3.6 the corresponding log-log correlogram, log-log periodogram, log-log variance-time and log-log R/S plots are drawn. Least squares regression in these plots yields naive estimates of β and $H = 1 - \beta/2$ (see Table 3.1). In the variance-time and R/S plots also the lines representing the short memory case are depicted (with slopes -1 and $1/2$, respectively). The estimates of β and H are all in the range of LRD and do not heavily depend on the method used to obtain them. However, it is natural that the estimates obtained using the log-log correlogram or the log-log periodogram are rather sensitive to the number of points included in the regression. Therefore, only the variance-time and R/S methods will be used in the sequel. Notice that the usual procedure to obtain confidence bands for the estimates of β and H cannot be used, since all four methods violate the assumption of uncorrelated residuals.

method	estimate β	estimate H
log-log correlogram	0.35	0.83
log-log periodogram	0.33	0.83
log-log variance-time	0.32	0.84
log-log R/S	0.24	0.88

Table 3.1: Estimates of the strength of LRD in **BC-pAug** using four naive methods.

3.2.5 Long-range dependence or non-stationarity?

From the definition of LRD it is clear that making a statement concerning the presence of LRD in a time series involves the assumption of stationarity. The

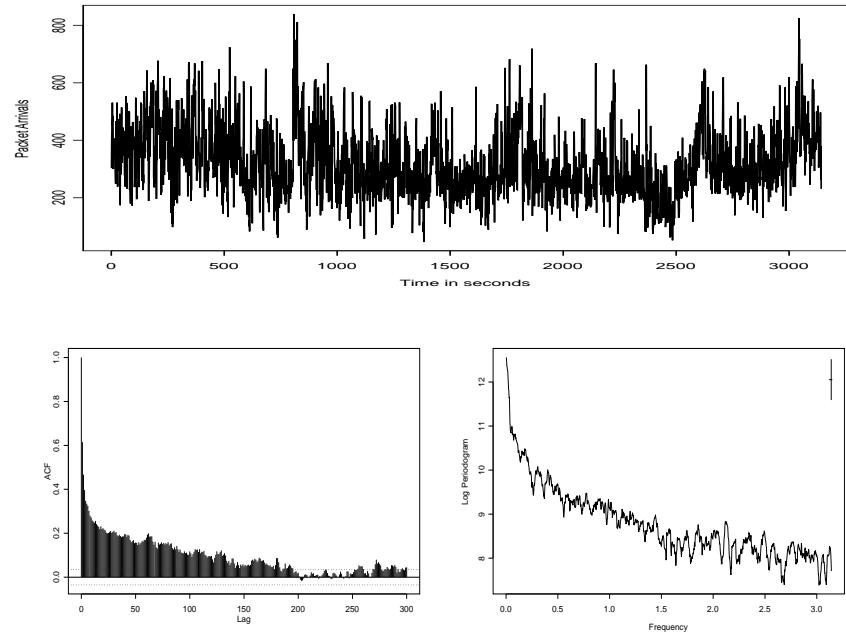


Figure 3.5: Time series **BC-pAug** (top) with sample autocorrelations (bottom-left) and smoothed log-periodogram (bottom-right).

question as to whether an observed time series can be considered as a realization from a stationary stochastic process is more philosophical than mathematical in nature. A stationary stochastic process is usually thought of as a process whose underlying distributional characteristics do not change over time. This concept has been made mathematically precise through the definitions of strict and weak stationarity (see e.g. Brockwell and Davis [15], Section 1.3). There is, however, no general statistical test available for the hypothesis of stationarity of a real-life time series. In practice, the assumption of stationarity is often justified by first fitting the data to a stationary model and then performing goodness-of-fit tests. Also, additional information can be used to decide whether it is reasonable to assume stationarity. However, as mentioned by Klemeš [54], any finite time series can be equally likely regarded as coming from a stationary or non-stationary model. Deviations from stationarity like trends, cycles or level shifts can be explained as ‘local non-stationarities’ in the first case or as ‘model characteristics’ in the second. An analogous reasoning applies to a time series with no apparent non-stationarities. It can either be seen as a ‘stationary looking’ realization from a non-stationary stochastic process or as coming from

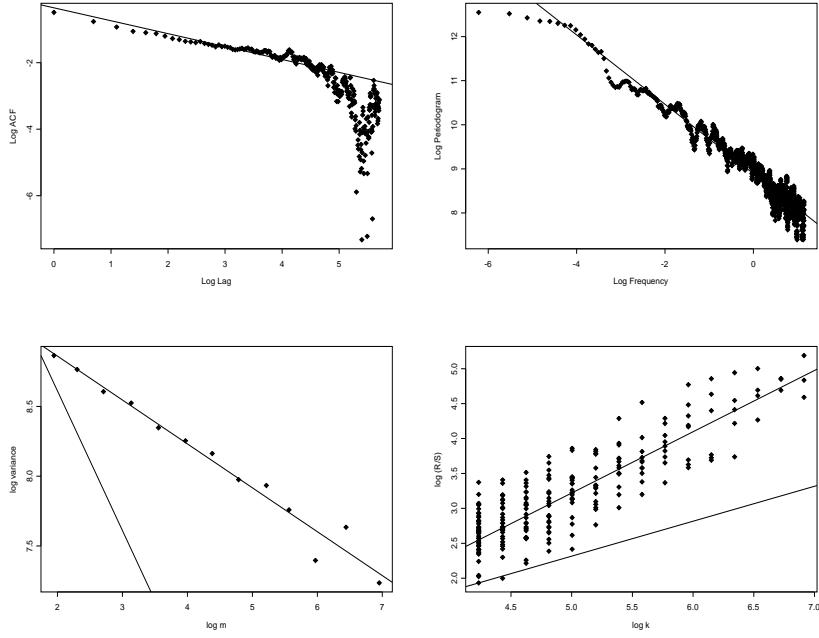


Figure 3.6: Exploratory plots to detect LRD in **BC-pAug**: log-log correlogram (top-left), log-log periodogram (top-right), log-log variance-time (bottom-left) and log-log R/S (bottom-right). Least squares regression yields heuristic estimates of β and H (see also Table 3.1). In the variance-time and R/S plot the lines representing the short memory case are depicted (with slopes -1 and $1/2$, respectively).

a truly stationary process. In this sense, the matter of stationarity or non-stationarity becomes one of personal belief.

Ever since Hurst [47] discovered features resembling LRD (also called the Hurst phenomenon) in the yearly minimal water levels of the Nile River, there has been much dispute about the stationarity assumption. Klemeš [54] shows that a realization from a sequence of independent random variables with shifts in the mean, can lead one to falsely conclude the presence of LRD on the basis of an R/S plot. See also Boes and Salas [11]. Another non-stationarity that can cause LRD-like phenomena is a slowly decaying trend, e.g. $ct^{-\nu}$ with $c > 0$ and $\nu \in (0, 1)$. This is studied in Bhattacharya et al. [6] (R/S plot), Künsch [56] (periodogram) and Teversonsky and Taqqu [102] (variance-time plot).

3.2.6 ARIMA models and long-range dependence

The class of ARMA processes is the class of models most frequently applied to time series that exhibit no apparent deviations from stationarity and have rapidly decreasing autocorrelations. ARMA processes are defined as follows. For non-negative integers p and q , let the polynomials ϕ and θ be given by

$$\phi(z) = 1 - \phi_1 z - \cdots - \phi_p z^p \quad \text{and} \quad \theta(z) = 1 + \theta_1 z + \cdots + \theta_q z^q.$$

Let the *backward shift operator* B be defined by

$$B^j X_t = X_{t-j}, \quad j = 0, 1, 2, \dots$$

The process $(X_t, t = 0, \pm 1, \pm 2, \dots)$ is an ARMA(p, q) process if (X_t) is stationary and if for every t

$$X_t - \phi_1 X_{t-1} - \cdots - \phi_p X_{t-p} = Z_t + \theta_1 Z_{t-1} + \cdots + \theta_q Z_{t-q}, \quad (3.15)$$

where $(Z_t, t = 0, \pm 1, \pm 2, \dots)$ is a sequence of *white noise*, i.e. uncorrelated and identically distributed random variables with mean 0 and variance σ_Z^2 .

The ARMA processes are a class of parametric processes defined by the set of linear difference equations (3.15). Equation (3.15) can be compactly written as $\phi(B)X_t = \theta(B)Z_t$. The name ARMA is an abbreviation of *autoregressive moving average*, the polynomial ϕ being the autoregressive part and θ the moving average part of the process. The parameters p and q are the *order* of the process. Notice that an ARMA process has mean zero. We say that (X_t) is an ARMA(p, q) process with mean μ if $(X_t - \mu)$ satisfies (3.15). If the polynomials ϕ and θ have no common zeros and ϕ has no zeros inside or on the unit circle in the complex plane, the process (X_t) is called *causal*. A causal ARMA process can be written as an infinite moving average $X_t = (\theta(B)/\phi(B))Z_t$, which is the unique stationary solution to (3.15). An ARMA process has short memory in the sense that the autocorrelation function $\rho(k)$ satisfies

$$|\rho(k)| \leq b a^k, \quad (3.16)$$

for some constants $b \in (0, \infty)$ and $a \in (0, 1)$. Clearly, (3.16) guarantees absolutely summable autocorrelations. More properties of ARMA process can be found in Brockwell and Davis [15], Chapter 3.

The ARMA class is very suitable for modeling stationary short memory time series. In practice, however, ‘stationarity’ of a time series is not always observed. Often, transformations are applied to “make the series look more stationary”. A frequently applied method is differencing, i.e. instead of the original series (X_t) one considers

$$(1 - B)X_t = X_t - X_{t-1}.$$

This will remove a global linear trend. Differencing twice, i.e. $(1 - B)^2 X_t$, will remove second degree polynomials. In practice the number of times differencing required to get an apparent stationary series is often not larger than one or two.

This is because, on a finite interval, many functions can be well approximated by polynomials of a reasonably low degree.

A process that, after differencing finitely many times, reduces to an ARMA process, is called an ARIMA process (autoregressive integrated moving average). An ARIMA process that is ARMA(p,q) after d times differencing is denoted by ARIMA(p,d,q). Here, d is a non-negative integer and is called the *order of differencing*. Hence, if (X_t) is ARIMA(p,d,q) then $(1 - B)^d X_t$ is ARMA(p,q). By definition, ARIMA($p,0,q$) is ARMA(p,q). Notice that if $d \geq 1$ the ARIMA process is not stationary.

It is well-known that the ACF of an ARMA process can show slow decay to zero if the polynomial ϕ has a zero close to the unit circle (see Brockwell and Davis [15], Section 3.3). An ARIMA(p,d,q) process has representation $\phi^*(B)X_t = \theta(B)Z_t$, where $\phi^*(B) = (1 - B)^d \phi(B)$. Hence, the polynomial $\phi^*(z)$ has a zero of order d at $z = 1$. This suggests that the sample ACF of a realization from an ARIMA process can show slowly decaying behavior. The next example confirms this and also considers the periodogram, variance-time and R/S plots. We simulated an ARIMA(1,1,1) model with $\phi_1 = 0.4$, $\theta_1 = -0.95$ and $\sigma_Z^2 = 8000$. The parameter values correspond to the estimated ARIMA model for **BC-pAug** in Section 4.2.2. The series length is $n = 3000$. In Figure 3.7 the simulated ARIMA series together with its sample autocorrelations, smoothed log-periodogram, log-log variance-time and log-log R/S plots are depicted. The autocorrelations show slow decay, the periodogram has a peak at zero and the variance-time and R/S plots yield a least squares line which differs significantly from the short memory case. The estimates of the Hurst parameter H in the latter two plots are 0.88 and 1.01, respectively. Hence, if the simulated series is interpreted as being stationary, these four exploratory methods are likely to yield the conclusion that LRD is present in the data.

3.2.7 Fractional Brownian motion

Here we define a stochastic process exhibiting LRD in the sense of (3.7). A process $(\sigma_0 B_H(t), t \geq 0)$ with $\sigma_0 > 0$ is called *fractional Brownian motion* if

- (1) $B_H(t) \sim N(0, \sigma^2 t^{2H})$
- (2) $\text{Cov}(B_H(s), B_H(t)) = \frac{1}{2} (s^{2H} + t^{2H} - |t - s|^{2H})$ for some $H \in (0, 1)$.
- (3) B_H has continuous sample paths a.s.

Since fractional Brownian motion is a Gaussian process, its finite-dimensional distributions are completely determined by (1) and (2). An application of Kolmogorov's existence theorem (see e.g. Billingsley [8], Section 36) guarantees that the process is properly defined in this way. Using (2), it can be shown that for $s < t$

$$\text{Var}(B_H(t) - B_H(s)) = \text{Var}(B_H(t - s)).$$

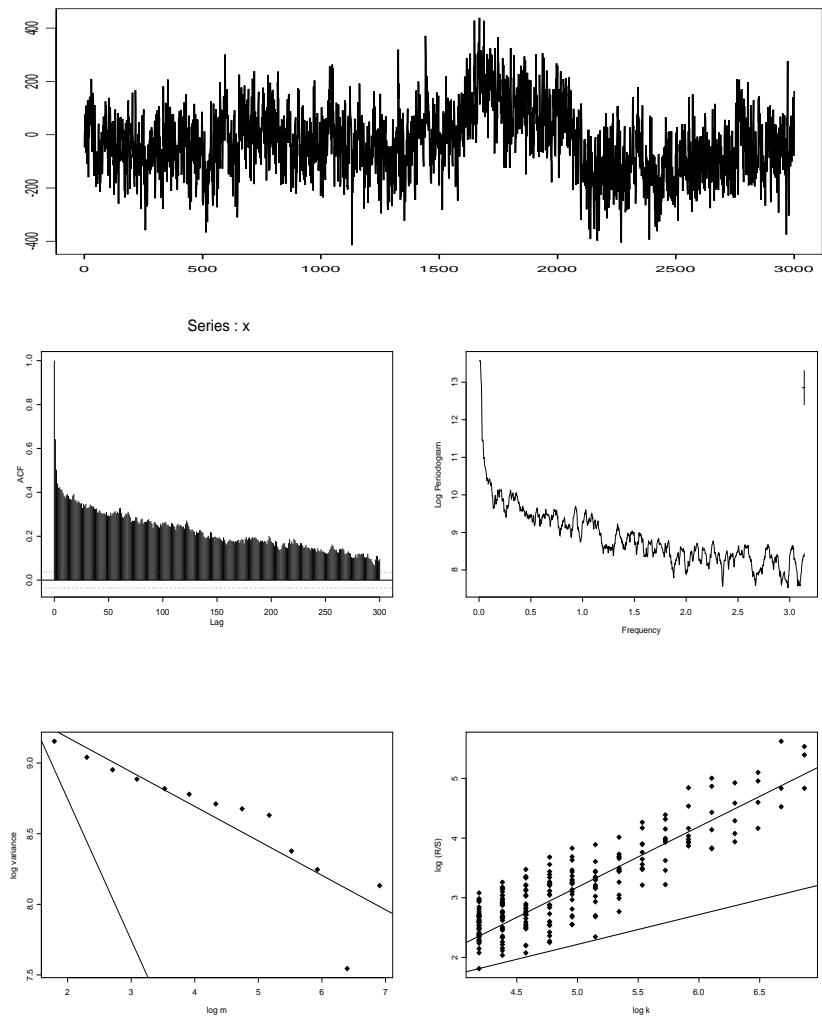


Figure 3.7: Simulated ARIMA(1,1,1) series (top) with sample autocorrelations (middle-left), smoothed log-periodogram (middle-right), log-log variance-time (bottom-left) and log-log R/S (bottom-right) plots.

Hence, fractional Brownian motion has stationary increments. Notice that for $H = 1/2$, B_H is ordinary Brownian motion. In Figure 3.8 sample paths of fractional Brownian motion are drawn for $H = 0.3$ and $H = 0.8$. Fractional Brownian motion has continuous sample paths which become smoother as H is increased.

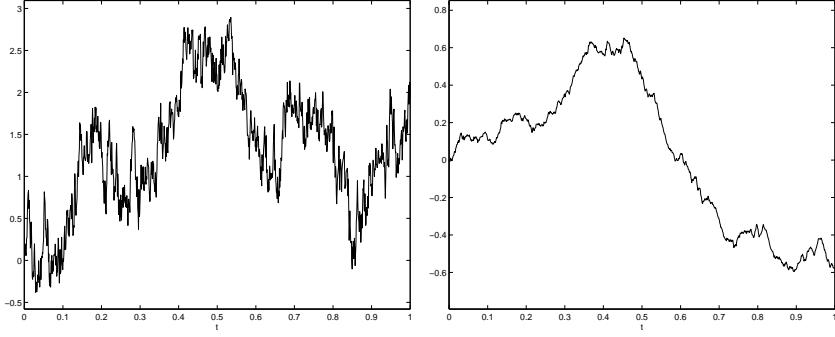


Figure 3.8: Sample paths of fractional Brownian motion with $H = 0.3$ (left) and $H = 0.8$ (right).

The increment process $Y_t = B_H(t) - B_H(t - 1)$, $t = 1, 2, \dots$, of fractional Brownian motion is called *fractional Gaussian noise*. It is a mean-zero stationary Gaussian process with ACF

$$\rho_Y(k) = \frac{1}{2} (|k + 1|^{2H} - 2|k|^{2H} + |k - 1|^{2H}) , \quad k = 0, 1, 2, \dots$$

For $H = 1/2$, Y is a sequence of iid $N(0, \sigma^2)$ variables. If $H \neq 1/2$, Y is a dependent sequence. It can be shown (see Samorodnitsky and Taqqu [89], Proposition 7.2.10) that for $H \neq 1/2$, as $k \rightarrow \infty$

$$\rho_Y(k) \sim H(2H - 1)k^{2H-2}.$$

Hence, for $H \in (1/2, 1)$ fractional Gaussian noise exhibits LRD in the sense of (3.7). If $H \in (0, 1/2)$ the autocorrelations $\rho_Y(k)$ are absolutely summable and, hence, the process has short memory. More details of fractional Brownian motion and fractional Gaussian noise can be found in Samorodnitsky and Taqqu [89], Section 7.2.

3.3 Self-similarity

3.3.1 Definition

A stochastic process $(X_t, t \geq 0)$ is said to be *self-similar* if the finite-dimensional distributions of (X_{at}) and $(a^H X_t)$ are identical for any $a > 0$ and some $H \in$

$(0, 1)$, i.e. if

$$(X_{at}, t \geq 0) \stackrel{d}{=} (a^H X_t, t \geq 0). \quad (3.17)$$

Hence, the distribution of a self-similar process is invariant under a particular scaling of time and space. The parameter H is called the *index of self-similarity*. It can be shown that α -stable Lévy motion is self-similar with $H = 1/\alpha$. Fractional Brownian motion is self-similar with index H and, naturally, for Brownian motion this is true with $H = 1/2$. For more self-similar processes we refer to Samorodnitsky and Taqqu [89], Chapter 7.

A different notion of self-similarity is given by Cox [19]. He considers a weakly stationary process $(Y_t, t \geq 0)$ and a new series $(Y_t^{(m)}, t \geq 0)$, given by

$$Y_t^{(m)} = \frac{1}{m}(Y_{tm-m+1} + \cdots + Y_{tm}),$$

where $m \geq 1$. Hence, $Y^{(m)}$ is formed by averaging the original series Y in adjacent blocks of size m , replacing each block by its mean. The process $Y^{(m)}$ is also weakly stationary. Cox calls the process Y *exactly second-order self-similar* if Y and $Y^{(m)}$ have the same ACF for each $m \geq 1$. It can be shown that fractional Gaussian noise satisfies this property. If the autocorrelation structures of Y and $Y^{(m)}$ coincide only when $m \rightarrow \infty$, the process Y is called *asymptotically second-order self-similar*. It can be shown that this is the case when the ACF of Y shows slow decay in the sense of (3.7). An example is fractional ARIMA. It has to be noted that these notions of self-similarity refer to the behavior of the ACF at large lags and, hence, must not be confused with (3.17), which concerns the whole distribution of the process. Moreover, Cox [19] considers a stationary process, while a process satisfying (3.17) is necessarily non-stationary.

3.3.2 Self-similarity ‘by picture’

It is tempting to argue that a sample path of a self-similar (with index H) process on $[0, 1]$ will look qualitatively the same as a sample path on $[0, 100]$, where the realizations are divided by 100^H . However, self-similarity means that the distribution of the process is invariant under the transformation and not necessarily the sample paths. The method we just described is also used in teletraffic research. In Leland et al. [60] and Willinger et al. [106] a time series of measured packet arrivals per time unit on the Ethernet LAN at Bellcore is plotted on five different time scales, the time units ranging from 0.01 up to 100 seconds. From the plots it can be seen that the relative variability of the arrival process remains roughly the same in four of the five plots. The authors conclude that evidence has been found of the “self-similar characteristics” of the measured Ethernet traffic. In the next example we demonstrate that the same conclusion can be drawn for an appropriate time series, which is not self-similar. We start with a realization from an ARIMA(1,1,1) model with $\phi_1 = 0.4$, $\theta_1 = -0.95$ and $\sigma_Z^2 = 8000$, which has length $n = 10.000.000$. The parameter values correspond to the estimated ARIMA model for **BC-pAug** in Section 4.2.2. From this sequence, we construct four new sequences by taking sums over

consecutive blocks of sizes 10, 100, 1000 and 10.000, respectively. In Figure 3.9 the series with block size 10.000 is shown in the top-left plot. It has length 1000. In the top-right plot we have ‘zoomed in’ on blocks 700 up to 800 of the top-left plot. In this new graph, the block size is 1000 and 1000 values are plotted. In this way we keep ‘zooming in’ on the graph until we have reached a block size of 10 in the bottom-right plot. This is about the same procedure as used by Leland et al. [60] and Willinger et al. [106] to investigate the self-similar nature of the packet arrivals.

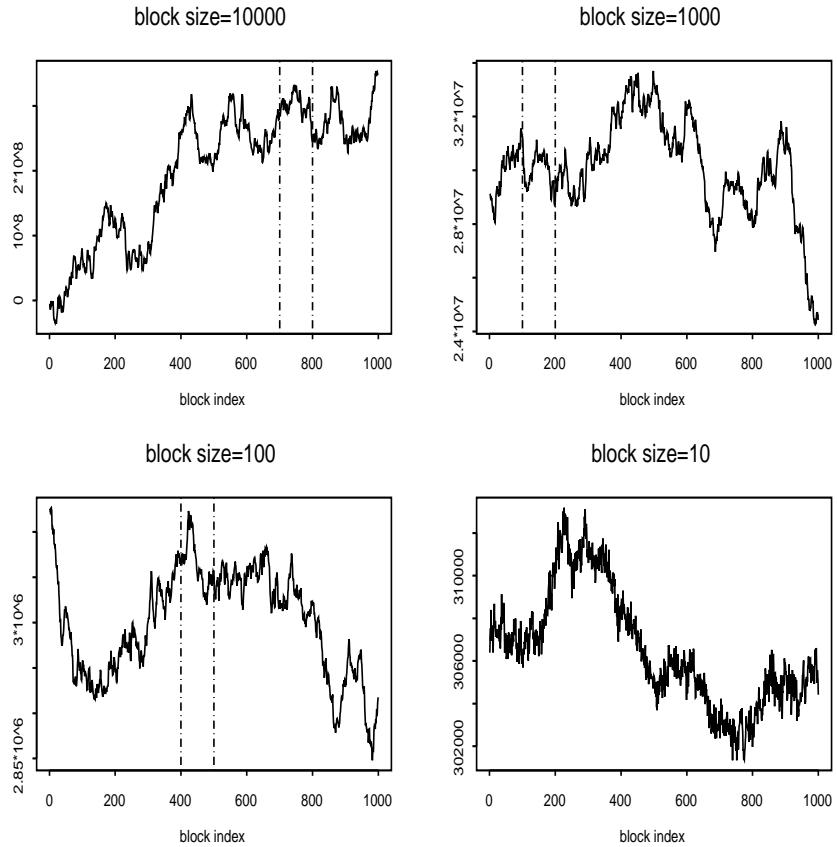


Figure 3.9: Consecutive ‘zooming in’ on a realization from an ARIMA model reveals its self-similar nature?

As we see from Figure 3.9, the relative variability of the four plots remains roughly the same. However, by taking this as an indication of the self-similar nature of the process we are far from the truth, since the realization is from an ARIMA model, which is not self-similar. As we have seen in Figure 3.7, the sam-

ple ACF of a realization from an ARIMA model with the specified parameter values above shows slowly decaying behavior. Since slowly decaying autocorrelations in the sense of (3.7) imply asymptotic second-order self-similarity in the sense of Cox [19], this may be an explanation for the phenomenon occurring in Figure 3.9.

Four

Traffic Data Analysis

In this chapter we focus on the statistical analysis of real-life teletraffic data. In Section 4.1 we give an overview of some of the most influential empirical studies of computer network traffic. Using the statistical tools described in Chapter 3, these studies suggest the presence of LRD and self-similarity in workload measurements and establish the existence of heavy tails in the empirical distributions of file lengths, transmission durations and idle times. In Section 4.2 we present our own analysis of some well-known measurement series available in the Internet Traffic Archive [49]. In Section 3.2.4 we mentioned that LRD-like phenomena can also be caused by non-stationarities. Even a realization of a non-stationary ARIMA process can, for appropriately chosen parameter values, induce LRD-like features in the autocorrelation function, periodogram, variance-time and R/S plots. We show that most of the workload measurements can be described by an appropriate ARIMA($p,1,q$) process, with fairly small values p and q . This suggests that non-stationarity rather than LRD is present in the workload measurements. The analysis in Section 4.2 was done in Stegeman [95]. In Section 4.3 we further investigate the hypothesis of non-stationarity by analyzing the sequence of packet sizes in the Bellcore data. We divide the packet sizes into five groups. Groups 1 up to 4 consist of 7 different packet sizes and make up 81 percent of all packets sent. Group 5 consists of all other packet sizes. It appears that the arrival processes of the packets in each group are (more or less) mutually uncorrelated. A possible explanation is that packets in Groups 1 up to 4 belong to four different processes running on the Bellcore Ethernet LAN during the measurement period. In the arrival processes of these groups clear non-stationarities are visible. Together with the fact that a computer network is a rather complex and rapidly changing environment, this makes the assumption of stationarity of the Bellcore measurements rather doubtful.

4.1 Overview

The statistical properties of computer network traffic significantly differ from those of voice traffic in the telephone system (see e.g. Fowler and Leland [34] or Willinger and Paxson [107]). Telephone calls, within an hour say, tend to arrive according to a homogeneous Poisson process, i.e. their inter-arrival times are exponentially distributed. The lengths of telephone calls have an exponentially bounded tail. Although variations from hour to hour are significant a pattern can be discerned which tends to repeat itself. Also, there are predictable seasonal variations. The amount of voice traffic grows from year to year but does so slowly and steadily. On a sufficiently large time scale the number of call arrivals is approximately equal to the long-term average, which is determined by the arrival rates of the Poisson processes governing the call arrivals in different time intervals. From an engineering perspective this is a very convenient property: everything boils down to the knowledge of the long-term arrival rates.

However, when computers instead of humans are communicating, everything changes. Unlike voice traffic, workload measurements in computer networks (e.g. the number of packets or bytes arriving per time unit) show a high level of variability on every time scale that is considered, from milliseconds to minutes. For the measurements performed in 1989 on the Ethernet LAN at Bellcore this conclusion is drawn by Leland and Wilson [59] and Fowler and Leland [34]. The ‘burstiness’ across an extremely wide range of time scales resembles a realization of a stochastic process with LRD. Indeed, periodic components with arbitrarily long periods seem to be present in the data. Also, the connection with the concept of second-order self-similarity of Cox [19] (see Section 3.3) has been made. Several empirical studies conclude that LRD is present in workload measurements on the basis of periodogram, variance-time and R/S plots. The most influential among these studies is due to Leland et al. [60], who provide a statistical analysis of the workload in the Ethernet LAN at Bellcore from 1989 until 1992. Variable-Bit-Rate (VBR) video traffic is analyzed by Beran et al. [5]. They conclude that LRD is present in 20 VBR video sequences, containing TV programs as well as movie scenes and video conferencing images. The data are obtained from Siemens, Alcatel and Bellcore. Paxson and Floyd [70] find evidence for LRD in Wide Area traffic measurements at Bellcore (1989), Digital Equipment Corporation (1993-1995) and Lawrence Berkeley Laboratory (1994). Abry and Veitch [1], using their wavelet based technique for estimating the Hurst parameter H , find that $H \approx 0.8$ for the Bellcore data, which is consistent with the LRD hypothesis. In Leland et al. [60] and Willinger et al. [106] the variability of the workload on the Bellcore Ethernet is shown to be roughly the same on five different time scales. This is believed to be ‘pictorial proof’ of self-similarity in the data (see Section 3.3.1).

Another difference between voice traffic in the telephone system and computer network traffic concerns the distributions of transmission durations and idle times. While in the telephone system these tend to be exponential, or at least exponentially tailed, in computer networks they appear to be heavy-tailed. This means that the Poisson process cannot be used to model arrival processes

of computer-generated traffic, a conclusion drawn by Paxson and Floyd [70] who analyze Telnet connections and FTP sessions on WANs. They find that Telnet packet inter-arrival times within a Telnet connection and FTP data burst lengths within an FTP session are heavy-tailed. The arrival processes of Telnet connections and FTP sessions, however, can be modeled by a Poisson process. This is explained by the fact that a Telnet connection is set up by human action, while the traffic generated within a Telnet connection is the result of a dialogue between computers. Willinger et al. [106] consider the traffic between individual source-destination pairs in the Ethernet at Bellcore. They define a so-called *OFF-period* as an interval of length greater than some threshold value in which no packet is sent. The interval between two OFF-periods is called an *ON-period*. For several choices of the threshold value and different source-destination pairs it is shown that the distributions of lengths of ON- and OFF-periods are heavy-tailed. Other influential studies are due to Crovella and Bestavros [20, 21]. They analyze WWW file requests from a file server at Boston University in 1995 and find that the distribution of transmission durations is heavy-tailed. Also, the distribution of lengths of files available on several WWW file servers is found to be heavy-tailed. The latter observation is seen by the authors as an explanation of the heavy-tailed transmission durations. Finally, Crovella and Bestavros consider the OFF-periods of individual workstations, an OFF-period being the time between the completion of the transmission of the last file requested from the files server and the arrival of the request for the next file at the files server. The lengths of the OFF-periods are shown to have a heavy-tailed distribution.

The fact that the distributions of transmission durations and idle times are heavy-tailed implies that extremely long transmission and idle periods occur with non-negligible probability. Intuitively, this explains why the workload on the network is dependent over a long range of time points and, hence, exhibits LRD. This is exactly the idea captured by the ON/OFF model for LAN traffic proposed by Willinger et al. [106]: individual sources with heavy-tailed ON- and OFF-periods generate a workload process with LRD. The ON/OFF model is discussed in detail in Section 5.1.1. Also, the so-called infinite source Poisson model for WAN traffic, with Poisson connection arrivals and heavy-tailed connection durations, yields an LRD workload process (see Section 5.1.3).

Because of the high variability of transmission durations and idle times and the burstiness of the workload across an extremely wide range of time scales, computer networks are a far greater challenge to an engineer than the telephone system. The effect on network performance of the traffic properties described above is discussed in Section 5.3.

4.2 LRD or ARIMA?

In Section 3.2.5 we mentioned that LRD-like phenomena can also be caused by non-stationarities, e.g. shifts in the mean or a slowly decaying trend. A way to distinguish between non-stationarities and LRD is to estimate the Hurst parameter H for various parts of a time series. If the changes in the estimated

value of H are ‘reasonably large’, it is not plausible that the series is stationary (see e.g. Duffield et al. [25]). Notice that this procedure is not statistically sound, since confidence bounds for the estimated value of H are not available. As an example, we consider four 500-second parts of **BC-pAug**. The parts are chosen such that they show significant differences in the qualitative behavior of **BC-pAug**. In Table 4.1 we give the estimated values of H based on the variance-time and R/S plots. It is evident that almost all estimates are significantly smaller than the overall values of 0.84 and 0.88, respectively. The estimates obtained from the R/S plot are very close to $H = 1/2$, which is the short memory case. If the variance-time method is used the estimated values increase as the behavior of **BC-pAug** becomes more wild and ‘less stationary’ (a time series plot of **BC-pAug** can be found in Figure 3.5). It is worthwhile mentioning, however, that the wavelet based estimator of Abry and Veitch [1] does find consistent values of $H \approx 0.8$ for different segments of **BC-pAug**. The fact that the estimates in Table 4.1 differ from those obtained from the whole series may, therefore, also result from a lack of robustness of the variance-time and R/S estimation procedures. This would also explain the large differences between the estimated values obtained by the variance-time and R/S methods.

part	log-log variance-time	log-log R/S
1-3142	0.84	0.88
51-550	0.64	0.49
851-1350	0.69	0.54
1601-2100	0.79	0.53
2643-3142	0.87	0.51

Table 4.1: Estimates of the Hurst parameter H for four 500-second parts of **BC-pAug**.

A common method to ‘get rid of non-stationarities’ in a time series is to consider the differenced series. We apply this approach to **BC-pAug**. In Figure 4.1 it can be seen that the sample autocorrelations and the smoothed log-periodogram of the (first order) differences of **BC-pAug** show no sign of LRD at all. The peak at zero of the periodogram has disappeared and beyond lag 5 almost all autocorrelations do not differ significantly from zero. From the shape of the periodogram we may conclude that, in the differenced series, low frequencies do not play a significant role while high frequencies are all almost equally important. Since low frequencies correspond to long cycle periods this indicates their complete absence in the differenced series. Sometimes, also the autocorrelations of the absolute and squared values of the series are computed. This procedure is used for the analysis of financial time series in order to detect non-linearities in the series, see e.g. Embrechts et al. [28], Section 8.4. If they are outside the 95 percent confidence bounds at large lags, this is seen as an indication that LRD is present in $(|X_t|)$ or (X_t^2) (this could, however, also be caused by non-stationarities; see Mikosch and Stărică [67]). However, from Figure 4.1 it is

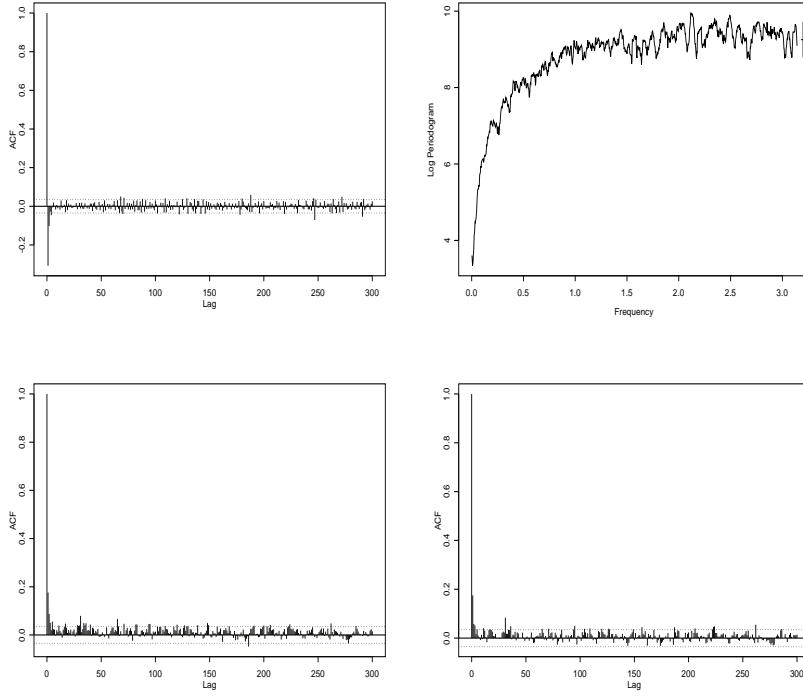


Figure 4.1: Sample autocorrelations (top-left) and smoothed log-periodogram (top-right) of the differences of **BC-pAug**. Bottom-left and bottom-right are the sample autocorrelations of the absolute and squared values of the differences, respectively.

clear that these autocorrelations are also small. Notice that the variance-time and R/S methods are useless to check whether $H \approx 1/2$, i.e. whether the differenced sequence has short memory, since the telescopic sum property of the differences results in artificially small cumulative sums.

In Section 3.2.6 we indicated that a realization of an appropriate non-stationary ARIMA process can exhibit LRD-like features. This suggests that ARIMA($p, 1, q$) may be a suitable model for **BC-pAug**. Before we fit **BC-pAug** to an ARIMA model, however, we first discuss two goodness-of-fit tests.

4.2.1 Two goodness-of-fit tests

Here we describe two methods for testing whether a time series is well modeled by a specified ARMA process. We assume that the orders p and q have been selected and that σ_Z^2 and the parameters of the polynomials ϕ and θ have been estimated by Gaussian maximum likelihood (for details, see Brockwell and Davis [15], Chapters 8 and 9).

The first method considers the sample autocorrelations of the estimated residuals. It can be shown (see Brockwell and Davis [15], Example 7.2.1) that the sample autocorrelations $\hat{\rho}_Z(k)$, $k \geq 1$, of an iid sequence Z_1, \dots, Z_n with $EZ_t^2 < \infty$ are for large n approximately iid with distribution $N(0, 1/n)$. Assuming that our fitted ARMA model is generated by an iid white noise sequence, the same approximation should be valid for the sample autocorrelations of the estimated residuals. Then the statistic $n \sum_{j=1}^k \hat{\rho}_Z^2(j)$ should be approximately chi-squared with k degrees of freedom. However, since the estimated residuals are a function of the maximum likelihood estimators of the parameters of the ARMA process, they do not constitute an iid sequence. This results in a reduction of the degrees of freedom by $p + q$ (see Brockwell and Davis [15], Section 9.4). Hence, a goodness-of-fit test statistic for the estimated ARMA model is

$$Q_Z(k) = n \sum_{j=1}^k \hat{\rho}_Z^2(j), \quad (4.1)$$

which is approximately chi-squared with $k - p - q$ degrees of freedom. The adequacy of the model is therefore rejected at level α if

$$Q_Z(k) > \chi_{1-\alpha}^2(k - p - q).$$

This test is known as the *Portmanteau test*. The number of lags k should not be chosen too large with respect to the length of the time series, since, for large lags, the sample autocorrelations are unreliable estimates.

The second method compares the periodogram $I(\lambda)$ of the time series with the theoretical spectral density $f(\lambda)$ of the estimated ARMA model. The spectral density of an ARMA process is given by

$$f(\lambda) = \frac{\sigma_Z^2}{2\pi} \frac{|\theta(e^{-i\lambda})|^2}{|\phi(e^{-i\lambda})|^2}, \quad \lambda \in [-\pi, \pi].$$

The idea, due to Bartlett [3], is to consider partial sums of the ratios $I(\lambda_j)/f(\lambda_j)$ for the positive Fourier frequencies λ_j . Self-normalization results in the statistics

$$T_h = \frac{\sum_{j=1}^{\lfloor nh/2 \rfloor} I(\lambda_j)/f(\lambda_j)}{\sum_{j=1}^{\lfloor n/2 \rfloor} I(\lambda_j)/f(\lambda_j)}, \quad h = 1/[n/2], 2/[n/2], \dots, 1.$$

It can be shown that, for large n , the process $\sqrt{n/2} (T_h - h)$ is approximately a standard Brownian bridge on $[0, 1]$ (see Priestley [76]). As a goodness-of-fit statistic one can take

$$T^* = \max_{1/[n/2] \leq h \leq 1} \sqrt{n/2} |T_h - h|, \quad (4.2)$$

which has the same limit distribution as the Kolmogorov-Smirnov statistic. The 95 and 99 percent quantiles of T^* are 1.36 and 1.63, respectively (see Shorack and Wellner [90], Table 1 on p. 143). For an overview of goodness-of-fit tests in the spectral domain we refer to Priestley [76], Section 6.2.6.

4.2.2 Fitting BC-pAug to an ARIMA model

We show that ARIMA(3,1,1) is an appropriate model for **BC-pAug**. Combined with the observation that a non-stationary ARIMA process can exhibit LRD-like phenomena (see Section 3.2.6), this is an argument in favor of the point of view that the LRD features of **BC-pAug** are actually caused by non-stationarities.

Using Gaussian maximum likelihood in the statistical package **Splus**, we fitted the differences of **BC-pAug** to an ARMA(3,1) model. The orders $p = 3$ and $q = 1$ are chosen using the Akaike Information Criterion (see Brockwell and Davis [15], Section 9.2). However, orders $(p, q) = (1, 1)$ until $(5, 1)$ all result in reasonably good fits. The estimated values of the parameters are $\hat{\phi}_1 = 0.42$, $\hat{\phi}_2 = 0.05$, $\hat{\phi}_3 = 0.01$, $\hat{\theta}_1 = -0.94$ and $\hat{\sigma}_Z^2 = 7703$. In Figure 4.2 the sample autocorrelations of the residuals and a qq-plot comparing the quantiles of the residuals with those from the density $f(x) = (\delta/2) \exp\{-\delta|x|\}$, with $\delta = 1/56$, are shown. From these plots it is reasonable to say that the residuals are uncorrelated and have roughly exponential tails. The value of the Portmanteau statistic $Q_Z(k)$ (see (4.1)) is smaller than the critical value $\chi_{0.95}^2(k-4)$ for lags k up to 300 (for $k = 300$ the P -value is still 0.12). The Bartlett statistic T^* (see (4.2)) has value 0.53, which is much smaller than the critical value of 1.36 at the 5 percent level. Hence, we may conclude that ARMA(3,1) is an appropriate model for the differences of **BC-pAug**. This implies, of course, that **BC-pAug** itself can be modeled by an ARIMA(3,1,1) process.

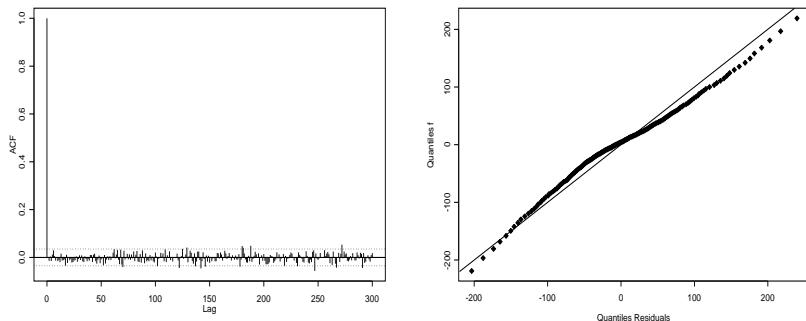


Figure 4.2: Sample autocorrelations (left) and a qq-plot (right) of the residuals from the ARIMA(3,1,1) fit to **BC-pAug**. In the qq-plot the quantiles of the residuals are compared to those of the density $f(x) = (\delta/2) \exp\{-\delta|x|\}$, with $\delta = 1/56$.

4.2.3 Analysis of other workload series

Next we consider other workload series which are studied in Leland et al. [60] and Paxson and Floyd [70] as well as more recent traces from an ATM network (Gogl [38]) and the famous Nile River data from Hurst [47]. Each of these series is believed to exhibit LRD. Below, we investigate as to whether evidence for the alternative hypothesis of non-stationarities can be found. As above, our strategy will be to show that the data can be modeled by an appropriate ARIMA process. For all series the sample autocorrelations of the (first order) differences and its absolute and squared values are small at large lags.

Bellcore data

The Bellcore data included in the Internet Traffic Archive [49] consists of four traces: **BC-pAug** and **BC-pOct** contain internal Bellcore traffic and **BC-OctExt** and **BC-OctExt4** contain external traffic, i.e. packet streams between Bellcore and the rest of the world. Each trace consists of timestamps and packet sizes of the first 1 million packet arrivals of a longer trace started in either August or October 1989. Packet sizes are between 64 and 1518 bytes, as imposed by the Ethernet protocol. For a detailed description of the traces we refer to Leland and Wilson [59], Fowler and Leland [34] and Leland et al. [60]. We created a workload series by computing either the number packets or the number of bytes that arrive per time unit, e.g. seconds. The traces of external traffic **BC-OctExt** and **BC-OctExt4** last for about 34 and 21 hours, respectively. They exhibit large peaks and **BC-OctExt** contains a clear diurnal cycle. Therefore, we do not regard these traces stationary (or ARIMA) and will not perform a detailed analysis.

Of the traces **BC-pAug** and **BC-pOct** we consider the packet and byte arrivals per second. The trace **BC-pOct** lasts for 1759 seconds. In Figure 4.3 **BC-pOct** is plotted (packets per second). There has been some dispute about the level shift that occurs between 1000 and 1200 seconds in the trace. According to Duffield et al. [25] this is a clear indication of non-stationarity. Their conclusion is based upon a comparison of the variance-time plots of the whole series and the segments 200-1000 and 1200-1759. In Table 4.2 the corresponding estimates of H can be found. As in the case of **BC-pAug** (see Table 4.1) the estimated values are smaller for the separate segments. However, Abry and Veitch [1], using their wavelet based estimator, find a consistent value of $H \approx 0.8$.

part	log-log variance-time	log-log R/S
1-1759	0.96	0.83
200-1000	0.80	0.62
1200-1759	0.64	0.52

Table 4.2: Estimates of the Hurst parameter H for two segments of **BC-pOct**.

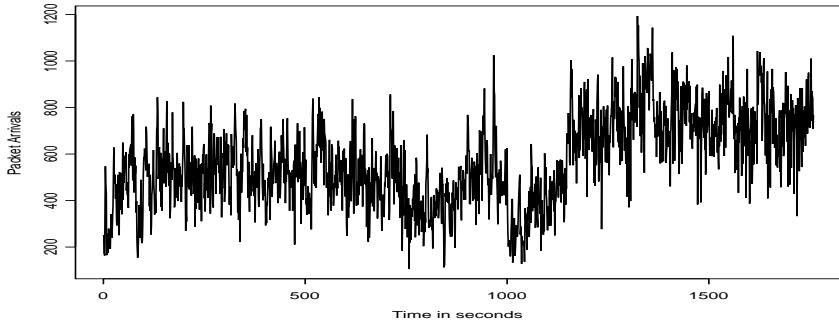


Figure 4.3: Packet arrivals per second in **BC-pOct**.

In Table 4.3 the estimated ARIMA models for the series **BC-pAug** and **BC-pOct** can be found. Since the estimated values of the parameters ϕ_j , $j \geq 2$, were smaller than 0.1 in absolute value, we have only included $\hat{\phi}_1$. Notice that the values of $\hat{\phi}_1$ and $\hat{\theta}_1$ are very similar for the four workload series. The fit is evaluated by the Portmanteau and Bartlett goodness-of-fit tests. The number of lags k included in the Portmanteau statistic $Q_Z(k)$ is chosen such that for larger lags the test fails at the 5 percent level (or k has reached $n/10$, beyond which the sample autocorrelations are less reliable). As we see, for both packets and bytes per second an appropriate ARIMA model can be found.

Wide Area traffic

In their study of WAN traffic Paxson and Floyd [70] use traces from the Digital Equipment Corporation (**DEC-PKT**) and the Lawrence Berkeley Laboratory (**LBL-PKT**), obtained in March 1995 and January 1994, respectively. Each of these traces consists of all wide area packets during one hour. Here, we consider the number of packet arrivals per second in **DEC-PKT-1** through **DEC-PKT-4** and **LBL-PKT-4** and **LBL-PKT-5**. Also, two hours' worth of TCP traffic between the Lawrence Berkeley Laboratory and the rest of the world, **LBL-TCP-3**, is included. The data can be obtained through the Internet Traffic Archive [49]. We will refer to **DEC-PKT**, **LBL-PKT** and **LBL-TCP-3** as **DEC**, **LBL** and **LBL-3**, respectively. Paxson and Floyd [70] find evidence for LRD in these series using a variance-time plot.

In Figure 4.4 the sample autocorrelations and periodogram of **DEC-1** are shown. From the two peaks in the periodogram (approximately at frequencies 0.2 and 0.4) and the periodic behavior of the autocorrelations we may conclude that periodicities are present in **DEC-1**. Hence, the assumption of stationarity becomes rather doubtful. The same holds for **DEC-j**, $j = 2, 3, 4$. From Table 4.3 we can conclude that the fit of an estimated ARIMA model and these series is

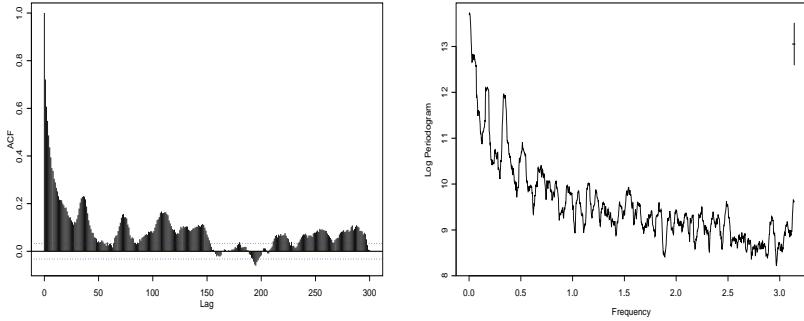


Figure 4.4: Sample autocorrelations and smoothed log-periodogram of **DEC-1**.

not very good. Especially the Portmanteau test gives poor results. In our opinion, however, Figure 4.4 shows that non-stationarity rather than LRD is present in **DEC**.

For **LBL-4** and **LBL-5** an appropriate ARIMA model can be found, although the order is higher than for the series analyzed so far. The fit of the estimated ARIMA model and **LBL-3** is less good. For $k < 200$ the Portmanteau test fails, but for $k \in [201, 402]$ the P -value is larger than 0.05. The Bartlett statistic is close to the critical value of 1.36. Considerably better results are obtained if only half of **LBL-3** is considered (which is still 3600 seconds long). In Table 4.3 the first and second half of **LBL-3** are denoted by **LBL-3(1)** and **LBL-3(2)**, respectively.

Munich University ATM data

Gogl [38] considers the traffic through an ATM link connecting the Technical University of Munich to the German Broadband Research Network. In both directions, called RX and TX, the number of passing ATM cells is measured in time units of 2 seconds. Here we consider two series of such measurements which were kindly supplied by Helmut Gogl. The series, called **ATM-RX** and **ATM-TX**, have length 43200 (corresponding to 24 hours) and contain measurements performed in March 1998. Since a clear diurnal cycle is present we restricted the analysis to the busy hours: 17:00 until 18:00 for **ATM-RX** and 12:30 until 13:30 for **ATM-TX**. From Table 4.3 we see that both series (also called **ATM-RX** and **ATM-TX**) can be modeled by an ARIMA process.

Nile River data

The last time series we consider is also the most famous one: the yearly minimal water levels of the Nile River at the Roda Gauge near Cairo for the years 622 – 1281. In 1951 Hurst [47] was the first to observe the LRD-like effect in

an R/S plot of this series. We obtained the data from Beran [4], Section 12.2. Again, Table 4.3 shows a nice fit with the estimated ARIMA model.

Most of the workload series we considered can be modeled by a non-stationary ARIMA($p,1,q$) process. This raises the question as to whether it is reasonable to assume that there is LRD in the workload measurements. A computer network is a rather complicated environment in which sending hosts compete for bandwidth (i.e. they are not independent), congestion control algorithms adapt sending rates to measured queue lengths and delays and routers try to find the fastest route for packets to travel. Also, workload characteristics may change when a new connection is set up or an active one is terminated. (See e.g. Cao et al. [16, 17] who study the influence of the number of active connections on traffic characteristics. See also Grasse et al. [39] who doubt the stationarity of VBR video traffic.) These observations support the opinion that stationarity of the workload in a computer network, even for short time periods, is rather unlikely.

trace	(p,q)	$\hat{\phi}_1$	$\hat{\phi}_2$	$\hat{\theta}_1$	$\hat{\theta}_2$	$\hat{\sigma}_Z^2$	Portmanteau	Bartlett
BC-pAug (pa)	(3, 1)	0.42	*	-0.94	-	$7.7 \cdot 10^3$	0.12 ($k = 300$)	0.53
BC-pAug (by)	(2, 1)	0.38	*	-0.94	-	$4.6 \cdot 10^9$	0.15 ($k = 75$)	0.52
BC-pOct (pa)	(4, 1)	0.41	*	-0.95	-	$12.8 \cdot 10^3$	0.08 ($k = 55$)	0.54
BC-pOct (by)	(4, 1)	0.35	*	-0.95	-	$8.3 \cdot 10^9$	0.19 ($k = 50$)	0.56
DEC-1 (pa)	(3, 1)	0.49	*	-0.93	-	$18.2 \cdot 10^3$	0.06 ($k = 12$)	0.66
DEC-2 (pa)	(3, 1)	0.46	*	-0.96	-	$21.8 \cdot 10^3$	0.11 ($k = 8$)	0.71
DEC-3 (pa)	(4, 1)	0.47	*	-0.96	-	$22.5 \cdot 10^3$	0.14 ($k = 23$)	0.48
DEC-4 (pa)	(3, 1)	0.38	*	-0.93	-	$25.7 \cdot 10^3$	0.11 ($k = 15$)	1.15
LBL-3 (pa)	(4, 1)	0.47	*	-0.94	-	$4.2 \cdot 10^3$	0.06 ($k = 402$)	1.24
LBL-3(1) (pa)	(3, 1)	0.43	*	-0.93	-	$4.3 \cdot 10^3$	0.15 ($k = 359$)	0.75
LBL-3(2) (pa)	(4, 1)	0.43	*	-0.91	-	$4.0 \cdot 10^3$	0.88 ($k = 359$)	0.73
LBL-4 (pa)	(5, 2)	-0.28	0.30	0.29	0.62	$5.0 \cdot 10^3$	0.28 ($k = 359$)	0.59
LBL-5 (pa)	(7, 2)	-0.26	0.45	0.22	0.72	$3.5 \cdot 10^3$	0.23 ($k = 359$)	0.51
ATM-RX (ce)	(2, 1)	0.34	*	-0.84	-	$3.8 \cdot 10^7$	0.55 ($k = 179$)	0.50
ATM-TX (ce)	(5, 1)	0.46	*	-0.92	-	$4.1 \cdot 10^7$	0.82 ($k = 179$)	0.39
Nile	(1, 1)	0.35	-	-0.91	-	$4.9 \cdot 10^3$	0.47 ($k = 66$)	0.56

Table 4.3: Parameter estimates and goodness-of-fit of estimated ARIMA($p,1,q$) models for various workload series; (pa)=packets per second, (by)=bytes per second, (ce)=cells per 2 seconds. For $q = 1$, the estimates $\hat{\phi}_j$, $j \geq 2$, are all smaller than 0.1 in absolute value (in this case a * represents $\hat{\phi}_2$). The value of the Bartlett test statistic T^* (see (4.2)) is given, as well as the P -value of the Portmanteau test (see (4.1)). The number of lags k included in the Portmanteau statistic $Q_Z(k)$ is chosen such that for larger lags the test fails (or k has reached $n/10$, beyond which the sample autocorrelations are less reliable). Both tests are performed at the 5 percent level, the critical value of T^* being 1.36.

4.3 Visualizing non-stationarity

Here we provide an illustration of non-stationary behavior in a computer network by means of an analysis of the sequence of packet sizes in the trace **BC-pAug**. The trace **BC-pAug** contains the first 1 million packet arrivals (about 3142.82 seconds) of the day-long trace started at 11:25 a.m., 29 August 1989, at the Bellcore Morristown Research and Engineering facility. Of each packet a timestamp and the packet size have been recorded. The data consists of only complete packets; fragments or collisions have not been recorded. About 95 percent of the packets were IP packets. The packet sizes are between 64 and 1518 bytes, as imposed by the Ethernet protocol. There are $1518 - 64 + 1 = 1455$ different packet sizes possible. Of these, 589 do not occur and 776 occur no more than 100 times. In fact, 6 packet sizes occur so often that they make up 78.4 percent of the data. This can be seen in Table 4.4 below.

packet size	frequency
64	105053
66	45496
162	203052
174	183907
1090	173802
1518	72676

Table 4.4: Packet sizes with the highest frequencies.

The smallest packets of 64 bytes are used by LAN stations to send information to other LAN stations concerning their state. The packets of 66 bytes are small TCP packets sent over the Ethernet by a TCP router, carrying information for other TCP routers. The packet size of 1090 bytes probably occurs so often because these packets traverse a network, other than the Ethernet, for which a maximum segment size is 1090 bytes. The largest packets of 1518 bytes make up files that were too long to transmit in one time. We do not know what the packets of 162 and 174 bytes represent. In Figure 4.5 a histogram of the packet sizes is depicted.

Next we consider the arrival processes of packets of a fixed size. We count the number of packet arrivals in time intervals of 4 seconds. In Figure 4.6 the number of arrivals per interval for the 6 most frequently occurring packet sizes (and 4 others) are plotted. In the plots of 162 and 174 bytes we see that a change occurs just before the 600th interval. For the 64 byte packets the number of arrivals in the first 100 intervals deviates from the rest of the arrival process. The arrivals of the 938 and 1518 bytes show a change after interval 650. Notice that the patterns in the 4 last plots are very dissimilar. This is probably due to the fact that these packet sizes are used by different processes running on the Ethernet at the time of the measurements. Between intervals 480 and 630 no packets of 1242 bytes were sent.

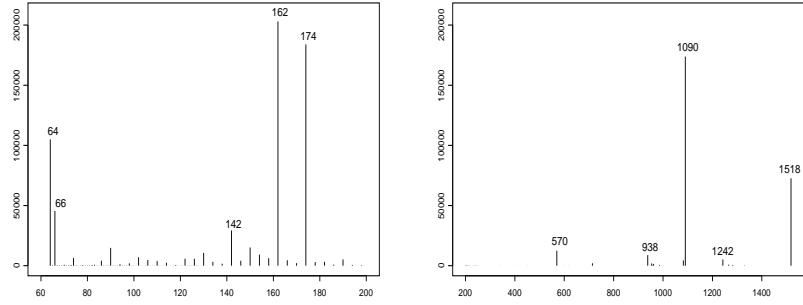


Figure 4.5: Histogram of the packet sizes (left: 64-200 bytes, right: 201-1518 bytes).

In Figure 4.6 it can be seen that the patterns in the graphs for 66 and 1090 bytes are similar. This also holds for 162 and 174 bytes and 938 and 1518 bytes. For these pairs of packet sizes scatter plots of the number of arrivals in the 4-second intervals are depicted in Figure 4.7.

Packet arrivals of 66 and 1090 bytes, 162 and 174 bytes and 938 and 1518 bytes show a strong positive correlation. There is no positive relation between the number of arrivals of 64 and 66 byte packets. We do, however, discern a slight negative correlation. This is due to the fact that these packets are competing with each other for bandwidth. Since the speed of the Ethernet LAN at Bellcore is 10 Mbps, the trace lasts for 3142.822 seconds and 434.292.031 bytes are sent during the measurements, the utilization of the network is

$$\frac{434.292.031 * 8}{10.000.000 * 3142.822} \approx 11.06.$$

The relatively low utilization of 11 percent explains that there is no strong negative relation between packet arrivals of 64 and 66 bytes. There is enough bandwidth available.

In Table 4.5 the correlation coefficients of the number arrivals of pairs of the 7 packet sizes in Figure 4.7 are computed. We made scatter plots and computed correlation coefficients of other pairs of packet sizes and discovered that the packet sizes can be divided in 5 (more or less) mutually uncorrelated groups: {64}, {66, 1090}, {162, 174}, {938, 1518} and a group consisting of the other packet sizes.

The exploratory analysis above shows that the traffic captured in **BC-pAug** is probably generated by a few processes which are roughly mutually uncorrelated. Figure 4.6 indicates that the arrival processes of the packets are not consistent with the assumption of stationarity. Therefore, stationarity of the workload series obtained from **BC-pAug** is rather doubtful.

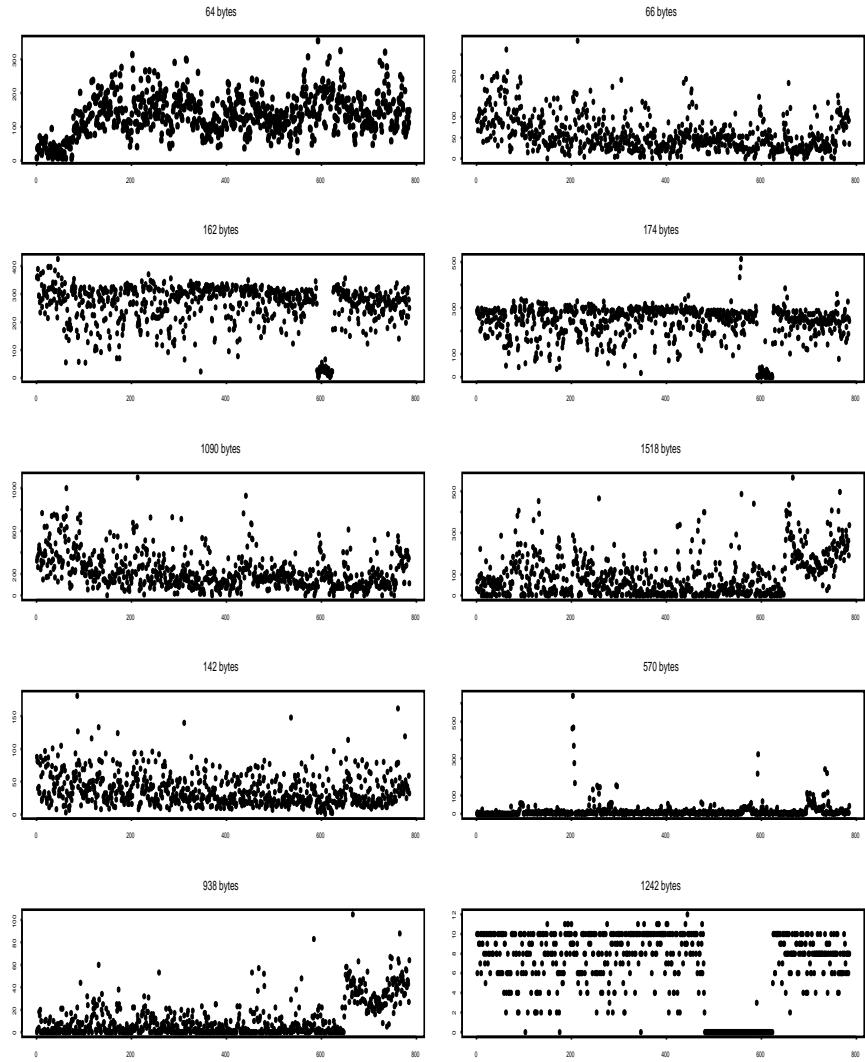


Figure 4.6: Number of packet arrivals in 4-second intervals for the most frequently occurring packet sizes.

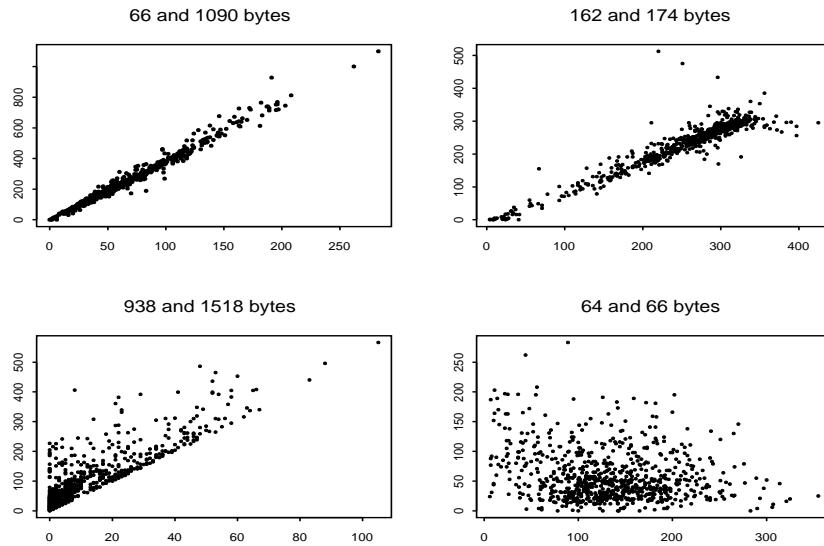


Figure 4.7: Scatter plots of packet arrivals in 4-second intervals.

	64	66	162	174	938	1090	1518
64	1.00	-0.21	-0.21	-0.18	0.07	-0.21	
66		1.00				0.99	0.09
162			1.00	0.94			-0.10
174				1.00			-0.10
938					1.00		0.88
1090						1.00	0.08
1518							1.00

Table 4.5: Significant correlation coefficients for pairs of arrival processes. The standard Pearson test was used at the 5 percent level with a two-sided alternative.

Five

Modeling Computer Network Traffic

LRD and self-similarity are observed in traffic measurements on a variety of networks, ranging from LANs and WANs to the global Internet. Computer networks have, over the past decade, undergone significant changes in their constituent traffic flows (from FTP-dominated to HTTP-dominated traffic), type of users (from professionals only to everybody online), transmission technologies (the invention of ATM technology) and scale (the Internet grows exponentially fast). Since the presence of LRD and self-similarity seems insensitive to changing network conditions, these phenomena are seen as *traffic invariants*. In this chapter we focus on underlying physical explanations for LRD- and self-similar features in workload measurements, making use of a mathematical modeling approach.

In Section 5.1 we discuss two popular models which link heavy-tailed transmission durations (and idle times) to LRD in the workload of the network. The first model, the so-called *ON/OFF model*, was introduced by Taqqu and Levy [98] and considered for modeling network traffic by Willinger et al. [106]. The ON/OFF model evaluates the workload generated by a finite number of independent ON/OFF sources. An individual ON/OFF source can be either ON or OFF, sending data at a constant rate when it is ON and remaining silent in the OFF state. The lengths of the periods that a source is ON are assumed to be realizations from a heavy-tailed distribution. Heath et al. [45] show that this yields an aggregated workload process exhibiting LRD. In the second model, network traffic consists of independent sessions whose arrivals are governed by a Poisson process. Within each session traffic is generated at unit rate. Session lengths have a heavy-tailed distribution and are independent of the arrival process. This model, introduced by Cox [19] as the immigration-death process, is a special case of the $M/G/\infty$ queuing model. In networking literature, it is known as the *infinite source Poisson model* (see e.g. Guerin et al. [41]). Cox [19] shows that the workload process generated by this model exhibits LRD. Since in the infinite source Poisson model sessions are the building blocks of network traffic it is appropriate for modeling WAN traffic. The ON/OFF model, with

its construction of explicit sources, is suitable to model traffic generation in a LAN.

In Section 5.2 we consider the cumulative workload process, both in the infinite source Poisson- and the ON/OFF model. In case of the ON/OFF model, it is shown in Willinger et al. [106] that if first the number of sources converges to infinity, followed by the time window, the properly centered and normalized cumulative workload process converges weakly to fractional Brownian motion. Hence, the cumulative workload process can be approximated by a self-similar process. Moreover, the LRD in the workload process is preserved in the increment process fractional Gaussian noise. Taqqu et al. [101] show that if the two limits of the number of sources and the time window are reversed (and a different normalization is used) stable Lévy motion is obtained as limiting process. In this case, the limit is still self-similar but has independent increments. Instead of these sequential limit regimes we also consider the case when the number of sources and the time window go to infinity simultaneously. If the number of sources grows ‘fast’ with respect to the time window, convergence of the cumulative workload to fractional Brownian motion can be shown. Alternatively, in the ‘slow’ growth regime stable Lévy motion appears as limit. Such simultaneous limit results were considered in Stegeman [93], Mikosch and Stegeman [65] and Mikosch et al. [66].

For the cumulative workload process in the infinite source Poisson model the same limit processes can be obtained. In this case, the parameters converging to infinity are the rate of the Poisson process and the time window. Taking only the latter limit, Resnick and van den Berg [80] show convergence to stable Lévy motion. Simultaneous limit regimes, yielding both fractional Brownian motion and stable Lévy motion as possible limit processes, were considered by Resnick and Rootzén and incorporated in Mikosch et al. [66].

In Section 5.3 we consider the effect on network performance of the heavy-tailed transmission durations incorporated in the ON/OFF- and infinite source Poisson models. Both theoretical queuing results as well as simulation studies are discussed.

5.1 Heavy tails as the cause of LRD

Here we define the ON/OFF model and the infinite source Poisson model in detail and discuss the physical explanation, in the form of heavy-tailed transmission (or session-) durations, they offer for the presence of LRD in workload measurements. Also, we address the question whether these rather simplistic models adequately describe real-life computer network traffic.

5.1.1 LAN traffic: the ON/OFF model

First, we consider a single ON/OFF source such as a workstation as described in Heath et al. [45]. During an ON-period, the source generates traffic at a constant rate 1, e.g. 1 byte per time unit. During an OFF-period, the source

remains silent and the input rate is 0. Let $X_{\text{on}}, X_1, X_2, \dots$ be iid non-negative random variables representing the lengths of ON-periods and $Y_{\text{off}}, Y_1, Y_2, \dots$ be iid non-negative random variables representing the lengths of OFF-periods. We also write

$$Z_i = X_i + Y_i, \quad i \geq 0.$$

The X - and Y -sequences are assumed independent. For any distribution function F we write $\bar{F} = 1 - F$ for the right tail. By $F_{\text{on}}(F_{\text{off}})$ we denote the common distribution of ON(OFF)-periods.

We assume that the distributions F_{on} and F_{off} are heavy-tailed, i.e.

$$\bar{F}_{\text{on}}(x) = x^{-\alpha_{\text{on}}} L_{\text{on}}(x) \quad \text{and} \quad \bar{F}_{\text{off}}(x) = x^{-\alpha_{\text{off}}} L_{\text{off}}(x), \quad x > 0, \quad (5.1)$$

where $\alpha_{\text{on}}, \alpha_{\text{off}} \in (1, 2)$ and $L_{\text{on}}, L_{\text{off}}$ are slowly varying at infinity. Hence, both distributions F_{on} and F_{off} have finite means μ_{on} and μ_{off} but their variances are infinite. Notice that the tail parameters α_{on} and α_{off} may be different, hence the extremes of the ON- and OFF-periods can differ significantly. If $\alpha_{\text{on}} = \alpha_{\text{off}}$, we assume that the tails \bar{F}_{on} and \bar{F}_{off} are balanced in the sense that

$$\ell = \lim_{x \rightarrow \infty} \frac{\bar{F}_{\text{on}}(x)}{\bar{F}_{\text{off}}(x)} \in (0, \infty),$$

i.e. $L_{\text{on}}(x) \sim \ell L_{\text{off}}(x)$ as $x \rightarrow \infty$.

Consider the renewal sequence generated by the alternating ON- and OFF-periods. Renewals happen at the beginnings of the ON-periods, the inter-arrival distribution is $F_{\text{on}} * F_{\text{off}}$ and the mean inter-arrival time

$$\mu = EZ_1 = \mu_{\text{on}} + \mu_{\text{off}}.$$

In order to make the renewal sequence stationary (see Resnick [78], p. 224, for a definition), a delay random variable T_0 is introduced which is independent of the X_i s and the Y_i s. A stationary version of the renewal sequence (T_n) is then given by

$$T_0, \quad T_n = T_0 + \sum_{i=1}^n Z_i, \quad n \geq 1. \quad (5.2)$$

One way to construct the delay variable T_0 (see Heath et al. [45]) is as follows. Let B , $X_{\text{on}}^{(0)}$ and $Y_{\text{off}}^{(0)}$ be independent random variables, independent of $\{Y_{\text{off}}, (X_n), (Y_n)\}$, such that B is Bernoulli with

$$P(B = 1) = \mu_{\text{on}}/\mu = 1 - P(B = 0),$$

and

$$P(X_{\text{on}}^{(0)} \leq x) = \frac{1}{\mu_{\text{on}}} \int_0^x \bar{F}_{\text{on}}(s) ds = F_{\text{on}}^{(0)}(x),$$

$$P(Y_{\text{off}}^{(0)} \leq x) = \frac{1}{\mu_{\text{off}}} \int_0^x \bar{F}_{\text{off}}(s) ds = F_{\text{off}}^{(0)}(x).$$

Define

$$T_0 = B (X_{\text{on}}^{(0)} + Y_{\text{off}}) + (1 - B) Y_{\text{off}}^{(0)}.$$

The renewal sequence (5.2) is then stationary. Define the corresponding renewal counting process

$$\xi_t = \sum_{n=0}^{\infty} 1_{[0,t]}(T_n) \quad \text{with mean } \mu_t = E\xi_t = t/\mu. \quad (5.3)$$

The *ON/OFF process of one source* is now defined as the indicator process

$$W(t) = B 1_{[0, X_{\text{on}}^{(0)}]}(t) + \sum_{n=0}^{\infty} 1_{[T_n, T_n + X_{n+1})}(t), \quad t \geq 0.$$

The ON/OFF process W is a binary process with $W(t) = 1$ if t is in an ON-period and $W(t) = 0$ if t is in an OFF-period. The stationarity of the renewal sequence (5.2) implies strict stationarity of the process W with mean

$$EW(t) = P(W(t) = 1) = \mu_{\text{on}}/\mu.$$

The precise rate of decay for the autocovariance function of the stationary process W , under the assumptions (5.1) and $\alpha_{\text{on}} \neq \alpha_{\text{off}}$ is given in Heath et al. [45]: defining

$$\alpha_{\min} = \min(\alpha_{\text{on}}, \alpha_{\text{off}}) \quad \text{and} \quad L_{\min} = \begin{cases} L_{\text{on}} & \text{if } \alpha_{\text{on}} < \alpha_{\text{off}}, \\ L_{\text{off}} & \text{if } \alpha_{\text{on}} \geq \alpha_{\text{off}}, \end{cases}$$

and

$$\mu_{\max} = \begin{cases} \mu_{\text{off}} & \text{if } \alpha_{\text{on}} < \alpha_{\text{off}}, \\ \mu_{\text{on}} & \text{if } \alpha_{\text{on}} > \alpha_{\text{off}}, \end{cases}$$

one has as $k \rightarrow \infty$,

$$\text{Cov}(W(t), W(t+k)) \sim \frac{\mu_{\max}^2}{(\alpha_{\min} - 1) \mu^3} k^{-(\alpha_{\min}-1)} L_{\min}(k). \quad (5.4)$$

Since $\alpha_{\min} \in (1, 2)$, the process W exhibits LRD in the sense that the auto-correlations are not absolutely summable (see Section 3.2). Intuitively, this can be explained as follows. Since the lengths of the ON- and OFF-periods follow a heavy-tailed distribution, they can assume extremely large values with non-negligible probability. Such an extremely large ON- or OFF-period may contain both $W(t)$ and $W(t+k)$, even if k is extremely large, yielding the non-negligible covariance in (5.4).

So far we have considered the ON/OFF process W of a single source. Willinger et al. [106] consider a network of M iid sources. Each source generates an ON/OFF process $W^{(m)}$. The total traffic in the network at time t is defined by

$$W_M(t) = \sum_{m=1}^M W^{(m)}(t), \quad t \geq 0.$$

We call W_M the *workload process*. Since the sources are iid, (5.4) yields that the workload process also exhibits LRD. Indeed, the stationary version of W_M satisfies

$$\text{Cov}(W_M(t), W_M(t+k)) = M \text{Cov}(W^{(1)}(t), W^{(1)}(t+k)).$$

The total traffic in the interval $[0, t]$ is then given by

$$CW_M(t) = \int_0^t \left(\sum_{m=1}^M W^{(m)}(u) \right) du, \quad t \geq 0. \quad (5.5)$$

We call CW_M the *cumulative workload process*.

5.1.2 The ON/OFF model and reality

The introduction of the ON/OFF model into the networking community by Willinger et al. [106] is accompanied by a detailed statistical analysis at the source level of traffic generated in the Ethernet LAN at Bellcore. As we described in Section 4.1, ON- and OFF-periods are defined for the traffic between individual source-destination pairs. It is found, using the LLCD plot and the Hill plot (see Section 3.1.2 for definitions), that the distributions of the lengths of ON- and OFF-periods are heavy-tailed with tail parameter between 1 and 2. Also, it is mentioned that no evidence is found for dependence in (or between) the sequences of ON- and OFF-periods. The other independence assumption, namely that of the M sources in the model, however, is less likely to hold in a real-life network. Indeed, sources, when simultaneously in the ON state, are usually competing for bandwidth and buffer-space if the utilization of the network is not extremely low. This was also illustrated in Section 4.3 (see Figure 4.7) where we analyzed Ethernet packet arrivals at Bellcore.

The ON/OFF model does not involve queuing or congestion control, which make it seem rather simplistic. Still, it is successful in capturing some characteristics of real-life LAN traffic. An explanation is offered by Park et al. [68]. In their simulation study, they consider a network consisting of 32 clients and 2 file servers. The clients request files from the servers. The file lengths of the files on the servers are taken from a heavy-tailed distribution. After the last byte of the requested file is received by the client, the time of the next request is taken from an exponential distribution. The simulated network also features queuing and congestion control, the latter being either TCP or UDP (see Section 2.3.3). Hence, the clients are not independent as the sources in the ON/OFF model. The workload in the network is measured at a bottleneck link between the clients and the two servers. It is found that LRD is present in the traffic measurements. Moreover, the intensity of LRD increases as the tail parameter of the heavy-tailed file size distribution is taken closer to 1, which is consistent with (5.4). Park et al. [68] focus on the influence of TCP and UDP on the observed LRD in the workload. In the UDP case there is little LRD, while TCP maintains the LRD induced by the heavy-tailed file lengths. This is explained

in [68] as follows. The reliable transmission and flow control mechanisms of TCP tend to ‘‘stretch a file in time’’, i.e. it takes longer for the transmission to be completed, while greedy unreliable UDP-based communication encourages traffic to be maximally ‘‘stretched out in space’’ by using a large portion of the available bandwidth (remember that UDP is used for applications for which speed is more important than reliability, e.g. video conferencing). Hence, the transmission of a file using TCP results in a longer ON-period than when UDP is used. In this way, TCP stretches the already heavy-tailed ON-periods and, hence, amplifies the LRD in the measured workload. Since TCP is the most frequently used transport protocol today, this explains the success of the simple ON/OFF model in describing network traffic.

The assumption that the ON/OFF sources are identical can be relaxed. Willinger et al. [106] also consider heterogeneous sources with different tail parameters of F_{on} and F_{off} . Also, a combination of sources with exponential ON- and OFF-periods and ‘heavy-tailed sources’ is possible, as well as sources with different sending rates. The assumption that the rate at which the ON/OFF source transmits data is constant is not very restrictive in the LAN context. This is confirmed by the so-called *textured plots* that are used to visualize the ON- and OFF-periods in the traffic between a source-destination pair (see Figure 2 in [106]).

5.1.3 WAN traffic: the infinite source Poisson model

First, we define the Poisson process governing the arrivals of the connections from individual sources in the network to a server. Let $(\Gamma_k, -\infty < k < \infty)$ be the points of a rate λ homogeneous Poisson process on \mathbb{R} , labeled so that $\Gamma_0 < 0 < \Gamma_1$. Hence, $\{-\Gamma_0, \Gamma_1, (\Gamma_{k+1} - \Gamma_k, k \neq 0)\}$ are iid exponentially distributed random variables with parameter λ . We imagine that the network has an infinite number of sources, and at time Γ_k a connection is made and some source begins a transmission at constant rate to the server. As a normalization, this constant rate is taken to be 1. The lengths of connections are random variables $X_k^{(\text{con})}$. We assume $X_{\text{on}}^{(\text{con})}, X_1^{(\text{con})}, X_2^{(\text{con})}, \dots$ are iid and independent of (Γ_k) and

$$P(X_{\text{on}}^{(\text{con})} > x) = \bar{F}_{\text{con}}(x) = x^{-\alpha_{\text{con}}} L_{\text{con}}(x), \quad x > 0, \quad (5.6)$$

where $\alpha_{\text{con}} \in (1, 2)$ and L_{con} is a function slowly varying at infinity. Hence, the variance of $X_{\text{on}}^{(\text{con})}$ is infinite and its mean μ_{con} is finite. Notice that

$$\nu = \sum_{k=-\infty}^{\infty} \epsilon_{(\Gamma_k, X_k^{(\text{con})})},$$

the counting function on $\mathbb{R} \times [0, \infty]$ corresponding to the points $\{(\Gamma_k, X_k^{(\text{con})})\}$, is a two dimensional Poisson process on $\mathbb{R} \times [0, \infty]$ with mean measure $\lambda \mathbb{L} \times F_{\text{con}}$ (see Resnick [77], Proposition 3.8), where \mathbb{L} denotes Lebesgue measure.

Since the sources transmit at unit rate, the workload $W^{(\text{Poi})}(t)$ in the network equals the number of active sources at time t . The workload has representation

$$W^{(\text{Poi})}(t) = \sum_{k=-\infty}^{\infty} 1_{[\Gamma_k \leq t < \Gamma_k + X_k]} = \nu(\{(s, y) \in \mathbb{R} \times (0, \infty] : s \leq t < s + y\}). \quad (5.7)$$

The second expression in (5.7) makes it clear that for each t , $W^{(\text{Poi})}(t)$ is a Poisson random variable with parameter

$$\begin{aligned} & \lambda \mathbb{L} \times F_{\text{con}}(\{(s, y) \in \mathbb{R} \times (0, \infty] : s \leq t < s + y\}) \\ &= \int_{s=-\infty}^t \int_{y=t-s}^{\infty} \lambda \mathbb{L}(ds) \times F_{\text{con}}(dy) = \lambda \int_{-\infty}^t \bar{F}_{\text{con}}(t-s) ds = \lambda \mu_{\text{con}}. \end{aligned}$$

Notice that due to the memoryless property of the exponential distribution and independence of the Poisson process and the connection lengths, the process $W^{(\text{Poi})}$ is stationary. The cumulative workload in $[0, t]$ is given by

$$CW^{(\text{Poi})}(t) = \int_0^t W^{(\text{Poi})}(s) ds. \quad (5.8)$$

Analogous to (5.4), we find that heavy-tailed connection lengths $X_k^{(\text{con})}$ induce LRD in $W^{(\text{Poi})}$. By means of a point process argument dating to Cox [19] one can show that as $k \rightarrow \infty$

$$\begin{aligned} \text{Cov}(W^{(\text{Poi})}(t), W^{(\text{Poi})}(t+k)) &= \lambda \int_k^{\infty} \bar{F}_{\text{con}}(v) dv \sim (\text{const}) k \bar{F}_{\text{con}}(k) \\ &= (\text{const}) k^{-(\alpha_{\text{con}}-1)} L_{\text{con}}(k), \end{aligned}$$

where \sim follows from Karamata's theorem (see e.g. Embrechts et al. [28], Theorem A3.6).

5.1.4 The infinite source Poisson model and reality

The construction of superimposed connections resembles traffic generation in a WAN: first a connection is set up after which data is transmitted. The assumptions of Poisson connection arrivals and heavy-tailed connection lengths are consistent with the findings of Paxson and Floyd [70], whose analysis of WAN traffic is described in Section 4.1. A more recent empirical study of WAN traffic is due to Guerin et al. [41]. They consider HTTP sessions at Boston University (1994-1995) and Berkeley (1996), traffic passing an ATM link at the Technical University of Munich (1997) and file transfers to and from a corporate WWW Ericsson server (1998). Although [41] find evidence for heavy-tailed connection lengths with tail parameter between 1 and 2, the data are inconclusive regarding the assumption of exponential inter-arrival times. As possible explanations [41] mention the difficulty of identifying Poisson time points in the data and the fact

that a Poisson process model is only expected to be a good approximation when activities by many humans are aggregated. Guerin et al. [41] do find the data consistent with the independence of the connection lengths.

In general, [41] state that the infinite source Poisson model does not adequately describe the datasets considered. They refer to the assumption of a constant transfer rate as the center of the problem. Indeed, the WAN traffic measurements analyzed by [41] show widely varying transfer rates. Moreover, it is known that two applications that contribute major portions to WAN traffic, namely FTP and HTTP, do not transmit their packets at a constant rate but in a highly bursty manner. This is mainly caused by congestion in the network and the TCP algorithm (see Section 2.3.3 for a detailed description). The assumption of a constant transfer rate is relaxed in Kurtz [57] and Konstantopoulos and Lin [55], where the cumulative workload generated by an individual connection is described by a general non-decreasing function. In this way, TCP dynamics can be taken into account by choosing a piecewise linear function, since TCP periodically adapts the transfer rate to network conditions. See also Resnick and van den Berg [80].

Finally, Resnick and Rootzén [82] consider the infinite source Poisson model in the case when the connection length $X_{\text{on}}^{(\text{con})}$ has a heavy-tailed distribution with tail parameter between 0 and 1.

5.2 Self-similar limits

In probability theory a self-similar process can be described as the limit (in the sense of the finite-dimensional distributions) of a sequence of stochastic processes (see Lamperti [58]). The role of self-similar processes is comparable to that of the Gaussian distribution in the Central Limit Theorem. In this section we show that the centered and properly normalized cumulative workload in both the ON/OFF- and infinite source Poisson model converges to a self-similar limiting process, the latter being either fractional Brownian motion or stable Lévy motion.

5.2.1 Convergence to fractional Brownian motion

Consider the cumulative workload in the ON/OFF model as defined in (5.5). We focus on the weak limit behavior of the sequence of processes $(CW_M(Tt))_{t \geq 0}$ as both the number of sources M and the time window T go to infinity. Here, we treat the case when first $M \rightarrow \infty$ and then $T \rightarrow \infty$.

Exploiting the notation introduced in Section 5.1.1, we write

$$d_{\text{on}} = \frac{\Gamma(2 - \alpha_{\text{on}})}{(\alpha_{\text{on}} - 1)} \quad \text{and} \quad d_{\text{off}} = \frac{\Gamma(2 - \alpha_{\text{off}})}{(\alpha_{\text{off}} - 1)},$$

and

$$d_{\min} = \begin{cases} d_{\text{on}} & \text{if } \alpha_{\text{on}} < \alpha_{\text{off}}, \\ d_{\text{off}} & \text{if } \alpha_{\text{off}} < \alpha_{\text{on}}. \end{cases}$$

Moreover, set

$$\sigma_0^2 = \begin{cases} \frac{2 \mu_{\max}^2 d_{\min}}{\mu^3 \Gamma(4 - \alpha_{\min})} & \text{if } \alpha_{\text{on}} \neq \alpha_{\text{off}}, \\ \frac{2 (\mu_{\text{off}}^2 d_{\text{on}} \ell + \mu_{\text{on}}^2 d_{\text{off}})}{\mu^3 \Gamma(4 - \alpha_{\min})} & \text{if } \alpha_{\text{on}} = \alpha_{\text{off}}. \end{cases}$$

Willinger et al. [106] show that if $T \rightarrow \infty$

$$\text{Var}(CW_1(T)) \sim \sigma_0^2 T^{3 - \alpha_{\min}} L_{\min}(T). \quad (5.9)$$

The main result of [106], formally stated below, is that, if first $M \rightarrow \infty$ and then $T \rightarrow \infty$, the finite-dimensional distributions of the centered and normalized process

$$V_{(M,T)}(t) = \frac{CW_M(Tt) - E[CW_M(Tt)]}{(MT^{3 - \alpha_{\min}} L_{\min}(T))^{1/2}} \quad (5.10)$$

converge to those of fractional Brownian motion. Notice that $E[CW_M(Tt)] = (\mu_{\text{on}}/\mu)MTt$.

Theorem 5.1 *Let F_{on} and F_{off} satisfy (5.1), $H = (3 - \alpha_{\min})/2$ and σ_0 be as above. Then for any $x_1, \dots, x_n \in \mathbb{R}$, $t_1, \dots, t_n \geq 0$ and $n \geq 1$*

$$\begin{aligned} & \lim_{T \rightarrow \infty} \lim_{M \rightarrow \infty} P(V_{(M,T)}(t_1) \leq x_1, \dots, V_{(M,T)}(t_n) \leq x_n) \\ &= P(\sigma_0 B_H(t_1) \leq x_1, \dots, \sigma_0 B_H(t_n) \leq x_n), \end{aligned}$$

where B_H is fractional Brownian motion as defined in Section 3.2.7. ♡

Usually, convergence in the sense of the finite-dimensional distributions involves only one limit. In this case, however, the process $V_{(M,T)}$ depends on two parameters, which both go to infinity.

Fractional Brownian motion is self-similar with parameter H . Under the conditions of the theorem, $H \in (1/2, 1)$ and so the corresponding fractional Gaussian noise sequence exhibits LRD (see Section 3.2.7). Hence, the LRD in the pre-limit workload process (see (5.4)) is preserved in the limit. Theorem 5.1 gives a probabilistic explanation of the observed self-similarity and LRD in computer network traffic provided one accepts that the limits of M and T are taken in the proposed order.

Brief sketch of the proof

The proof of Theorem 5.1 is as follows. Due to the Multivariate Central Limit Theorem and the factor $M^{1/2}$ in the normalization of (5.10) the first limit $M \rightarrow \infty$ guarantees a mean-zero Gaussian limiting process with stationary increments. Afterwards, as $T \rightarrow \infty$, the variance of the limit is stabilized by

choosing the normalization equal to the square root of the right-hand side of (5.9). Notice that (5.9) follows from

$$\text{Var}(CW_1(T)) = 2 \int_0^T \int_0^s \text{Cov}(W(0), W(u)) \, du \, ds,$$

combined with (5.4) and Karamata's theorem.

Since two sequential limits are taken it is not clear whether an extension of the result to weak convergence in the space of continuous functions $\mathbb{C}[0, \infty)$ (see e.g. Whitt [108]) is possible.

5.2.2 Convergence to stable Lévy motion

Here we see what happens when the limits in Theorem 5.1 are reversed, i.e. we consider the convergence of the finite-dimensional distributions of the cumulative workload process $(CW_M(Tt))_{t \geq 0}$, if first $T \rightarrow \infty$ and then $M \rightarrow \infty$. In this case, stable Lévy motion appears as limiting process.

Again exploiting the notation introduced in Section 5.1.1, we write

$$\sigma_{\min} = C_{\alpha_{\min}}^{-1/\alpha_{\min}},$$

where C_α is given by (3.2). If $\alpha_{\text{on}} \neq \alpha_{\text{off}}$, set

$$c = \frac{\mu_{\max}}{\mu^{1+1/\alpha_{\min}}} \quad \text{and} \quad \beta = \begin{cases} 1 & \text{if } \alpha_{\text{on}} < \alpha_{\text{off}}, \\ -1 & \text{if } \alpha_{\text{off}} < \alpha_{\text{on}}. \end{cases}$$

If $\alpha_{\text{on}} = \alpha_{\text{off}} = \alpha$, set

$$c = \frac{(\mu_{\text{off}}^\alpha \ell + \mu_{\text{on}}^\alpha)^{1/\alpha}}{\mu^{1+1/\alpha}} \quad \text{and} \quad \beta = \frac{\mu_{\text{off}}^\alpha \ell - \mu_{\text{on}}^\alpha}{\mu_{\text{off}}^\alpha \ell + \mu_{\text{on}}^\alpha}.$$

Let $F_{\min}(x) = 1 - x^{-\alpha_{\min}} L_{\min}(x)$, i.e.

$$F_{\min} = \begin{cases} F_{\text{on}} & \text{if } \alpha_{\text{on}} < \alpha_{\text{off}}, \\ F_{\text{off}} & \text{if } \alpha_{\text{off}} \leq \alpha_{\text{on}}, \end{cases}$$

and define its quantile function

$$b(x) = (1/\bar{F}_{\min}(x))^\leftarrow, \quad x > 0. \quad (5.11)$$

For a given non-decreasing function g we define the left-continuous generalized inverse of g as

$$g^\leftarrow(y) = \inf \{x : g(x) \geq y\}.$$

Notice that $b(x) = x^{1/\alpha_{\min}} \tilde{L}_{\min}(x)$, where \tilde{L}_{\min} is a slowly varying function satisfying

$$\lim_{x \rightarrow \infty} \frac{L_{\min}(x^{1/\alpha_{\min}} \tilde{L}_{\min}(x))}{\tilde{L}_{\min}^{\alpha_{\min}}(x)} = 1,$$

(see Ibragimov and Linnik [48], relation (2.6.4)).

The following theorem, due to Taqqu et al. [101], states that if first $T \rightarrow \infty$, the finite-dimensional distributions of the process

$$\tilde{V}_{(M,T)}(t) = \frac{CW_M(Tt) - E[CW_M(Tt)]}{M^{1/\alpha_{\min}} b(T)}$$

converge to those of stable Lévy motion.

Theorem 5.2 *Let F_{on} and F_{off} satisfy (5.1) and c , σ_{\min} and β be as above. Then for any $x_1, \dots, x_n \in \mathbb{R}$, $t_1, \dots, t_n \geq 0$ and $n \geq 1$*

$$\begin{aligned} \lim_{M \rightarrow \infty} \lim_{T \rightarrow \infty} P(\tilde{V}_{(M,T)}(t_1) \leq x_1, \dots, \tilde{V}_{(M,T)}(t_n) \leq x_n) \\ = P(c \Lambda_{\alpha_{\min}, \sigma_{\min}, \beta}(t_1) \leq x_1, \dots, c \Lambda_{\alpha_{\min}, \sigma_{\min}, \beta}(t_n) \leq x_n), \end{aligned}$$

where $\Lambda_{\alpha_{\min}, \sigma_{\min}, \beta}$ is stable Lévy motion as defined in Section 3.1.3. \heartsuit

Comparing this result with Theorem 5.1 one may conclude the following: if the limits are reversed and a different normalization is used, another self-similar process, infinite variance stable Lévy motion, appears as limit process. In contrast to the fractional Brownian motion of Theorem 5.1, the LRD is completely lost: the limit process has independent increments.

A related result is obtained by Levy and Taqqu [61], who consider a renewal reward process for which both the inter-renewal distribution and the distribution of the rewards are heavy-tailed. Convergence of the total reward in $[0, Tt]$ to a stable self-similar process is shown as $T \rightarrow \infty$. See also Pipiras and Taqqu [73].

Brief sketch of the proof

The process $\tilde{V}_{(M,T)}$ can be decomposed into

$$\tilde{V}_{(M,T)}(t) = [M^{1/\alpha_{\min}} b(T)]^{-1} \sum_{m=1}^M \sum_{k=1}^{\xi_{Tt}^{(m)}} J_k^{(m)} + o_P(1), \quad (5.12)$$

where ξ is the counting process defined by (5.3), J_k is the centered contribution of the renewal interval $[T_{k-1}, T_k]$, i.e.

$$J_k = X_k - P(W=1) Z_k = (\mu_{\text{off}}/\mu)(X_k - \mu_{\text{on}}) - (\mu_{\text{on}}/\mu)(Y_k - \mu_{\text{off}}),$$

and $o_P(1)$ converges to zero in probability as $T \rightarrow \infty$. The term $o_P(1)$ contains the contributions of the intervals $[0, T_0]$ and $[T_{\xi_{Tt}-1}, Tt]$ and the difference between the centering used in $\tilde{V}_{(M,T)}$ and the random sum in (5.12). An application of Theorem V2.2 in Gut [43] yields that, for $T \rightarrow \infty$ and $M \geq 1$ fixed, the random sum in (5.12) converges weakly to the sum of M iid α_{\min} -stable Lévy motions. The weak convergence takes place in the space $\mathbb{D}[0, \infty)$ of càdlàg

functions, equipped with the J_1 -metric (see e.g. Pollard [75]). Since the sum of M iid mean-zero α -stable random variables $\tilde{X}_1, \dots, \tilde{X}_M$ has the same distribution as $M^{1/\alpha} \tilde{X}_1$ for $1 < \alpha \leq 2$ (see Samorodnitsky and Taqqu [89], Corollary 1.2.9), the limit $M \rightarrow \infty$ is trivial and Theorem 5.2 holds for any fixed $M \geq 1$.

The convergence in Theorem 5.2 cannot be extended to weak convergence in the J_1 -topology. This is proved by Konstantopoulos and Lin [55]. They consider a sequence of random elements with continuous paths, converging in the sense of finite-dimensional distributions, to a random element in $\mathbb{D}[0, 1]$. If the limit has discontinuous paths with positive probability, then there is no weak convergence under the J_1 -topology. Since the process $V_{(M, T)}$ has continuous paths and stable Lévy motion has discontinuous paths a.s. the result of [55] applies. An alternative proof of the impossibility of J_1 -convergence, using an extreme value argument, is given in Stegeman [92]. At this moment it is not clear whether weak convergence holds in a weaker topology on $\mathbb{D}[0, \infty)$, e.g. the M_1 -topology (see Skorohod [91]).

Finite variance case: $\alpha_{\min} = 2$

The main reason for the different results of Theorems 5.1 and 5.2 is the infinite variance of the lengths of the ON- and OFF-periods. Indeed, suppose the variances of both X and Y are finite. Then Willinger et al. [106] show that (with $\alpha_{\min} = 2$) Theorem 5.1 yields fractional Brownian motion with $H = 1/2$ as limit process. Clearly, this is Brownian motion.

In the proof of Theorem 5.2 one can use a functional limit theorem (Theorem V2.2 in Gut [43]), which states that a random sum of iid heavy-tailed random variables converges weakly to stable Lévy motion. If the random variables have finite variance, Theorem V2.1 in Gut [43] shows that the limit process will be Brownian motion.

Brownian motion is self-similar and has independent increments. In this case the pre-limit workload process has short-range dependence in the sense that the corresponding autocorrelations are summable.

Equivalent result for the infinite source Poisson model

Convergence of the cumulative workload process $CW^{(\text{Poi})}(Tt)$ to stable Lévy motion is shown by Resnick and van den Berg [80]. Below we state their result. Define the quantile function of F_{con} by

$$b_{\text{con}}(x) = (1/\bar{F}_{\text{con}}(x))^\leftarrow, \quad x > 0, \quad (5.13)$$

and set

$$\sigma_{\text{con}} = C_{\alpha_{\text{con}}}^{-1/\alpha_{\text{con}}},$$

where C_α is given by (3.2).

Theorem 5.3 *Let F_{con} satisfy (5.6). Then, as $T \rightarrow \infty$,*

$$\frac{CW^{(\text{Poi})}(Tt) - E[CW^{(\text{Poi})}(Tt)]}{\lambda^{1/\alpha_{\text{con}}} b_{\text{con}}(T)} \xrightarrow{M_1} \Lambda_{\alpha_{\text{con}}, \sigma_{\text{con}}, 1}(t),$$

where $\Lambda_{\alpha_{\text{con}}, \sigma_{\text{con}}, 1}$ is stable Lévy motion as defined in Section 3.1.3 and $\xrightarrow{M_1}$ denotes weak convergence in $\mathbb{D}[0, \infty)$ equipped with the M_1 -topology. \heartsuit

Notice that $E[CW^{(\text{Poi})}(Tt)] = \lambda \mu_{\text{con}} Tt$. In the case $\alpha_{\text{on}} < \alpha_{\text{off}}$, the result is analogous to Theorem 5.2. Theorem 5.3 holds for any fixed $\lambda > 0$; the limit $\lambda \rightarrow \infty$ is omitted. The only difference with Theorem 5.2 is that now the convergence of the finite-dimensional distributions is strengthened to M_1 -convergence. The M_1 -topology is weaker than the J_1 -topology (see Skorohod [91]). As in Theorem 5.2, J_1 -convergence is impossible. Resnick and van den Berg [80] prove Theorem 5.3 for the more general case in which the cumulative workload generated by an individual session is given by a regularly varying function ζ . A unit transmission rate corresponds to $\zeta(t) = t$.

5.2.3 Simultaneous limit regimes

There is no particular (practical or theoretical) reason why we should prefer the limit regime in Theorem 5.1 to that in Theorem 5.2. In order to better understand the interplay of the roles of the limits of M and T to infinity, we study the weak limit behavior of the cumulative workload process in the ON/OFF model for simultaneous limit regimes in which M and T go to infinity at the same time. In particular, we assume that $M = M(T)$ is some integer-valued function such that

$$M(T) \text{ is non-decreasing in } T \text{ and } \lim_{T \rightarrow \infty} M(T) = \infty.$$

For ease of presentation we usually suppress the dependence of M on T . For example, we write $W_M = W_{M(T)}$ and $CW_M = CW_{M(T)}$ for the workload- and cumulative workload process, respectively.

We will show that if M exhibits ‘slow growth’, the finite-dimensional distributions of $(CW_M(Tt))_{t \geq 0}$ have an infinite variance stable Lévy motion as limit. If M grows ‘fast’ the limiting process will be fractional Brownian motion. The following conditions express ‘slow’ and ‘fast’ growth of M in terms of the quantile function b defined in (5.11).

$$\text{Slow Growth Condition 1: } \lim_{T \rightarrow \infty} \frac{b(MT)}{T} = 0,$$

$$\text{Fast Growth Condition 2: } \lim_{T \rightarrow \infty} \frac{b(MT)}{T} = \infty.$$

The next lemma provides alternate ways to express the conditions and can be found in Mikosch et al. [66].

Lemma 5.4

1. *The Slow Growth Condition 1 is equivalent to*

$$\lim_{T \rightarrow \infty} MT \bar{F}_{\min}(T) = 0 \quad \text{or} \quad \lim_{T \rightarrow \infty} \text{Cov}(W_M(0), W_M(T)) = 0.$$

2. The Fast Growth Condition 2 is equivalent to

$$\lim_{T \rightarrow \infty} MT \overline{F}_{\min}(T) = \infty \quad \text{or} \quad \lim_{T \rightarrow \infty} \text{Cov}(W_M(0), W_M(T)) = \infty.$$

Proof. In the case of Condition 1, there exists a function $0 < \epsilon_T \rightarrow 0$ such that $T\epsilon_T \rightarrow \infty$ and $b(MT) = T\epsilon_T$. Thus

$$MT \sim 1/\overline{F}_{\min}(T\epsilon_T).$$

Therefore, Condition 1 implies

$$MT \overline{F}_{\min}(T) \sim \overline{F}_{\min}(T)/\overline{F}_{\min}(T\epsilon_T) \rightarrow 0. \quad (5.14)$$

Conversely, if $\delta_T = MT\overline{F}_{\min}(T) \rightarrow 0$, then using $b^{\leftarrow}(T) \sim 1/\overline{F}_{\min}(T)$, we get

$$\frac{b(MT)}{T} \sim \frac{b(\delta_T b^{\leftarrow}(T))}{b(b^{\leftarrow}(T))} \rightarrow 0,$$

and so Condition 1 and (5.14) are equivalent. The proof for Condition 2 is similar. The equivalence in terms of covariances follows from

$$\text{Cov}(W_M(0), W_M(T)) = MC\text{Cov}(W(0), W(T)) \sim (\text{const})MT \overline{F}_{\min}(T),$$

where we used (5.4) in the last step. \heartsuit

From Lemma 5.4 it follows that $M \sim (\text{const})T^{\alpha_{\min}-1}L_{\min}^{-1}(T)$ is the critical growth rate in between the ‘slow’ and ‘fast’ growth situations.

The following theorem gives the limiting process of the cumulative workload, depending on whether Condition 1 or 2 holds. Define

$$V_T(t) = \frac{CW_M(Tt) - E[CW_M(Tt)]}{(MT^{3-\alpha_{\min}}L_{\min}(T))^{1/2}},$$

and

$$\tilde{V}_T(t) = \frac{CW_M(Tt) - E[CW_M(Tt)]}{b(MT)}.$$

Theorem 5.5 Let F_{on} and F_{off} satisfy (5.1). Let σ_0 be as in Theorem 5.1 and $H = (3 - \alpha_{\min})/2$. Let c , σ_{\min} and β be as in Theorem 5.2.

1. If Slow Growth Condition 1 holds, then as $T \rightarrow \infty$

$$\tilde{V}_T(t) \xrightarrow{fidi} c \Lambda_{\alpha_{\min}, \sigma_{\min}, \beta}(t),$$

where \xrightarrow{fidi} denotes convergence of the finite-dimensional distributions and $\Lambda_{\alpha_{\min}, \sigma_{\min}, \beta}$ is stable Lévy motion as defined in Section 3.1.3.

2. If Fast Growth Condition 2 holds, then as $T \rightarrow \infty$

$$V_T(t) \xrightarrow{d} \sigma_0 B_H(t),$$

where \xrightarrow{d} denotes weak convergence in the space $\mathbb{C}[0, \infty)$ equipped with the uniform topology and B_H is fractional Brownian motion as defined in Section 3.2.7.

♡

As in Theorem 5.2 J_1 -convergence to stable Lévy motion is impossible. It is not clear whether M_1 -convergence holds. In the case when F_{on} and F_{off} have finite variances Theorem 5.5 holds with $\alpha_{\min} = 2$ and yields Brownian motion as limiting process for both limit regimes. The proof of Theorem 5.5 is presented in Sections 5.A and 5.B.

It seems surprising that although a nearly continuous transition between the ‘slow’ and ‘fast’ growth cases exists, the dependence structures of the limiting processes are completely different. Under Slow Growth Condition 1 the limiting increment sequence consists of iid random variables, while under Fast Growth Condition 2 it has LRD. Intuitively, this can be explained as follows. The effect of $T \rightarrow \infty$ is that the time scale is blown up. Values of t that were close to each other are far apart as T becomes large. Keeping M fixed, this destroys the dependence structure within an individual ON/OFF process (as in Theorem 5.2). Moreover, Lemma 5.4 shows that also under Slow Growth Condition 1 the covariances of the workload process converge to zero. Therefore, Fast Growth Condition 2 is necessary to preserve the dependence structure of the workload process. In fact, Lemma 5.4 tells us that in this case the covariances blow up to infinity. Hence, a stabilizing normalization is needed, which according to (5.9) should be $(MT^{3-\alpha_{\min}} L_{\min}(T))^{1/2}$.

A result related to Theorem 5.5 is due to Pipiras et al. [74], who consider M iid renewal reward processes with a heavy-tailed inter-renewal distribution (tail index $\alpha \in (1, 2)$) and a heavy-tailed reward distribution (tail index $\beta \in (1, 2)$). They investigate the weak limit behavior of the total reward in $[0, Tt]$ for $M = M(T)$ as $T \rightarrow \infty$, using the same growth conditions as in Theorem 5.5. If $\alpha < \beta$, then Slow Growth Condition 1 yields the α -stable Lévy motion from Theorem 5.5 as limit, while Fast Growth Condition 2 results in a self-similar β -stable process with dependent increments. For $\beta < \alpha$, β -stable Lévy motion is obtained as limit, regardless of the growth rate of M . If the rewards have finite variance and Fast Growth Condition 2 holds, the limit is the fractional Brownian motion from Theorem 5.5.

Equivalent result for the infinite source Poisson model

A result analogous to Theorem 5.5 for the cumulative workload $CW^{(\text{Poi})}$ in the infinite source Poisson model can be found in Mikosch et al. [66]. The role of M is now played by the Poisson parameter $\lambda = \lambda(T)$. Recall the definition of the quantile function b_{con} from (5.13) and set

$$\tilde{\sigma}_{\text{con}} = \frac{1}{3 - \alpha_{\text{con}}} \left[\frac{\alpha_{\text{con}}}{2 - \alpha_{\text{con}}} + \frac{2}{\mu_{\text{con}}} \right].$$

Theorem 5.6 *Let F_{con} satisfy (5.6), $H = (3 - \alpha_{\text{con}})/2$ and $\tilde{\sigma}_{\text{con}}$ be as above.*

1. *If $b_{\text{con}}(\lambda T) = o(T)$, then as $T \rightarrow \infty$*

$$\frac{CW^{(\text{Poi})}(Tt) - E[CW^{(\text{Poi})}(Tt)]}{b_{\text{con}}(\lambda T)} \xrightarrow{\text{fidi}} \Lambda_{\alpha_{\text{con}}, 1, 1}(t),$$

where $\xrightarrow{\text{fidi}}$ denotes convergence of the finite-dimensional distributions and $\Lambda_{\alpha_{\text{con}}, 1, 1}$ is stable Lévy motion as defined in Section 3.1.3.

2. *If $T = o(b_{\text{con}}(\lambda T))$, then as $T \rightarrow \infty$*

$$\frac{CW^{(\text{Poi})}(Tt) - E[CW^{(\text{Poi})}(Tt)]}{(\lambda T^{3 - \alpha_{\text{con}}} L_{\text{con}}(T))^{1/2}} \xrightarrow{d} \tilde{\sigma}_{\text{con}} B_H(t),$$

where \xrightarrow{d} denotes weak convergence in the space $\mathbb{C}[0, \infty)$ equipped with the uniform topology and B_H is fractional Brownian motion as defined in Section 3.2.7.

♡

Since the conditions on λ are the same as Conditions 1 and 2 on M , Lemma 5.4 holds with λ and b_{con} replacing M and b , respectively. The remarks following Theorem 5.5 also apply here.

Practical relevance

Considering a real-life computer network, what is the practical relevance of the simultaneous limits in Theorem 5.5? Although the parameters M and T can be interpreted as the number of connections or hosts in the network and the time window, respectively, it is not clear which ratios M/T correspond to the Slow Growth Condition and which to the Fast Growth Condition. However, it is generally agreed upon that workload measurements contain LRD and certainly do not consist of independent observations. This would indicate that “in practice the fast growth regime applies”. This is confirmed by Cao et al. [16, 17] who study the effect of the number of active connections on traffic characteristics. Analyzing a large number of Ethernet and ATM traffic traces, they observe

that the variability of the aggregated workload becomes smaller as the number of active connections increases. Moreover, the LRD (present in all workload series in [16, 17]) is preserved. Another issue is the question of Gaussian versus stable marginal distributions. Guerin et al. [41], analyzing several Wide Area traffic traces (see also Section 5.1.4), do not find conclusive evidence for one or the other. Heyman [46], however, observes that the effect of multiplexing a large number of TCP sources is that the workload smooths out and its marginals tend towards a Gaussian distribution.

An anonymous referee of the paper [66] made two important remarks about the relation between the two limit regimes and actual networking practice. First, most measured traffic traces have been collected from relatively low bandwidth links. This limits the range over which tail behavior can be observed, i.e. extremely large workloads are not observed due to bandwidth limitations. Second, the vast majority of measured WAN traffic consists of TCP/IP packets. From Section 2.3.3 we know that TCP does not allow individual connections to grab the lion's share of the available bandwidth on a given link, but aims at giving each connection its fair share. Even for high-speed links of 100 Mbps (=Megabits per second) or more, TCP imposes a maximum window size that drastically limits the rates at which the different connections can send their packets.

From the comments of the referee we may conclude that in the workload measurements a regime of small M and large T , i.e. 'slow growth', is not likely, either due to a relatively low bandwidth of the considered link or the nature of TCP. This is consistent with the observed LRD and Gaussianity in network traffic, which are properties of the limiting fractional Brownian motion under 'fast growth'. The tendency of TCP to amplify the LRD in an inputted traffic stream is also observed in the simulation study by Park et al. [68] (see Section 5.1.2). Networking arguments and modeling efforts strengthening this hypothesis can be found in Heyman [46], Guo et al. [42], Veres and Boda [104] and Veres et al. [105].

The relevance of Theorem 5.5 for networking practice is not limited to today's networks. In the future, available link bandwidth will range from Mbps to Gbps (=Gigabits per second) and beyond. Moreover, we can expect protocols that, in contrast to today's TCP, will allow individual connections to grab a highly variable amount of the available bandwidth almost instantaneously. Hence, for future network traffic a regime of large M and small T belongs to the possibilities. In this sense, Theorem 5.5 predicts the dynamics of future network traffic depending on the developments in protocol design and networking technology.

5.3 Network performance – a discussion

The heavy-tailed transmission durations incorporated in the ON/OFF- and infinite source Poisson models are responsible for a significant decrease in network performance. This conclusion has been drawn in a number of theoretical and

simulation studies we will discuss here.

5.3.1 Queuing results

Here we consider a buffer fed by the workload generated by M ON/OFF sources in the ON/OFF model as defined in Section 5.1.1. The number of sources M is taken to be fixed and instead of (5.1) we assume

$$\overline{F}_{\text{on}}(x) = x^{-\alpha} L(x) \quad \text{and} \quad \overline{F}_{\text{off}}(x) = o(\overline{F}_{\text{on}}(x)), \quad (5.15)$$

where $\alpha > 1$ and L is slowly varying at infinity. Hence, \overline{F}_{on} is regularly varying and F_{off} has a lighter tail than F_{on} . The buffer has infinite capacity and the *outflow rate* is rM , where r satisfies

$$\frac{\mu_{\text{on}}}{\mu} < r < 1. \quad (5.16)$$

It is clear that if $r \geq 1$ the buffer would always be empty. The left inequality in (5.16) is a stability condition guaranteeing a ‘tendency’ towards an empty buffer. This will be made more precise later on. We denote the buffer content at time $t \geq 0$ by $Q(t)$. Following Asmussen [2], Section III.8, we write

$$Q(t) = \max \left\{ Q(0) + CW_M[0, t] - rMt, \sup_{0 \leq s \leq t} (CW_M[s, t] - rM(t-s)) \right\},$$

where $CW_M[0, t] = CW_M(t)$ is the cumulative workload in $[0, t]$ and $Q(0)$ is an arbitrary initial state. Since the ON/OFF processes are stationary, we have

$$\sup_{0 \leq s \leq t} (CW_M[s, t] - rM(t-s)) \stackrel{d}{=} \sup_{0 \leq s \leq t} (CW_M[0, s] - rMs).$$

Since $E(CW_M[0, t]) = Mt\mu_{\text{on}}/\mu$, the law of large numbers and (5.16) imply that as $t \rightarrow \infty$

$$CW_M[0, t] - rMt \rightarrow -\infty \quad \text{a.s.} \quad (5.17)$$

which means that there is a negative drift in the buffer content process $Q(t)$. Next, define

$$Q = \sup_{t \geq 0} (CW_M[0, t] - rMt). \quad (5.18)$$

Combining (5.17) and Asmussen [2], Proposition III.8.2, we obtain

$$P(Q < \infty) = 1 \quad \text{and} \quad Q(t) \xrightarrow{d} Q, \quad \text{as } t \rightarrow \infty.$$

We are interested in the rate of decay of the probability tail

$$P(Q > u), \quad \text{as } u \rightarrow \infty.$$

The following result holds.

Proposition 5.7 Suppose $M = 1$, F_{on} and F_{off} satisfy (5.15) and r satisfies (5.16). Then, as $u \rightarrow \infty$

$$P(Q > u) \sim (\text{const}) \int_{u/(1-r)}^{\infty} \overline{F}_{\text{on}}(s) \, ds \sim (\text{const}) u^{-(\alpha-1)} L(u). \quad (5.19)$$

♡

Hence, for $\alpha \in (1, 3)$ the distribution of the buffer content is heavy-tailed. Notice that the second equivalence in (5.19) is just Karamata's theorem. A proof of the first equivalence in (5.19) involving queuing theory can be found in Jelenković and Lazar [51]. An alternative proof can be given by using a result by Embrechts and Veraverbeke [27] concerning the supremum of heavy-tailed random walks.

For $M \geq 1$ the exact rate of $P(Q > u)$ is given in Zwart et al. [109]. Their result is as follows.

Proposition 5.8 Suppose F_{on} and F_{off} satisfy (5.15) and r satisfies (5.16). Set

$$n = \left[M \frac{r - \mu_{\text{on}}/\mu}{1 - \mu_{\text{on}}/\mu} \right] + 1, \quad (5.20)$$

where $[x]$ denotes the integer part of x . Assume that

$$M \left(\frac{r - \mu_{\text{on}}/\mu}{1 - \mu_{\text{on}}/\mu} \right) \text{ is not an integer.} \quad (5.21)$$

Then, for a certain constant $K > 0$, as $u \rightarrow \infty$

$$P(Q > u) \sim \binom{M}{n} \frac{K}{\mu^n} \left(\int_{x_u}^{\infty} \overline{F}_{\text{on}}(s) \, ds \right)^n, \quad (5.22)$$

where

$$x_u = \frac{u}{n + (M - n)(\mu_{\text{on}}/\mu) - rM}.$$

♡

From (5.22) we see that $P(Q > u)$ is regularly varying with index $n(\alpha - 1)$. Comparing (5.19) and (5.22) we see that for M sources $P(Q > u)$ is of the same order as the product of n probability tails in the case $M = 1$. It seems that an extremely high buffer content is the result of n of the M sources being extremely active simultaneously. The factor $\binom{M}{n}$ gives the total number of n -subsets of $\{1, \dots, M\}$. Intuitively, this can be explained as follows. Suppose one source has an extremely long ON-period and the other $M - 1$ sources have average activity. During this long ON-period the average net-inflow into the buffer is

$$a_1 = 1 + (M - 1) \frac{\mu_{\text{on}}}{\mu} - rM.$$

If $a_1 \leq 0$ the high activity of a single source cannot cause an increase in the buffer content. For $P(Q > u)$ this event is negligible. If, however, $a_1 > 0$ one extremely active source can be responsible for an extremely high buffer content. If j sources are extremely active simultaneously over a period of time, the average net-inflow is

$$a_j = j + (M - j) \frac{\mu_{\text{on}}}{\mu} - rM.$$

Since $r < 1$ we have $a_M > 0$. Following the principle “*the most unlikely events happen in the most likely way*,” we assume that an extremely high buffer content is caused by simultaneous extreme activity of n sources, where

$$n = \min\{j \in \{1, \dots, M\} : a_j > 0\}. \quad (5.23)$$

This n is the same as in (5.20). Condition (5.21) implies that $a_j = 0$ cannot occur. If $n \geq 2$ then $P(Q > u)$ decays faster than in the case $M = 1$, i.e. multiplexing improves the performance of the queuing system. These multiplexing gains may even be larger if some sources have exponentially bounded tails of their distribution of ON-period lengths. In this case $P(Q > u)$ may decrease exponentially, even though there are sources with heavy-tailed ON-period distributions (see e.g. Boxma [12], Boxma and Dumas [13], Jelenković and Lazar [51], Dumas and Simonian [26] and Zwart et al. [109]). When the ON-period distributions of all sources have exponentially bounded tails $P(Q > u)$ decreases at an exponential rate. From a queuing perspective, this indicates the size of the effect of heavy-tailed ON-periods on network performance.

Queuing results concerning the infinite source Poisson model can be found in Zwart [110], Chapter 8.

5.3.2 Simulation studies

In Section 5.1.2 we discussed the simulation study by Park et al. [68] in which a client-server network is considered with 32 clients and 2 file servers. The file lengths are taken from a heavy-tailed distribution. The workload is measured at a bottleneck link between the clients and the servers. It is found that LRD is present in the traffic measurements. Moreover, the intensity of LRD increases as the tail parameter α of the heavy-tailed file size distribution is taken closer to 1, which is consistent with formula (5.4) expressing the LRD in the workload described by the ON/OFF model. Park et al. [68, 69] also consider several performance measures of the link: throughput (in bytes per second), packet loss rate (in percents), buffer utilization (in bytes) and average queue length (in bytes). The effect of changes in the tail parameter α , bandwidth of the link, buffer size and choice of transport mechanism (TCP or UDP, see Section 2.3.3) on the performance of the network is studied. It is found that a decrease in α , i.e. a ‘heavier’ file size distribution, results in deteriorating network performance, while more bandwidth or buffer size increases network performance. These effects are even bigger when instead of the reliable transport mechanism TCP the unreliable UDP is used.

The effect of LRD on network performance is also studied by Erramilli et al. [29]. They consider a single server queuing system with an infinite buffer and deterministic service times. The packet arrivals are governed by a sequence of inter-arrival times taken from real-life Ethernet traffic measurements. Erramilli et al. [29] construct several *shuffled* versions of the original Ethernet input trace as follows. They divide the sequence of inter-arrival times in blocks of size m and consider the *internally shuffled* trace in which the values within each block are shuffled but the order of the blocks does not change, and the *externally shuffled* trace in which the blocks themselves are shuffled but the order of the values within each block is unaffected. For the original and shuffled traces the average delay in seconds is plotted against the level of utilization of the queue. It appears that for m relatively small the plot of the internally shuffled trace is close to the original, while the average delay is significantly less using the externally shuffled trace. Hence, the destruction of the correlation structure up to lag m in the sequence of inter-arrival times does not affect the performance of the queuing system. Moreover, correlations at lags larger than m are responsible for the delay characteristics of the queue.

Although network performance is decreased by LRD, it must be remembered that LRD is inherently an asymptotic notion: autocorrelations at large lags are non-negligible. In practice, however, buffers have finite capacity and congestion control algorithms consider only finite time horizons. Ryu and Elwalid [88] conclude on the basis of a simulated queuing system with a finite buffer that short-term correlations dominate long-term correlations in their influence on the packet loss rate, even in the presence of LRD in the input stream. In this sense, small buffer capacity combined with large bandwidth delimits the scope of influence of LRD on network performance. See also Grossglauser and Bolot [40].

A different approach in dealing with the presence of LRD is to explicitly take into account larger time scales for detecting persistent shifts in the overall network contention. This so-called *multiple time scale congestion control* (MTSC) framework is introduced by Tuan and Park [103]. In a simulated client-server network (similar to the one considered by Park et al. [68, 69]) the application of MTSC results in significant throughput gains. These gains increase with the intensity of LRD in network traffic.

Appendix

5.A Proof of Theorem 5.5: Slow Growth

5.A.1 The basic decomposition

We will use the following decomposition of the cumulative workload:

$$CW_M(T) = I(T) + II(T) + III(T), \quad (\text{A.1})$$

where

$$\begin{aligned} I(T) &= \sum_{m=1}^M B^{(m)} \min(T, (X_{\text{on}}^{(0)})^{(m)}) , \\ II(T) &= \sum_{m=1}^M \sum_{k=1}^{\xi_T^{(m)}} X_k^{(m)} , \\ III(T) &= - \sum_{m=1}^M \max(0, T_{\xi_T^{(m)}-1}^{(m)} + X_{\xi_T^{(m)}}^{(m)} - T) \mathbf{1}_{[\xi_T^{(m)} \geq 1]} . \end{aligned}$$

Here, $I(T)$ is the contribution of the 0th renewal intervals $[0, T_0^{(m)}]$ while $II(T)$ captures $[T_{k-1}^{(m)}, T_k^{(m)}]$ for $k = 1, \dots, \xi_T^{(m)}$. The term $III(T)$ is needed to delete the contribution of $[T, T_{\xi_T^{(m)}}^{(m)}]$.

The basic idea of the proof is first to show that the terms $I(T)$ and $III(T)$ are asymptotically negligible. This is done in Section 5.A.2. The next step consists of replacing the counting processes $\xi_T^{(m)}$ in $II(T)$ simultaneously by their identical means μ_T . After the replacement, the resulting process is a sum of iid random variables and so classical limit theory for sums of iid random variables comes in. The replacement described above is provided by a large deviation result given in Section 5.C. In Section 5.D we present a bound for regularly varying functions which is frequently used in the proof of Theorem 5.5.

5.A.2 Vanishing remainder terms

Before dealing with $I(T)$, we need the following lemma.

Lemma 5.A.1 *If Slow Growth Condition 1 holds, then*

$$\lim_{T \rightarrow \infty} \frac{MT^2 \bar{F}_{\min}(T)}{b(MT)} = 0, \quad (\text{A.2})$$

and if Fast Growth Condition 2 holds, this limit is infinite.

Proof. Assume that Condition 1 holds. Set $\epsilon_T = b(MT)/T \rightarrow 0$ so that $\epsilon_T T \rightarrow \infty$. Denoting the ratio in (A.2) by r_T , we see that

$$r_T \sim \frac{\bar{F}_{\min}(T)}{\epsilon_T \bar{F}_{\min}(T\epsilon_T)},$$

and using the Karamata representation of a regularly varying function (see Bingham et al. [9]), we obtain

$$r_T \sim [\epsilon_T]^{-1} \exp \left\{ - \int_{T\epsilon_T}^T u^{-1} \alpha(u) du \right\} \quad (\text{A.3})$$

for some function $\alpha(u) \rightarrow \alpha_{\min}$, as $u \rightarrow \infty$. Since $1 < \alpha_{\min} < 2$, we may pick δ so small that $\alpha_{\min} - \delta > 1$ and since $T\epsilon_T \rightarrow \infty$, we have for T sufficiently large, that the right-hand side in (A.3) is bounded from above by

$$\epsilon_T^{-1} \exp \{-(\alpha_{\min} - \delta) \log(1/\epsilon_T)\} = \epsilon_T^{\alpha_{\min} - \delta - 1},$$

and the right-hand side converges to zero as $T \rightarrow \infty$. The proof of an infinite limit under Condition 2 is similar. \heartsuit

Proposition 5.A.2 *For every $t \geq 0$*

$$\frac{I(Tt) - EI(Tt)}{b(MT)} \xrightarrow{P} 0, \quad T \rightarrow \infty.$$

Proof. We have

$$[b(MT)]^{-1} EI(T) \leq [b(MT)]^{-1} M E \min(T, X_{on}^{(0)}).$$

Using Karamata's theorem, we obtain

$$[b(MT)]^{-1} M \int_0^T P(X_{on}^{(0)} > x) dx \leq (\text{const}) \frac{MT^2}{b(MT)} \frac{\bar{F}_{\min}(T)}{\bar{F}_{\min}(T)} \frac{\bar{F}_{on}(T)}{\bar{F}_{\min}(T)}.$$

The first term is $o(1)$ by Lemma 5.A.1. The second term is equal to 1 if $\alpha_{on} < \alpha_{off}$. If $\alpha_{on} \geq \alpha_{off}$ it converges either to zero or to ℓ . This completes the proof. \heartsuit

Next, we show that $III(T)$ is asymptotically negligible. By virtue of the Slow Growth Condition $b(MT) = o(T)$, we can find a function $\epsilon_T \rightarrow 0$ such that

$$b(MT) = o(\epsilon_T T) \quad \text{and} \quad 1/\log(T) = o(\epsilon_T) \quad \text{as } T \rightarrow \infty. \quad (\text{A.4})$$

For example, we could let

$$\epsilon_T = \max \left([b(MT)/T]^{1/2}, [\log T]^{-1/2} \right).$$

Lemma 5.A.3 *Assume that ϵ_T satisfies (A.4). Then*

$$M P(|\xi_T - \mu_T| > \epsilon_T \mu_T) = o(1) \quad \text{as } T \rightarrow \infty.$$

Proof. First we treat the case $\xi_T > (1 + \epsilon_T)\mu_T$. Since $Z_i = X_i + Y_i$ has a regularly varying right tail there exist iid mean-zero random variables E_i concentrated on $[-EZ_1, \infty)$ and a positive number x_0 such that for some $\beta > 0$

$$P(Z_1 - EZ_1 > x) \geq P(E_1 > x) \quad \text{for } x \geq -EZ_1,$$

and

$$P(E_1 > x) = e^{-\beta x}, \quad x \geq x_0.$$

Then a stochastic domination argument shows that with $m_T = [(1 + \epsilon_T)\mu_T]$,

$$\begin{aligned} P(\xi_T > (1 + \epsilon_T)\mu_T) &= P(T_0 + Z_1 + \dots + Z_{m_T} \leq T) \\ &\leq P(Z_1 + \dots + Z_{m_T} - m_T \mu \leq T - m_T \mu) \\ &\leq P(E_1 + \dots + E_{m_T} \leq T - m_T \mu) \\ &= P((m_T \text{Var}(E_1))^{-1/2}(E_1 + \dots + E_{m_T}) \leq -a_T) \\ &= p_T, \end{aligned}$$

where

$$a_T = (m_T \text{Var}(E_1))^{-1/2} (m_T \mu - T).$$

Since $\mu_T = T/\mu$, we have for some $|\theta_T| \leq 1$, that

$$a_T \sim (\text{const}) \frac{\epsilon_T T + \theta_T}{\sqrt{T}} \sim (\text{const}) \epsilon_T T^{1/2},$$

and hence for all large T

$$a_T \geq T^{1/6},$$

since $\epsilon_T T^{1/2} \geq (\log T)^{-1/2} T^{1/2} \geq T^{1/6}$. The classical Cramér result on large deviations for sums of iid random variables with moment generating function existing in a neighborhood of the origin (see Petrov [71], Theorem 3 in Chapter VIII) gives for large T ,

$$\begin{aligned} p_T &\leq P((m_T \text{Var}(E_1))^{-1/2}(E_1 + \dots + E_{m_T}) \leq -T^{1/6}) \\ &\leq (\text{const}) \Phi(-T^{1/6}) \leq (\text{const}) e^{-T^{1/3}/4}, \end{aligned}$$

where Φ is the standard normal distribution function. Finally, since the Slow Growth Condition on M implies $M = o(T)$ we have

$$M P(\xi_T > (1 + \epsilon_T)\mu_T) \leq M e^{-T^{1/3}/4} = o(1). \quad (\text{A.5})$$

Next we treat the case $\xi_T < (1 - \epsilon_T)\mu_T$. Choose

$$\beta_T = T - [(1 - \epsilon_T)\mu_T] \mu \sim \epsilon_T T,$$

with ϵ_T obeying (A.4). Notice that as $T \rightarrow \infty$,

$$\beta_T^{-1}(Z_1 + \dots + Z_{[(1-\epsilon_T)\mu_T]} - [(1-\epsilon_T)\mu_T]\mu) \xrightarrow{P} 0.$$

Moreover, condition (C.2) in Section 5.C is satisfied for $B_n = [\beta_n, \infty)$ and the sequence $(Z_i - \mu)_{i \geq 1}$ since for x large enough

$$x^{-2} E(Z_i^2 1_{\{Z_i \leq x\}}) \leq (\text{const}) P(Z_i > x),$$

as a consequence of Karamata's theorem. An application of Corollary 5.C.2 now shows that

$$\begin{aligned} P(\xi_T < (1-\epsilon_T)\mu_T) &= P(T_0 + Z_1 + \dots + Z_{[(1-\epsilon_T)\mu_T]} > T) \\ &\sim P(Z_1 + \dots + Z_{[(1-\epsilon_T)\mu_T]} - [(1-\epsilon_T)\mu_T]\mu > \beta_T) \\ &\sim [(1-\epsilon_T)\mu_T] P(Z > \beta_T) \\ &\sim \mu_T P(Z > \epsilon_T T) = (\text{const}) T P(Z > \epsilon_T T) \\ &\sim \frac{(\text{const})}{M} \frac{\bar{F}_{\min}(\epsilon_T T)}{\bar{F}_{\min}(b(MT))} = o(M^{-1}). \end{aligned}$$

In the last step we used the fact that \bar{F}_{\min} is regularly varying with index $-\alpha_{\min}$, Proposition 0.8 (iii) in Resnick [77] and (A.4). \heartsuit

We need another auxiliary result.

Lemma 5.A.4 *For all $\delta > 0$,*

$$M [b(MT)]^{-1} E X_{\xi_T} 1_{[X_{\xi_T} > \delta b(MT)]} 1_{[\xi_T \geq 1]} \rightarrow 0 \quad \text{as } T \rightarrow \infty.$$

Proof. Choose $\epsilon_T \rightarrow 0$ such that (A.4) holds. Using Karamata's theorem, we have for large T ,

$$\begin{aligned} &M [b(MT)]^{-1} \int_{\delta b(MT)}^{\infty} P(X_{\xi_T} > x, |\xi_T - \mu_T| \leq \epsilon_T \mu_T, \xi_T \geq 1) dx \\ &\leq M [b(MT)]^{-1} \int_{\delta b(MT)}^{\infty} P\left(\max_{i \geq 1, |i - \mu_T| \leq \epsilon_T \mu_T} X_i > x\right) dx \\ &\leq (\text{const}) M [b(MT)]^{-1} \epsilon_T \mu_T \int_{\delta b(MT)}^{\infty} \bar{F}_{\min}(x) dx \\ &\leq (\text{const}) \delta^{1-\alpha_{\min}} [b(MT)]^{\alpha_{\min} - \alpha_{\min}} \epsilon_T = o(1). \end{aligned}$$

Choose $c_T \rightarrow \infty$ such that $b(MT) = o(c_T^{-1} \epsilon_T T)$. It follows from the proof of Lemma 5.A.3 that

$$M P(|\xi_T - \mu_T| > \epsilon_T \mu_T) = o(c_T^{-\alpha_{\min}}). \quad (\text{A.6})$$

Let $K > 0$ be a constant so large that $T^K > c_T b(MT)$ for large T . The following bound is straightforward:

$$\begin{aligned}
& \int_{\delta b(MT)}^{\infty} P(X_{\xi_T} 1_{[\xi_T \geq 1]} > x, |\xi_T - \mu_T| > \epsilon_T \mu_T) \, dx \\
& \leq \int_{\delta b(MT)}^{c_T b(MT)} P(|\xi_T - \mu_T| > \epsilon_T \mu_T) \, dx \\
& \quad + \int_{c_T b(MT)}^{\infty} P(X_{\xi_T} > x, 1 \leq \xi_T < (1 - \epsilon_T) \mu_T) \, dx \\
& \quad + \int_{c_T b(MT)}^{T^K} P(\xi_T > (1 + \epsilon_T) \mu_T) \, dx + \int_{T^K}^{\infty} P(X_{\xi_T} 1_{[\xi_T \geq 1]} > x) \, dx \\
& = I_1 + I_2 + I_3 + I_4.
\end{aligned}$$

Obviously, by (A.6),

$$M[b(MT)]^{-1} I_1 = (c_T - \delta) M P(|\xi_T - \mu_T| > \epsilon_T \mu_T) = o(1). \quad (\text{A.7})$$

Moreover, by Karamata's theorem,

$$\begin{aligned}
M[b(MT)]^{-1} I_2 & \leq M[b(MT)]^{-1} \int_{c_T b(MT)}^{\infty} P\left(\max_{1 \leq i \leq (1 - \epsilon_T) \mu_T} X_i > x\right) \, dx \\
& \leq (\text{const}) M[b(MT)]^{-1} \mu_T \int_{c_T b(MT)}^{\infty} \bar{F}_{\text{on}}(x) \, dx \\
& \sim (\text{const}) [b(MT)]^{\alpha_{\min} - \alpha_{\text{on}}} c_T^{1 + \alpha_{\min} - \alpha_{\text{on}}} \frac{\bar{F}_{\min}(c_T b(MT))}{\bar{F}_{\min}(b(MT))}.
\end{aligned}$$

Using the right-hand inequality in Proposition 5.D.1, with $x = c_T$, $t = b(MT)$ and $\varepsilon = \alpha_{\min} - 1 - \delta > 0$ for some small $\delta > 0$, gives that there is a fixed t_0 such that for $x \geq 1$ and $t \geq t_0$

$$\frac{\bar{F}_{\min}(c_T b(MT))}{\bar{F}_{\min}(b(MT))} \leq (\alpha_{\min} - \delta) c_T^{-(1+\delta)}.$$

This shows that $M[b(MT)]^{-1} I_2 \rightarrow 0$.

As for the proof of Lemma 5.A.3 (see (A.5)) we conclude that

$$M[b(MT)]^{-1} I_3 \leq M[b(MT)]^{-1} \int_{c_T b(MT)}^{T^K} e^{-T^{1/3}/4} \, dx = o(1) \quad \text{as } T \rightarrow \infty,$$

since the Slow Growth Condition on M holds. Using Markov's inequality and (8.12) in Theorem I8.1 of Gut [43], we have for $\epsilon \in (0, \alpha_{\text{on}} - 1)$ and K sufficiently

large,

$$\begin{aligned}
I_4 &\leq EX_{\xi_T}^{\alpha_{\text{on}}-\epsilon} 1_{[\xi_T \geq 1]} \int_{T^K}^{\infty} x^{-\alpha_{\text{on}}+\epsilon} dx \\
&= (\text{const}) EX_{\xi_T}^{\alpha_{\text{on}}-\epsilon} 1_{[\xi_T \geq 1]} T^{-K(\alpha_{\text{on}}-1-\epsilon)} \\
&= o(T^{-K(\alpha_{\text{on}}-1-\epsilon)+1}).
\end{aligned}$$

The Slow Growth Condition on M implies that $M = o(T^{\alpha_{\text{min}}-1+\epsilon})$. Therefore,

$$M[b(MT)]^{-1} I_4 = o(T^{-K(\alpha_{\text{on}}-1-\epsilon)+\alpha_{\text{min}}+\epsilon}) = o(1),$$

provided K is chosen so large that $K > (\alpha_{\text{min}} + \epsilon)/(\alpha_{\text{on}} - 1 - \epsilon)$. Combining all the estimates above, we finally proved the statement of the lemma. \heartsuit

Now we are ready to deal with $III(T)$.

Proposition 5.A.5 *For every $t \geq 0$, as $T \rightarrow \infty$*

$$\frac{III(Tt) - EIII(Tt)}{b(MT)} \xrightarrow{P} 0.$$

Proof. Fix $\delta > 0$. Define the iid random variables

$$\hat{X}_T^{(m)} = \max \left(0, T_{\xi_T^{(m)}-1}^{(m)} + X_{\xi_T^{(m)}}^{(m)} - T \right) 1_{[\xi_T^{(m)} \geq 1]},$$

and their truncated versions

$$\tilde{X}_T^{(m)} = \hat{X}_T^{(m)} I_{[\hat{X}_T^{(m)} \leq \delta b(MT)]}.$$

By virtue of Lemma 5.A.4, it suffices to show that

$$[b(MT)]^{-1} \sum_{m=1}^M \left(\tilde{X}_T^{(m)} - E\tilde{X}_T^{(m)} \right) \xrightarrow{P} 0.$$

The variance of the sum on the left-hand side is given by

$$M[b(MT)]^{-2} \text{Var}(\tilde{X}_T) \leq M[b(MT)]^{-2} E\tilde{X}_T^2,$$

and so it suffices to show that the right-hand side converges to zero. Assume $\epsilon_T \rightarrow 0$ satisfies (A.4). Then we have

$$\begin{aligned}
E\tilde{X}_T^2 &\leq \delta^2 [b(MT)]^2 P(|\xi_T - \mu_T| > \epsilon_T \mu_T) \\
&\quad + \int_0^{\delta^2 [b(MT)]^2} P(X_{\xi_T} > \sqrt{x}, |\xi_T - \mu_T| \leq \epsilon_T \mu_T) dx \\
&= I_1 + I_2.
\end{aligned}$$

By Lemma 5.A.3 we have

$$M[b(MT)]^{-2}I_1 = o(1).$$

An application of Karamata's theorem yields that

$$\begin{aligned} M[b(MT)]^{-2}I_2 &\leq M[b(MT)]^{-2} \int_0^{\delta^2[b(MT)]^2} P\left(\max_{|i-\mu_T| \leq \epsilon_T \mu_T} X_i > \sqrt{x}\right) dx \\ &\leq (\text{const}) M[b(MT)]^{-2} \epsilon_T \mu_T \int_0^{\delta^2[b(MT)]^2} \bar{F}_{\text{on}}(\sqrt{x}) dx \\ &\sim (\text{const}) \epsilon_T \delta^2 M T \bar{F}_{\text{on}}(\delta b(MT)) \\ &\sim (\text{const}) \delta^{2-\alpha_{\text{on}}} [b(MT)]^{\alpha_{\text{min}}-\alpha_{\text{on}}} \epsilon_T = o(1). \end{aligned}$$

This completes the proof. \heartsuit

5.A.3 Convergence of the marginal distributions

In this section we show that the random variables $II(T)$ converge weakly to a stable distribution as $T \rightarrow \infty$. This fact and the results of the previous section, together with a Slutsky argument, prove the convergence of the one-dimensional distributions of \tilde{V}_T in Theorem 5.5.

Introduce the iid mean zero random variables

$$J_k^{(m)} = X_k^{(m)} - r_{\text{off}} Z_k^{(m)} = r_{\text{on}}(X_k^{(m)} - \mu_{\text{on}}) - r_{\text{off}}(Y_k^{(m)} - \mu_{\text{off}}),$$

where

$$r_{\text{on}} = \mu_{\text{off}}/\mu \quad \text{and} \quad r_{\text{off}} = \mu_{\text{on}}/\mu.$$

The tails of the J_k 's are regularly varying: as $x \rightarrow \infty$

$$P(J_k > x) \sim r_{\text{on}}^{\alpha_{\text{on}}} \bar{F}_{\text{on}}(x) \quad \text{and} \quad P(J_k \leq -x) \sim r_{\text{off}}^{\alpha_{\text{off}}} \bar{F}_{\text{off}}(x).$$

Write

$$S_{T,m} = \sum_{k=1}^{\xi_T^{(m)}} J_k^{(m)}.$$

The following decomposition will be useful:

$$\begin{aligned} II(T) &= \sum_{m=1}^M S_{T,m} + r_{\text{off}} \sum_{m=1}^M T_{\xi_T^{(m)}}^{(m)} \mathbf{1}_{[\xi_T^{(m)} \geq 1]} - r_{\text{off}} \sum_{m=1}^M T_0^{(m)} \mathbf{1}_{[\xi_T^{(m)} \geq 1]} \\ &= II_1(T) + II_2(T) + II_3(T). \end{aligned}$$

In what follows, we show that II_1 has an α_{min} -stable limit whereas II_2 and II_3 are asymptotically negligible. By Proposition 5.A.2 it follows that

$$[b(MT)]^{-1} E II_3(T) \rightarrow 0.$$

Following the argument for Theorem I5.3 in Gut [43], we obtain

$$E(T_{\xi_T} 1_{\{\xi_T \geq 1\}}) = E\left(\sum_{i=1}^{\xi_T} Z_i + T_0 1_{[\xi_T \geq 1]}\right) = T + E(T_0 1_{[\xi_T \geq 1]}) . \quad (\text{A.8})$$

By virtue of Lemma 5.A.1, we have for large T

$$\frac{M}{b(MT)} E(T_{\xi_T} - T) = \frac{M}{b(MT)} E(T_0 1_{[\xi_T \geq 1]}) \leq (\text{const}) \frac{MT^2 \bar{F}_{\min}(T)}{b(MT)} \frac{F_{\text{on}}(T)}{\bar{F}_{\min}(T)} ,$$

which is $o(1)$ by the same arguments as in Proposition 5.A.2. Hence,

$$\frac{II_2(T) - EII_2(T)}{b(MT)} = \frac{r_{\text{off}}}{b(MT)} \sum_{m=1}^M [(T_{\xi_T^{(m)}} - T) - E(T_{\xi_T} - T)] \xrightarrow{P} 0 .$$

Again using the argument for Theorem I5.3 of [43], we obtain

$$EII_1(T) = M E\xi_T EJ_1 = 0 .$$

In the remainder of this section we prove that II_1 has an α_{\min} -stable limit. Let c , σ_{\min} and β be as in Theorem 5.5. Theorem 8 in Chapter IV of Petrov [71] gives the following necessary and sufficient conditions for the sums of row-wise iid random variables $S_{T,m}$, $m = 1, \dots, M$, to converge weakly to the stable distribution $S_{\alpha_{\min}}(c\sigma_{\min}, \beta, 0)$: as $T \rightarrow \infty$, for all $x > 0$:

$$\begin{aligned} \text{(RT)} \quad M P(S_{T,1} > x b(MT)) &\rightarrow C_{\alpha_{\min}} \frac{1+\beta}{2} (c\sigma_{\min})^{\alpha_{\min}} x^{-\alpha_{\min}} , \\ \text{(LT)} \quad M P(S_{T,1} \leq -x b(MT)) &\rightarrow C_{\alpha_{\min}} \frac{1-\beta}{2} (c\sigma_{\min})^{\alpha_{\min}} x^{-\alpha_{\min}} , \\ \text{(VA)} \quad \lim_{\epsilon \downarrow 0} \limsup_{T \rightarrow \infty} M [b(MT)]^{-2} \text{Var}(S_{T,1} 1_{\{|S_{T,1}| < \epsilon b(MT)\}}) &= 0 , \end{aligned}$$

where C_{α} is given by (3.2). Notice that $C_{\alpha_{\min}} \sigma_{\min}^{\alpha_{\min}} = 1$.

We have $S_{T,1} = S^{(1)}(T) - S^{(2)}(T)$, where

$$S^{(1)}(T) = r_{\text{on}} \sum_{k=1}^{\xi_T} (X_k - \mu_{\text{on}}) \quad \text{and} \quad S^{(2)}(T) = r_{\text{off}} \sum_{k=1}^{\xi_T} (Y_k - \mu_{\text{off}}) .$$

Define

$$S_n = \sum_{k=1}^n J_k , \quad S_n^{(1)} = r_{\text{on}} \sum_{k=1}^n (X_k - \mu_{\text{on}}) , \quad S_n^{(2)} = r_{\text{off}} \sum_{k=1}^n (Y_k - \mu_{\text{off}}) .$$

The proofs of (RT), (LT), and (VA) are now presented via a series of lemmas.

Lemma 5.A.6 *If $\alpha_{\text{on}} < \alpha_{\text{off}}$, then for all $x > 0$*

$$M P(-S_{[\mu_T]}^{(1)} > x b(MT)) = o(1), \quad \text{as } T \rightarrow \infty.$$

Correspondingly, if $\alpha_{\text{off}} < \alpha_{\text{on}}$, then for all $x > 0$

$$M P(-S_{[\mu_T]}^{(2)} > x b(MT)) = o(1), \quad \text{as } T \rightarrow \infty.$$

Proof. We will only prove the first statement. The second statement follows by interchanging the roles of X_i and Y_i . Suppose $\alpha_{\text{on}} < \alpha_{\text{off}}$. Let $D = D_T$ be a positive function such that $D \rightarrow 0$ and as $T \rightarrow \infty$,

$$D M \rightarrow \infty \quad \text{and} \quad D b(MT) \rightarrow \infty. \quad (\text{A.9})$$

Let

$$\tilde{X}_k = X_k 1_{[X_k \leq D b(MT)]} \quad \text{and} \quad \tilde{S}_n^{(1)} = \sum_{k=1}^n (\tilde{X}_k - E \tilde{X}),$$

and assume without loss of generality that $r_{\text{on}} = 1$. We have

$$p_T = P(-S_{[\mu_T]}^{(1)} > x b(MT)) \leq P(-\tilde{S}_{[\mu_T]}^{(1)} > x b(MT) - \mu_T E(X - \tilde{X})).$$

Using Karamata's theorem, we have

$$\frac{\mu_T E(X - \tilde{X})}{b(MT)} \sim (\text{const}) \frac{D}{M} \frac{\overline{F}_{\text{on}}(D b(MT))}{\overline{F}_{\text{on}}(b(MT))}. \quad (\text{A.10})$$

Using the left-hand inequality of Proposition 5.D.1, with $x = 1/D$, $t = D b(MT)$ and $\varepsilon = 2 - \alpha_{\text{on}}$ gives that there is a fixed t_0 such that for $x \geq 1$ and $t \geq t_0$ the right-hand side of (A.10) is bounded by

$$\frac{(\text{const})}{\alpha_{\text{on}} - 1} \frac{1}{DM},$$

which is $o(1)$ by (A.9). So we may bound the probability p_T for large T from above by

$$p_T \leq P\left(-[\text{Var}(\tilde{X})\mu_T]^{-1/2} \tilde{S}_{[\mu_T]}^{(1)} > a_T(x)\right), \quad \text{where} \quad a_T(x) = \frac{x b(MT)/2}{[\text{Var}(\tilde{X})\mu_T]^{1/2}}.$$

Using a non-uniform Berry–Esséen estimate in the central limit theorem (see Petrov [72], Theorem 5.16) the right-hand side is bounded by

$$\overline{\Phi}(a_T(x)) + (\text{const}) \frac{E|\tilde{X} - E\tilde{X}|^3}{\mu_T^{1/2} [\text{Var}(\tilde{X})]^{3/2} (1 + a_T(x))^3}, \quad (\text{A.11})$$

where $\overline{\Phi}$ denotes the right tail of the standard normal distribution. Notice that

$$a_T(x) \sim (\text{const}) \left(\frac{M}{D^2} \frac{\overline{F}_{\text{on}}(b(MT))}{\overline{F}_{\text{on}}(D b(MT))} \right)^{1/2}.$$

As above we apply the left-hand inequality of Proposition 5.D.1, to obtain that for large T

$$a_T(x) > (\text{const}) (\alpha_{\text{on}} - 1) M^{1/2},$$

so the first term in (A.11) decreases at an exponential (in M) rate and hence, is $o(M^{-1})$. The second term behaves asymptotically as

$$(\text{const}) x^{-3} \frac{D^3}{M} \frac{\bar{F}_{\text{on}}(D b(MT))}{\bar{F}_{\text{on}}(b(MT))}. \quad (\text{A.12})$$

Using the left-hand inequality of Proposition 5.D.1, as before, gives that the right-hand side of (A.12) is bounded from above by

$$\frac{(\text{const})}{\alpha_{\text{on}} - 1} x^{-3} \frac{D}{M} = o(M^{-1}).$$

This completes the proof. \heartsuit

For $\epsilon_T \rightarrow 0$ satisfying (A.4), define the event

$$\Theta_T = \{|\xi_T - \mu_T| \leq \epsilon_T \mu_T\}. \quad (\text{A.13})$$

Lemma 5.A.7 *For all $x > 0$,*

$$M P(|S_{T,1} - S_{[\mu_T]}| > x b(MT), \Theta_T) = o(1) \quad \text{as } T \rightarrow \infty.$$

Proof. Using Theorem 2.3 in Petrov [72], we have

$$\begin{aligned} P(|S_{T,1} - S_{[\mu_T]}| > x b(MT), \Theta_T) &\leq P\left(\max_{|j - \mu_T| \leq \epsilon_T \mu_T} |S_j - S_{[\mu_T]}| > x b(MT)\right) \\ &\leq (\text{const}) P(|S_{[\epsilon_T \mu_T]}| > x b(MT)/2). \end{aligned}$$

Applying the same result, we also see that

$$\begin{aligned} P(|S_{[\epsilon_T \mu_T]}| > x b(MT)/2) &= P(|S_{[\epsilon_T \mu_T]}^{(1)} - S_{[\epsilon_T \mu_T]}^{(2)}| > x b(MT)/2) \\ &\leq (\text{const}) \sum_{i=1}^2 P\left(S_{[\epsilon_T \mu_T]}^{(i)} - \hat{S}_{[\epsilon_T \mu_T]}^{(i)} > \frac{x}{4} b(MT)\right), \end{aligned}$$

where $\hat{S}^{(1)}$ and $\hat{S}^{(2)}$ are independent copies of $S^{(1)}$ and $S^{(2)}$. Using Corollary 5.C.2, we see that the two probabilities on the right-hand side multiplied by M are asymptotic to

$$\begin{aligned} &M \epsilon_T \mu_T \left[(\text{const}) x^{-\alpha_{\text{on}}} \bar{F}_{\text{on}}(b(MT)) + (\text{const}) x^{-\alpha_{\text{off}}} \bar{F}_{\text{off}}(b(MT)) \right] \\ &\sim \epsilon_T \left[(\text{const}) x^{-\alpha_{\text{on}}} [b(MT)]^{\alpha_{\text{min}} - \alpha_{\text{on}}} + (\text{const}) x^{-\alpha_{\text{off}}} [b(MT)]^{\alpha_{\text{min}} - \alpha_{\text{off}}} \right], \end{aligned}$$

which is $o(1)$. This completes the proof. \heartsuit

Lemma 5.A.8 *If $\alpha_{\text{on}} < \alpha_{\text{off}}$, then for all $x > 0$*

$$M P(S_{[\mu_T]} \leq -x b(MT)) = o(1) \quad \text{as } T \rightarrow \infty.$$

Equivalently, if $\alpha_{\text{off}} < \alpha_{\text{on}}$, then for all $x > 0$

$$M P(S_{[\mu_T]} > x b(MT)) = o(1) \quad \text{as } T \rightarrow \infty.$$

Proof. Suppose $\alpha_{\text{on}} < \alpha_{\text{off}}$. We have

$$P(S_{[\mu_T]} \leq -x b(MT)) \leq P(-S_{[\mu_T]}^{(1)} > x b(MT)/2) + P(S_{[\mu_T]}^{(2)} > x b(MT)/2).$$

The first probability is $o(M^{-1})$ by Lemma 5.A.6. The second probability can be treated as follows. Let $\delta > 0$ such that $\alpha_{\text{on}} + \delta < \alpha_{\text{off}}$. Using Markov's inequality and a bound for the $(\alpha_{\text{on}} + \delta)$ th moment of sums of independent mean-zero random variables (see Petrov [71], p.60), we obtain

$$\begin{aligned} M P(S_{[\mu_T]}^{(2)} > x b(MT)/2) &\leq (\text{const}) \frac{M}{[x b(MT)]^{\alpha_{\text{on}} + \delta}} E|S_{[\mu_T]}^{(2)}|^{\alpha_{\text{on}} + \delta} \\ &\leq (\text{const}) \frac{M \mu_T}{[b(MT)]^{\alpha_{\text{on}} + \delta}} \frac{E|Y_{\text{off}} - EY_{\text{off}}|^{\alpha_{\text{on}} + \delta}}{x^{\alpha_{\text{on}} + \delta}}, \end{aligned}$$

which is $o(1)$ since $T/b(T)$ is regularly varying with index $-\delta/\alpha_{\text{on}}$.

Next, suppose $\alpha_{\text{off}} < \alpha_{\text{on}}$. Again, we have

$$P(S_{[\mu_T]} > x b(MT)) \leq P(S_{[\mu_T]}^{(1)} > x b(MT)/2) + P(-S_{[\mu_T]}^{(2)} > x b(MT)/2).$$

The second probability is $o(M^{-1})$ by Lemma 5.A.6. As above we use Markov's inequality for the first probability. This completes the proof. \heartsuit

Lemma 5.A.9 *If $\alpha_{\text{off}} < \alpha_{\text{on}}$, then*

$$\lim_{\epsilon \downarrow 0} \limsup_{T \rightarrow \infty} M [b(MT)]^{-2} \int_0^{\epsilon^2 [b(MT)]^2} P(S_{[\mu_T]} > \sqrt{x}/2) dx = o(1) \quad \text{as } T \rightarrow \infty.$$

Proof. As in Lemma 5.A.8, we have

$$P(S_{[\mu_T]} > \sqrt{x}/2) \leq P(S_{[\mu_T]}^{(1)} > \sqrt{x}/4) + P(-S_{[\mu_T]}^{(2)} > \sqrt{x}/4).$$

Let $\eta > 0$, such that $\alpha_{\text{off}} + \eta < \alpha_{\text{on}}$. Using Markov's inequality (as in the proof of Lemma 5.A.8), we obtain

$$\begin{aligned} M [b(MT)]^{-2} \int_0^{\epsilon^2 [b(MT)]^2} P(S_{[\mu_T]}^{(1)} > \sqrt{x}/4) dx \\ &\leq M [b(MT)]^{-2} \int_0^{\epsilon^2 [b(MT)]^2} \frac{\mu_T E|X_{\text{on}} - EX_{\text{on}}|^{\alpha_{\text{off}} + \eta}}{(\sqrt{x}/4)^{\alpha_{\text{off}} + \eta}} dx \\ &\sim (\text{const}) M T [b(MT)]^{-2} (\epsilon^2 [b(MT)]^2)^{1 - (\alpha_{\text{off}} + \eta)/2} \\ &= (\text{const}) M T [b(MT)]^{-\alpha_{\text{off}} - \eta} \epsilon^{2 - (\alpha_{\text{off}} + \eta)} = o(1), \end{aligned}$$

as $T \rightarrow \infty$. Using the same approach we can show that for some small $\delta > 0$, as $T \rightarrow \infty$

$$M [b(MT)]^{-2} \int_0^{\epsilon^2[b(MT)]^2/M^{1-\delta}} P(-S_{[\mu_T]}^{(2)} > \sqrt{x}/4) dx = o(1). \quad (\text{A.14})$$

Indeed, for some small $\eta > 0$ we have the upper-bound

$$\begin{aligned} M [b(MT)]^{-2} & \int_0^{\epsilon^2[b(MT)]^2/M^{1-\delta}} \frac{\mu_T E|Y_{\text{off}} - EY_{\text{off}}|^{\alpha_{\text{off}}-\eta}}{(\sqrt{x}/4)^{\alpha_{\text{off}}-\eta}} dx \\ & \sim (\text{const}) M T [b(MT)]^{-2} \left(\frac{\epsilon^2[b(MT)]^2}{M^{1-\delta}} \right)^{1-(\alpha_{\text{off}}-\eta)/2} \\ & = (\text{const}) \frac{MT}{[b(MT)]^{\alpha_{\text{off}}-\eta}} M^{(\delta-1)(1-(\alpha_{\text{off}}-\eta)/2)} \epsilon^{2-(\alpha_{\text{off}}-\eta)} = d_T. \end{aligned}$$

By virtue of the Slow Growth Condition $MT[b(MT)]^{-\alpha_{\text{off}}} = o(T^\nu)$ and $M = o(T^{\alpha_{\text{off}}-1+\nu})$ for some small $\nu > 0$. Hence, as $T \rightarrow \infty$

$$d_T = o(T^K),$$

with

$$\begin{aligned} K &= \nu + \eta - (\alpha_{\text{off}} - 1 + \nu) (1 - \delta) (1 - (\alpha_{\text{off}} - \eta)/2) \\ &\leq \nu + \eta - (\alpha_{\text{off}} - 1) (1 - \delta - \alpha_{\text{off}}/2). \end{aligned}$$

Choosing $\delta < 1 - \alpha_{\text{off}}/2$ and ν and η small such that $K \leq 0$ proves (A.14).

It remains to consider

$$M [b(MT)]^{-2} \int_{\epsilon^2[b(MT)]^2/M^{1-\delta}}^{\epsilon^2[b(MT)]^2} P(-S_{[\mu_T]}^{(2)} > \sqrt{x}/4) dx.$$

The approach is the same as in the proof of Lemma 5.A.6. Let $D = D_T$ be a positive function such that $D \rightarrow 0$ as $T \rightarrow \infty$. In particular, we take

$$D = M^{-1/2}. \quad (\text{A.15})$$

Notice that this implies

$$D b(MT) \rightarrow \infty. \quad (\text{A.16})$$

Let

$$\tilde{Y}_k = Y_k 1_{[Y_k \leq D b(MT)]} \quad \text{and} \quad \tilde{S}_n^{(2)} = \sum_{k=1}^n (\tilde{Y}_k - E(\tilde{Y})),$$

and assume without loss of generality that $r_{\text{off}} = 1$. We have

$$P(-S_{[\mu_T]}^{(2)} > \sqrt{x}/4) \leq P(-\tilde{S}_{[\mu_T]}^{(2)} > \sqrt{x}/4 - \mu_T E(Y - \tilde{Y})).$$

Using (A.16), we get

$$\frac{\mu_T E(Y - \tilde{Y})}{b(MT)/M^{(1-\delta)/2}} \sim (\text{const}) \frac{D}{M^{(1+\delta)/2}} \frac{\bar{F}_{\text{off}}(D b(MT))}{\bar{F}_{\text{off}}(b(MT))}. \quad (\text{A.17})$$

The left-hand inequality of Proposition 5.D.1, with $x = 1/D$, $t = D b(MT)$ and $\varepsilon = 2 - \alpha_{\text{off}}$ gives that there is a fixed t_0 such that for $x \geq 1$ and $t \geq t_0$ the right-hand side of (A.17) is bounded from above by

$$\frac{(\text{const})}{\alpha_{\text{off}} - 1} \frac{1}{D M^{(1+\delta)/2}},$$

which is $o(1)$ by (A.15). Hence, we have the following upper-bound for large T :

$$\begin{aligned} M [b(MT)]^{-2} \int_{\epsilon^2 [b(MT)]^2 / M^{(1-\delta)}}^{\epsilon^2 [b(MT)]^2} P(-S_{[\mu_T]}^{(2)} > \sqrt{x}/4) dx \\ \leq M [b(MT)]^{-2} \int_{\epsilon^2 [b(MT)]^2 / M^{(1-\delta)}}^{\epsilon^2 [b(MT)]^2} P(-\tilde{S}_{[\mu_T]}^{(2)} > \sqrt{x}/8) dx \\ = M [b(MT)]^{-2} \int_{\epsilon^2 [b(MT)]^2 / M^{(1-\delta)}}^{\epsilon^2 [b(MT)]^2} P\left(\frac{-\tilde{S}_{[\mu_T]}^{(2)}}{[\text{Var}(\tilde{Y})\mu_T]^{1/2}} > a_T(x)\right) dx, \end{aligned}$$

where

$$a_T(x) = \frac{\sqrt{x}/8}{[\text{Var}(\tilde{Y})\mu_T]^{1/2}}.$$

Using Petrov [72], Theorem 5.16, the right-hand side is bounded by

$$\begin{aligned} \frac{M}{[b(MT)]^2} \int_{\epsilon^2 [b(MT)]^2 / M^{(1-\delta)}}^{\epsilon^2 [b(MT)]^2} \bar{\Phi}(a_T(x)) dx \\ + \frac{(\text{const}) M}{[b(MT)]^2} \int_{\epsilon^2 [b(MT)]^2 / M^{(1-\delta)}}^{\epsilon^2 [b(MT)]^2} \left[\frac{E|\tilde{Y} - E\tilde{Y}|^3 (\text{Var}(\tilde{Y}))^{-3/2}}{(\mu_T)^{1/2} (1 + a_T(x))^3} \right] dx, \end{aligned} \quad (\text{A.18})$$

where $\bar{\Phi}$ denotes the right tail of the standard normal distribution. Notice that

$$\frac{b(MT)}{(M^{(1-\delta)} \text{Var}(\tilde{Y})\mu_T)^{1/2}} \sim (\text{const}) \left(\frac{M^\delta}{D^2} \frac{\bar{F}_{\text{off}}(b(MT))}{\bar{F}_{\text{off}}(D b(MT))} \right)^{1/2}. \quad (\text{A.19})$$

As above we apply the left-hand inequality of Proposition 5.D.1, with $x = 1/D$, $t = D b(MT)$ and $\varepsilon = 2 - \alpha_{\text{off}} - \eta$ for some small $\eta > 0$. This gives that for large T , the right-hand side of (A.19) is bounded from below by

$$(\text{const}) \left(\frac{(\alpha_{\text{off}} + \eta - 1) M^\delta}{D^\eta} \right)^{1/2} \rightarrow \infty.$$

Therefore, the first term in (A.18) is bounded (for large T) by

$$\begin{aligned} M \epsilon^2 \bar{\Phi} \left(\frac{\epsilon b(MT)}{8(M^{(1-\delta)} \text{Var}(\tilde{Y}) \mu_T)^{1/2}} \right) \\ \leq (\text{const}) M \exp\{-(\text{const}) (\alpha_{\text{off}} + \eta - 1)^{1/2} M^{\delta/2} D^{-\eta/2}\}, \end{aligned}$$

which is $o(1)$ by (A.15). The second term in (A.18) is asymptotic to

$$(\text{const}) \frac{M T}{[b(MT)]^2} [D b(MT)]^3 \bar{F}_{\text{off}}(D b(MT)) \int_{\epsilon^2 [b(MT)]^2 / M^{(1-\delta)}}^{\epsilon^2 [b(MT)]^2} x^{-3/2} dx = c_T.$$

Notice that the integral has upper-bound

$$(\text{const}) \frac{M^{(1-\delta)/2}}{\epsilon b(MT)}.$$

Therefore,

$$c_T \leq (\text{const}) D^3 \epsilon^{-1} M^{(1-\delta)/2} \frac{\bar{F}_{\text{off}}(D b(MT))}{\bar{F}_{\text{off}}(b(MT))}. \quad (\text{A.20})$$

Using the left-hand inequality of Proposition 5.D.1, with $x = 1/D$, $t = D b(MT)$ and $\varepsilon = 2 - \alpha_{\text{off}}$, gives that for large T the right-hand side of (A.20) is bounded from above by

$$\frac{(\text{const})}{\epsilon (\alpha_{\text{off}} - 1)} D M^{(1-\delta)/2},$$

which is $o(1)$ by (A.15). This completes the proof. \heartsuit

The following proposition shows the weak convergence of $II_1(T)$ to the marginal distribution (at $t = 1$) of the α_{min} -stable Lévy motion in Theorem 5.5.

Proposition 5.A.10 *Let c , σ_{min} and β be as in Theorem 5.5. We have*

$$[b(MT)]^{-1} II_1(T) \xrightarrow{d} c \Lambda_{\alpha_{\text{min}}, \sigma_{\text{min}}, \beta}(1) \quad \text{as } T \rightarrow \infty. \quad (\text{A.21})$$

Proof. We have to show **(RT)**, **(LT)** and **(VA)** defined at the beginning of Section 5.A.3.

Proof of (RT)

We have to show that for all $x > 0$

$$M P(S_{T,1} > x b(MT)) \rightarrow c^{\alpha_{\text{min}}} \frac{1+\beta}{2} x^{-\alpha_{\text{min}}}, \quad \text{as } T \rightarrow \infty.$$

Recall the definition of Θ_T from (A.13). By Lemma 5.A.3 it suffices to consider the intersection of $\{S_{T,1} > x b(MT)\}$ with Θ_T . For $\delta \in (0, 1)$ we have

$$\begin{aligned} P(S_{T,1} > x b(MT), \Theta_T) &\leq P(S_{T,1} - S_{[\mu_T]} > \delta x b(MT), \Theta_T) \\ &\quad + P(S_{[\mu_T]} > (1 - \delta)x b(MT)). \end{aligned}$$

The first probability on the right-hand side is $o(M^{-1})$ by Lemma 5.A.7. If $\alpha_{\text{off}} < \alpha_{\text{on}}$, then the second probability is $o(M^{-1})$ by Lemma 5.A.8. If $\alpha_{\text{on}} \leq \alpha_{\text{off}}$ then, using Corollary 5.C.2, we have

$$M P(S_{[\mu_T]} > (1-\delta)x b(MT)) \sim \begin{cases} r_{\text{on}}^{\alpha_{\text{on}}} \mu^{-1} (1-\delta)^{-\alpha_{\text{on}}} x^{-\alpha_{\text{on}}} & \text{if } \alpha_{\text{on}} < \alpha_{\text{off}}, \\ r_{\text{on}}^{\alpha_{\text{on}}} \ell \mu^{-1} (1-\delta)^{-\alpha_{\text{on}}} x^{-\alpha_{\text{on}}} & \text{if } \alpha_{\text{on}} = \alpha_{\text{off}}. \end{cases} \quad (\text{A.22})$$

A lower bound is given by

$$\begin{aligned} & P(S_{T,1} > x b(MT), \Theta_T) \\ & \geq P(S_{T,1} - S_{[\mu_T]} > -\delta x b(MT), S_{[\mu_T]} > (1+\delta)x b(MT), \Theta_T) \\ & \geq P(S_{[\mu_T]} > (1+\delta)x b(MT)) - P(S_{T,1} - S_{[\mu_T]} \leq -\delta x b(MT), \Theta_T) \\ & \quad - P(\Theta_T^c). \end{aligned}$$

The second and third probabilities on the right-hand side are $o(M^{-1})$ by Lemmas 5.A.7 and 5.A.3. If $\alpha_{\text{off}} < \alpha_{\text{on}}$, then the first probability is $o(M^{-1})$ by Lemma 5.A.8. If $\alpha_{\text{on}} \leq \alpha_{\text{off}}$ then using Corollary 5.C.2 gives

$$M P(S_{[\mu_T]} > (1+\delta)x b(MT)) \sim \begin{cases} r_{\text{on}}^{\alpha_{\text{on}}} \mu^{-1} (1+\delta)^{-\alpha_{\text{on}}} x^{-\alpha_{\text{on}}} & \text{if } \alpha_{\text{on}} < \alpha_{\text{off}}, \\ r_{\text{on}}^{\alpha_{\text{on}}} \ell \mu^{-1} (1+\delta)^{-\alpha_{\text{on}}} x^{-\alpha_{\text{on}}} & \text{if } \alpha_{\text{on}} = \alpha_{\text{off}}. \end{cases} \quad (\text{A.23})$$

Notice that the different limit for $\alpha_{\text{on}} = \alpha_{\text{off}}$ is due to the fact that b is the quantile function of F_{off} in this case.

Letting $\delta \rightarrow 0$ in (A.22) and (A.23) yields, as $T \rightarrow \infty$

$$M P(S_{T,1} > x b(MT)) \sim \begin{cases} o(1) & \text{if } \alpha_{\text{on}} > \alpha_{\text{off}}, \\ r_{\text{on}}^{\alpha_{\text{on}}} \ell \mu^{-1} x^{-\alpha_{\text{on}}} & \text{if } \alpha_{\text{on}} = \alpha_{\text{off}}, \\ r_{\text{on}}^{\alpha_{\text{on}}} \mu^{-1} x^{-\alpha_{\text{on}}} & \text{if } \alpha_{\text{on}} < \alpha_{\text{off}}. \end{cases}$$

This completes the proof of (RT).

Proof of (LT)

We have to show that for all $x > 0$

$$M P(S_{T,1} \leq -x b(MT)) \rightarrow c^{\alpha_{\min}} \frac{1-\beta}{2} x^{-\alpha_{\min}}, \quad \text{as } T \rightarrow \infty.$$

As before, we consider the intersection of $\{S_{T,1} \leq -x b(MT)\}$ with Θ_T . For $\delta \in (0, 1)$ we have

$$\begin{aligned} P(S_{T,1} \leq -x b(MT), \Theta_T) & \leq P(S_{T,1} - S_{[\mu_T]} \leq -\delta x b(MT), \Theta_T) \\ & \quad + P(S_{[\mu_T]} \leq -(1-\delta)x b(MT)). \end{aligned}$$

The first probability on the right-hand side is $o(M^{-1})$ by Lemma 5.A.7. If $\alpha_{\text{on}} < \alpha_{\text{off}}$, then the second probability is $o(M^{-1})$ by Lemma 5.A.8. If $\alpha_{\text{off}} \leq \alpha_{\text{on}}$, then applying Corollary 5.C.2 gives

$$M P(S_{[\mu_T]} \leq -(1 - \delta)x b(MT)) \sim r_{\text{off}}^{\alpha_{\text{off}}} \mu^{-1} (1 - \delta)^{-\alpha_{\text{off}}} x^{-\alpha_{\text{off}}}. \quad (\text{A.24})$$

On the other hand,

$$\begin{aligned} P(S_{T,1} \leq -x b(MT), \Theta_T) &\geq P(S_{T,1} - S_{[\mu_T]} \leq \delta x b(MT), S_{[\mu_T]} \leq -(1 + \delta)x b(MT), \Theta_T) \\ &\geq P(S_{[\mu_T]} \leq -(1 + \delta)x b(MT)) - P(S_{T,1} - S_{[\mu_T]} > \delta x b(MT), \Theta_T) \\ &\quad - P(\Theta_T^c). \end{aligned}$$

The second and third probabilities on the right-hand side are $o(M^{-1})$ by Lemmas 5.A.7 and 5.A.3. If $\alpha_{\text{on}} < \alpha_{\text{off}}$, then the first probability is $o(M^{-1})$ by Lemma 5.A.8. If $\alpha_{\text{off}} \leq \alpha_{\text{on}}$, then applying Corollary 5.C.2 gives

$$M P(S_{[\mu_T]} \leq -(1 + \delta)x b(MT)) \sim r_{\text{off}}^{\alpha_{\text{off}}} \mu^{-1} (1 + \delta)^{-\alpha_{\text{off}}} x^{-\alpha_{\text{off}}}. \quad (\text{A.25})$$

Letting $\delta \rightarrow 0$ in (A.24) and (A.25) yields, as $T \rightarrow \infty$

$$M P(S_{T,1} \leq -x b(MT)) \sim \begin{cases} o(1) & \text{if } \alpha_{\text{on}} < \alpha_{\text{off}}, \\ r_{\text{off}}^{\alpha_{\text{off}}} \mu^{-1} x^{-\alpha_{\text{off}}} & \text{if } \alpha_{\text{off}} \leq \alpha_{\text{on}}. \end{cases}$$

This completes the proof of **(LT)**.

Proof of (VA)

We have to show that

$$\lim_{\epsilon \downarrow 0} \limsup_{T \rightarrow \infty} M [b(MT)]^{-2} \text{Var}(S_{T,1} \mathbf{1}_{\{|S_{T,1}| < \epsilon b(MT)\}}) = 0.$$

We have

$$\begin{aligned} \text{Var}(S_{T,1} \mathbf{1}_{\{|S_{T,1}| < \epsilon b(MT)\}}) &\leq E([S_{T,1}]^2 \mathbf{1}_{\{|S_{T,1}| < \epsilon b(MT)\}}) \\ &= \int_0^{\epsilon^2 [b(MT)]^2} P([S_{T,1}]^2 > x) dx \\ &= \int_0^{\epsilon^2 [b(MT)]^2} P(S_{T,1} > \sqrt{x}) dx \\ &\quad + \int_0^{\epsilon^2 [b(MT)]^2} P(S_{T,1} \leq -\sqrt{x}) dx. \end{aligned}$$

In our proof we will only consider the first integral. The second one can be treated analogously. It suffices to intersect $\{S_{T,1} > \sqrt{x}\}$ with Θ_T by virtue of

Lemma 5.A.3. Combining the upper-bound in the proof of **(RT)** with the proof of Lemma 5.A.7, we obtain the following upper-bound:

$$\begin{aligned} & \frac{M}{[b(MT)]^2} \int_0^{\epsilon^2[b(MT)]^2} (\text{const}) P(|S_{[\epsilon_T \mu_T]}| > \sqrt{x}/4) dx \\ & + \frac{M}{[b(MT)]^2} \int_0^{\epsilon^2[b(MT)]^2} P(S_{[\mu_T]} > \sqrt{x}/2) dx. \end{aligned}$$

As in the proof of Lemma 5.A.7, we can use a symmetrization inequality and Corollary 5.C.2 to show that the first term is $o(1)$ as $T \rightarrow \infty$. If $\alpha_{\text{on}} \leq \alpha_{\text{off}}$ then, using Corollary 5.C.2, we have

$$\begin{aligned} & \frac{M}{[b(MT)]^2} \int_0^{\epsilon^2[b(MT)]^2} P(S_{[\mu_T]} > \sqrt{x}/2) dx \\ & \sim (\text{const}) \frac{MT}{[b(MT)]^2} \epsilon^{2-\alpha_{\text{on}}} [b(MT)]^2 \bar{F}_{\min}(b(MT)) \\ & \sim (\text{const}) \epsilon^{2-\alpha_{\text{on}}} \rightarrow 0, \end{aligned}$$

as $\epsilon \downarrow 0$. If $\alpha_{\text{off}} < \alpha_{\text{on}}$, then the second term vanishes by Lemma 5.A.9. This completes the proof of **(VA)**. \heartsuit

Now an appeal to the decomposition (A.1), Propositions 5.A.2, 5.A.5 and 5.A.10 in combination with a Slutsky argument proves the convergence of the one-dimensional distributions of the normalized cumulative workload process \tilde{V}_T to α_{\min} -stable distributions.

5.A.4 Convergence of the finite-dimensional distributions

In this section we complete the proof of the first part of Theorem 5.5 by showing that the finite-dimensional distributions of $II_1(Tt)$ converge to α_{\min} -stable Lévy motion. By the Cramér-Wold theorem, we have to show the following. Let $n \geq 1$, $a_1, a_2, \dots, a_n \in \mathbb{R}$ and $t_1, t_2, \dots, t_n \geq 0$. Then

$$\sum_{i=1}^n a_i II_1(Tt_i) \xrightarrow{d} c \sum_{i=1}^n a_i \Lambda_{\alpha_{\min}, \sigma_{\min}, \beta}(t_i) \quad \text{as } T \rightarrow \infty.$$

It is easier to show convergence for the increments of the $II_1(Tt)$ process. For illustrational purposes we only consider $II_1(Tt_1)$ and $II_1(Tt_2) - II_1(Tt_1)$. First we prove the following lemma.

Lemma 5.A.11 *Let $a_1, a_2 \in \mathbb{R}$, $t_2 \geq t_1 \geq 0$. Define*

$$Z_T^{(1)} = a_1 \sum_{k=1}^{[\mu_{Tt_1}]} J_k^{(m)} \quad \text{and} \quad Z_T^{(2)} = a_2 \sum_{k=[\mu_{Tt_1}]+1}^{[\mu_{Tt_2}]} J_k^{(m)}.$$

Then we have as $T \rightarrow \infty$, for all $x > 0$

$$\begin{aligned} M P(Z_T^{(1)} + Z_T^{(2)} > x b(MT)) &\sim M P(Z_T^{(1)} > x b(MT)) \\ &\quad + M P(Z_T^{(2)} > x b(MT)). \end{aligned} \quad (\text{A.26})$$

Proof. Suppose that $a_1 > 0$ and $a_2 > 0$. If $\alpha_{\text{off}} < \alpha_{\text{on}}$, it follows from the proof of **(RT)** that both the left-hand side and right-hand side of (A.26) are $o(1)$.

Suppose $\alpha_{\text{on}} \leq \alpha_{\text{off}}$. Since, for $\delta \in (0, 1/2)$,

$$\begin{aligned} P(Z_T^{(1)} + Z_T^{(2)} > x b(MT)) &\leq P(Z_T^{(1)} > (1 - \delta)x b(MT)) \\ &\quad + P(Z_T^{(2)} > (1 - \delta)x b(MT)) \\ &\quad + P(Z_T^{(1)} > \delta x b(MT)) P(Z_T^{(2)} > \delta x b(MT)), \end{aligned}$$

we have by Corollary 5.C.2

$$\begin{aligned} &\limsup_{T \rightarrow \infty} \frac{M P(Z_T^{(1)} + Z_T^{(2)} > x b(MT))}{M P(Z_T^{(1)} > x b(MT)) + M P(Z_T^{(2)} > x b(MT))} \\ &\leq \limsup_{T \rightarrow \infty} \frac{M P(Z_T^{(1)} > (1 - \delta)x b(MT)) + M P(Z_T^{(2)} > (1 - \delta)x b(MT))}{M P(Z_T^{(1)} > x b(MT)) + M P(Z_T^{(2)} > x b(MT))} \\ &= (1 - \delta)^{-\alpha_{\text{on}}}. \end{aligned} \quad (\text{A.27})$$

On the other hand, for $\delta \in (0, 1/2)$

$$\begin{aligned} P(Z_T^{(1)} + Z_T^{(2)} > x b(MT)) &\geq P(Z_T^{(1)} > (1 + \delta)x b(MT)) P(|Z_T^{(2)}| \leq \delta x b(MT)) \\ &\quad + P(Z_T^{(2)} > (1 + \delta)x b(MT)) P(|Z_T^{(1)}| \leq \delta x b(MT)), \end{aligned}$$

so we have by Corollary 5.C.2

$$\begin{aligned} &\liminf_{T \rightarrow \infty} \frac{M P(Z_T^{(1)} + Z_T^{(2)} > x b(MT))}{M P(Z_T^{(1)} > x b(MT)) + M P(Z_T^{(2)} > x b(MT))} \\ &\geq \liminf_{T \rightarrow \infty} \frac{M P(Z_T^{(1)} > (1 + \delta)x b(MT)) + M P(Z_T^{(2)} > (1 + \delta)x b(MT))}{M P(Z_T^{(1)} > x b(MT)) + M P(Z_T^{(2)} > x b(MT))} \\ &= (1 + \delta)^{-\alpha_{\text{on}}}. \end{aligned} \quad (\text{A.28})$$

Letting $\delta \rightarrow 0$ in (A.27) and (A.28) completes the proof. For negative a_i the proof is similar. \heartsuit

Proposition 5.A.12 *Let $a_1, a_2 \in \mathbb{R}$ and $t_2 \geq t_1 \geq 0$. Then, as $T \rightarrow \infty$*

$$\begin{aligned} & a_1 II_1(Tt_1) + a_2 (II_1(Tt_2) - II_1(Tt_1)) \\ & \xrightarrow{d} a_1 c \Lambda_{\alpha_{\min}, \sigma_{\min}, \beta}(t_1) \\ & + a_2 (c \Lambda_{\alpha_{\min}, \sigma_{\min}, \beta}(t_2) - c \Lambda_{\alpha_{\min}, \sigma_{\min}, \beta}(t_1)). \end{aligned}$$

Proof. Assume that $a_1 > 0$ and $a_2 > 0$. Define

$$Z_{T,m} = a_1 [b(MT)]^{-1} \sum_{k=1}^{\xi_{Tt_1}^{(m)}} J_k^{(m)} + a_2 [b(MT)]^{-1} \sum_{k=\xi_{Tt_1}^{(m)}+1}^{\xi_{Tt_2}^{(m)}} J_k^{(m)}.$$

According to Petrov [71], Theorem 8 in Chapter IV, we have to show that as $T \rightarrow \infty$, for all $x > 0$

$$(\mathbf{RT2}) \quad M P(Z_{T,1} > x) \rightarrow c^{\alpha_{\min}} \frac{1+\beta}{2} [a_1^{\alpha_{\min}} t_1 + a_2^{\alpha_{\min}} (t_2 - t_1)] x^{-\alpha_{\min}},$$

$$(\mathbf{LT2}) \quad M P(Z_{T,1} \leq x) \rightarrow c^{\alpha_{\min}} \frac{1-\beta}{2} [a_1^{\alpha_{\min}} t_1 + a_2^{\alpha_{\min}} (t_2 - t_1)] x^{-\alpha_{\min}},$$

$$(\mathbf{VA2}) \quad \lim_{\epsilon \downarrow 0} \limsup_{T \rightarrow \infty} M \operatorname{Var}(Z_{T,1} \mathbf{1}_{\{|Z_{T,1}| < \epsilon\}}) = 0.$$

We will only give the proof of **(RT2)**. The proof of **(LT2)** is analogous and **(VA2)** follows in the same way as in the proof of Proposition 5.A.10. Let $\epsilon_T \rightarrow 0$ satisfy (A.4). Since we know from Lemma 5.A.3 that as $T \rightarrow \infty$

$$M P(|xi_{Tt_j} - \mu_{Tt_j}| > \epsilon_T \mu_{Tt_j}) = o(1), \quad j = 1, 2,$$

we only have to consider the intersection of $\{Z_{T,1} > x\}$ with the event

$$\Theta_T = \{|xi_{Tt_j} - \mu_{Tt_j}| \leq \epsilon_T \mu_{Tt_j}, j = 1, 2\}.$$

For $\delta \in (0, 1)$ we have

$$\begin{aligned} P(Z_{T,1} > x, \Theta_T) & \leq P\left(a_1 \left[\sum_{k=1}^{\xi_{Tt_1}} J_k - \sum_{k=1}^{[\mu_{Tt_1}]} J_k \right] > \frac{\delta}{2} x b(MT), \Theta_T\right) \\ & + P\left(a_2 \left[\sum_{k=\xi_{Tt_1}+1}^{\xi_{Tt_2}} J_k - \sum_{k=[\mu_{Tt_1}]+1}^{[\mu_{Tt_2}]} J_k \right] > \frac{\delta}{2} x b(MT), \Theta_T\right) \\ & + P\left(a_1 \sum_{k=1}^{[\mu_{Tt_1}]} J_k + a_2 \sum_{k=[\mu_{Tt_1}]+1}^{[\mu_{Tt_2}]} J_k > (1-\delta) x b(MT)\right). \end{aligned}$$

The first and second probabilities on the right-hand side are $o(M^{-1})$ by Lemma 5.A.7. By Lemma 5.A.11, the third probability is asymptotically equivalent to

$$P\left(a_1 \sum_{k=1}^{[\mu_{Tt_1}]} J_k > (1-\delta)x b(MT)\right) + P\left(a_2 \sum_{k=1}^{[\mu_{Tt_2}]-[\mu_{Tt_1}]} J_k > (1-\delta)x b(MT)\right).$$

We apply Corollary 5.C.2 to M times the above probabilities and find the estimate

$$\limsup_{T \rightarrow \infty} M P(Z_{T,1} > x) \leq \begin{cases} o(1) & \text{if } \alpha_{\text{on}} > \alpha_{\text{off}}, \\ r_{\text{on}}^{\alpha_{\text{on}}} \ell \mu^{-1} C_{1,2} x^{-\alpha_{\text{on}}} & \text{if } \alpha_{\text{on}} = \alpha_{\text{off}}, \\ r_{\text{on}}^{\alpha_{\text{on}}} \mu^{-1} C_{1,2} x^{-\alpha_{\text{on}}} & \text{if } \alpha_{\text{on}} < \alpha_{\text{off}}, \end{cases} \quad (\text{A.29})$$

where

$$C_{1,2} = a_1^{\alpha_{\text{on}}} t_1 + a_2^{\alpha_{\text{on}}} (t_2 - t_1).$$

The \geq -analogue of (A.29) can be found in the same way as in the proof of Proposition 5.A.10. This completes the proof. If a_i are not both positive, then the right-hand side of (A.29) is different of course. The proof, however, is similar. \heartsuit

The limit process has independent increments and stable marginal distributions. From the proof of Proposition 5.A.12 it follows that it has stationary increments. Therefore, the limit process is stable Lévy motion.

Remark Notice that the independent increments of the limit process arise in a very natural way. We used the decomposition (A.1) and showed that the contributions of the 0th renewal interval and the remainder term $III(T)$ vanish in the limit. We are left with $II_1(T)$, which is a random sum of iid heavy tailed random variables. Basically, the counting process ξ_T can be replaced by its mean μ_T , since values far from the mean are asymptotically negligible. This means that for large T , $II_1(T)$ behaves like a sum of iid random variables and, hence, has independent increments.

5.B Proof of Theorem 5.5: Fast Growth

Set

$$d_T = [MT^{3-\alpha_{\min}} L_{\min}(T)]^{1/2},$$

and

$$G_T^{(m)} = \int_0^T (W_u^{(m)} - EW_u^{(m)}) du.$$

By Lemma 5.4, Fast Growth Condition 2 is equivalent to $o(d_T) = T$, since

$$\frac{d_T}{T} = [M T \bar{F}_{\min}(T)]^{1/2}.$$

Now

$$V_T(t) = d_T^{-1} \sum_{m=1}^M G_{Tt}^{(m)}.$$

Convergence of the one-dimensional distributions is established in the following proposition.

Proposition 5.B.1 *Let σ_0 be as in Theorem 5.5, $H = (3 - \alpha_{\min})/2$ and B_H be fractional Brownian motion as defined in Section 3.2.6. Then, for every $t \geq 0$,*

$$V_T(t) \xrightarrow{d} N(0, \sigma_0^2 t^{3-\alpha_{\min}}) \stackrel{d}{=} \sigma_0 B_H(t). \quad (\text{B.1})$$

Proof. In Petrov [72], Theorem 4.2, we find the following necessary and sufficient conditions for (B.1): as $T \rightarrow \infty$

$$(\text{PRO}) \quad M P(|G_{Tt}| \geq \epsilon d_T) \rightarrow 0 \quad \text{for all } \epsilon > 0,$$

$$(\text{VAR}) \quad M d_T^{-2} \text{Var}(G_{Tt} 1_{[|G_{Tt}| \leq \tau d_T]}) \rightarrow \sigma_0^2 t^{3-\alpha_{\min}} \quad \text{for some } \tau > 0,$$

$$(\text{EXP}) \quad M d_T^{-1} E(G_{Tt} 1_{[|G_{Tt}| \leq \tau d_T]}) \rightarrow 0 \quad \text{for some } \tau > 0.$$

(**PRO**) follows from the fact that $P(|G_T| \geq \epsilon d_T) = 0$ for large T , since $T = o(d_T)$ and $|G_T| \leq T$ a.s. From the same observation it follows that the indicator functions in (**VAR**) and (**EXP**) are equal to one for T large enough. The proof of (**EXP**) now follows from $EG_{Tt} = 0$. Finally, (**VAR**) follows from (5.9). \heartsuit

Now, it is only a small step to prove convergence of the finite dimensional distributions of V_T . We only consider 2-dimensional convergence, since the general case is completely analogous. We have to show that, for $a_1, a_2 \in \mathbb{R}$ and $t_2 \geq t_1 \geq 0$,

$$d_T^{-1} \sum_{m=1}^M [a_1 G_{Tt_1}^{(m)} + a_2 G_{Tt_2}^{(m)}] \xrightarrow{d} a_1 \sigma_0 B_H(t_1) + a_2 \sigma_0 B_H(t_2).$$

Again using Theorem 4.2 in Petrov [72], one has to show the statements corresponding to (**PRO**), (**VAR**) and (**EXP**) above. The proofs of (**PRO**) and (**EXP**) follow in the same way as in Proposition 5.B.1. For (**VAR**) we have to show that for $t_1 \leq t_2$, as $T \rightarrow \infty$

$$\begin{aligned} M d_T^{-2} \text{Cov}(G_{Tt_1}, G_{Tt_2}) &\rightarrow \frac{\sigma_0^2}{2} [t_1^{2H} + t_2^{2H} - (t_2 - t_1)^{2H}] \quad (\text{B.2}) \\ &= \text{Cov}(\sigma_0 B_H(t_1), \sigma_0 B_H(t_2)). \end{aligned}$$

But this follows from

$$\text{Cov}(G_{Tt_1}, G_{Tt_2}) = \frac{1}{2} [\text{Var}(G_{Tt_1}) + \text{Var}(G_{Tt_2}) - \text{Var}(G_{Tt_2} - G_{Tt_1})],$$

the fact that G has stationary increments and (5.9). Therefore, the finite dimensional distributions of V_T converge to those of fractional Brownian motion.

It remains to show that the family of stochastic processes V_T is tight in $\mathbb{C}[0, K]$ for any fixed $K > 0$. We will show that for small $u > 0$ and $T \geq T^*$

$$E \left| d_T^{-1} \sum_{m=1}^M G_{Tu}^{(m)} \right|^2 \leq (\text{const}) u^{1+\varepsilon},$$

for some small $\varepsilon > 0$. Then Theorem 12.3 in Billingsley [7] gives the result.

According to (5.9) we have for T large enough

$$E \left| d_T^{-1} \sum_{m=1}^M G_{Tu}^{(m)} \right|^2 = \frac{M}{d_T^2} EG_{Tu}^2 = \frac{EG_{Tu}^2}{T^{3-\alpha_{\min}} L_{\min}(T)} \leq 2 \sigma_0^2 \frac{EG_{Tu}^2}{EG_T^2}.$$

By (5.9) we know that the function EG_x^2 is regularly varying with index $3 - \alpha_{\min}$. Using the left-hand inequality of Proposition 5.D.1, with $x = 1/u$, $t = Tu$ and some small $\varepsilon > 0$ such that $3 - \alpha_{\min} - 2\varepsilon > 1$, gives that there is a fixed t_0 such that for $u \leq 1$ and $Tu \geq t_0$

$$\frac{EG_{Tu}^2}{EG_T^2} < \frac{1}{1-\varepsilon} u^{3-\alpha_{\min}-\varepsilon}.$$

For $Tu < t_0$ we have for large enough T

$$\begin{aligned} \frac{EG_{Tu}^2}{T^{3-\alpha_{\min}} L_{\min}(T)} &\leq \frac{(Tu)^2}{T^{3-\alpha_{\min}} L_{\min}(T)} \\ &\leq \frac{(Tu)^{1+\varepsilon} t_0^{1-\varepsilon}}{T^{3-\alpha_{\min}} L_{\min}(T)} \\ &= \frac{T^{1-(3-\alpha_{\min}-\varepsilon)}}{L_{\min}(T)} t_0^{1-\varepsilon} u^{1+\varepsilon} \\ &\leq t_0^{1-\varepsilon} u^{1+\varepsilon}. \end{aligned}$$

Since $3 - \alpha_{\min} - \varepsilon > 1 + \varepsilon$ we have for T large enough and $u \leq 1$

$$E \left| d_T^{-1} \sum_{m=1}^M G_{Tu}^{(m)} \right|^2 \leq \max(2\sigma_0^2/(1-\varepsilon), t_0^{1-\varepsilon}) u^{1+\varepsilon}.$$

This completes the proof. \heartsuit

5.C Large deviations of heavy-tailed sums

We present a large deviation result which is frequently used in the proof of Theorem 5.5. Let $(\tilde{X}_k, k \geq 1)$ be iid random variables with distribution F such

that

$$\bar{F}(x) = x^{-\alpha_1} L_1(x), \quad x > 0, \quad \text{for some } \alpha_1 > 0 \text{ and } L_1 \text{ slowly varying, (C.1)}$$

and denote by

$$S_n = \tilde{X}_1 + \cdots + \tilde{X}_n, \quad n \geq 1,$$

the corresponding partial sums. Define

$$\mu_2(x) = x^{-2} \int_{|u| \leq x} u^2 dF(u).$$

The following large deviation result is due to Cline and Hsing [18].

Proposition 5.C.1 *Let $\beta_n \rightarrow \infty$ such that $S_n/\beta_n \xrightarrow{P} 0$. Suppose $B_n \subset [\beta_n, \infty)$. If the condition*

$$\lim_{n \rightarrow \infty} \sup_{x \in B_n} \left| n \mu_2(x) \ln(n \bar{F}(x)) \right| = 0 \quad (\text{C.2})$$

holds, then

$$\lim_{n \rightarrow \infty} \sup_{x \in B_n} \left| \frac{P(S_n > x)}{n \bar{F}(x)} - 1 \right| = 0. \quad (\text{C.3})$$

♡

Remark. Writing $M_n = \max_{k=1, \dots, n} \tilde{X}_k$ for the partial maxima of the \tilde{X} -sequence, we see that we can replace $n \bar{F}(x)$ in (C.3) by $P(M_n > x)$. This means that the large deviation $\{S_n > x\}$ is essentially due to the event $\{M_n > x\}$.

A consequence is the following result.

Corollary 5.C.2 *In addition to (C.1) assume that $E\tilde{X} = 0$ and either*

$$F(-x) = x^{-\alpha_2} L_2(x), \quad x > 0, \quad \alpha_2 > \alpha_1, \quad \alpha_1 \in (1, 2), \quad L_2 \text{ slowly varying,}$$

or

$$F(-x) = 0 \quad \text{for } x > x_0, \text{ some } x_0 > 0.$$

Then (C.3) holds with $\beta_n = a_n h_n$ and $B_n = [\beta_n, \infty)$ where (h_n) is any sequence with $h_n \uparrow \infty$ and (a_n) satisfies $n \bar{F}(a_n) \sim 1$.

Proof. Since $(a_n^{-1} S_n)$ weakly converges to an α_1 -stable distribution relation $\beta_n^{-1} S_n \xrightarrow{P} 0$ is immediate. Moreover, by Karamata's theorem

$$\mu_2(x) \leq (\text{const}) P(|\tilde{X}| > x), \quad x > 0,$$

and so (C.2) is satisfied since

$$n \mu_2(x) \ln(n P(|\tilde{X}| > \beta_n)) \leq (\text{const}) n P(|\tilde{X}| > \beta_n) \ln(n P(|\tilde{X}| > \beta_n)) \rightarrow 0.$$

This concludes the proof. ♡

5.D Bounds for regularly varying functions

Let $U(x)$ be a regularly varying function with index $\rho \in \mathbb{R}$, i.e. for $x > 0$

$$\lim_{t \rightarrow \infty} \frac{U(tx)}{U(t)} = x^\rho.$$

The following result can be found in Resnick [77], Proposition 0.8 (ii).

Proposition 5.D.1 *Take $\varepsilon > 0$. Then there is a fixed t_0 such that for $x \geq 1$ and $t \geq t_0$*

$$(1 - \varepsilon) x^{\rho - \varepsilon} < \frac{U(tx)}{U(t)} < (1 + \varepsilon) x^{\rho + \varepsilon}.$$

♡

In Bingham et al. [9] these bounds are called the *Potter bounds*.

Six

Extremal Behavior of ON-Periods

In this chapter, we will use the framework of the ON/OFF model to study the number of exceedances by the sequence of ON-periods $(X_i^{(m)})$. We use a threshold x_T which has an infinite limit in T and adopt the notation introduced in Section 5.1.1. Instead of (5.1), however, we assume that as $x \rightarrow \infty$

$$\overline{F}_{\text{on}}(x) = x^{-\alpha} L(x) \quad \text{and} \quad \overline{F}_{\text{off}}(x) = o(\overline{F}_{\text{on}}(x)), \quad (6.1)$$

where $\alpha > 1$ and L is slowly varying at infinity. Hence, \overline{F}_{on} is regularly varying and F_{off} has a lighter tail than F_{on} . As before, ON- and OFF-periods have means μ_{on} and μ_{off} , respectively, and $\mu = \mu_{\text{on}} + \mu_{\text{off}}$.

The number of exceedances is counted up to time T , which means that we only consider the completed ON-periods $X_1^{(m)}, \dots, X_{\xi_T^{(m)}-1}^{(m)}$ and $\min(T - S_{\xi_T^{(m)}-1}^{(m)}, X_{\xi_T^{(m)}}^{(m)})$, for $m = 1, \dots, M$. The total number of exceedances up to time T is given by

$$A_T = \sum_{m=1}^M \left[\sum_{i=1}^{\xi_T^{(m)}-1} 1_{[x_T, \infty)}(X_i^{(m)}) + 1_{[x_T, \infty)}(\min(T - S_{\xi_T^{(m)}-1}^{(m)}, X_{\xi_T^{(m)}}^{(m)})) \right]. \quad (6.2)$$

We are interested in the limit distribution of A_T as $T \rightarrow \infty$. It is clear that the ON-periods $X_i^{(m)}$ and the counting process $\xi_T^{(m)}$ are heavily dependent. From the definition of $\xi_T^{(m)}$ in (5.3) it follows that

$$X_1^{(m)}, \dots, X_{\xi_T^{(m)}-1}^{(m)}, \min(T - S_{\xi_T^{(m)}-1}^{(m)}, X_{\xi_T^{(m)}}^{(m)}) \leq T, \quad m = 1, \dots, M.$$

This implies that the threshold x_T must be less than T to obtain a non-degenerate limit for A_T .

We study the case $M = M_T$ and distinguish between fast and slow growth of M . Recall the Slow- and Fast Growth Condition from Section 5.2.3, where b is the quantile function of F_{on} . If M satisfies the Slow Growth Condition we obtain a Poisson limit for A_T , for certain thresholds x_T . The intuition behind the proof is as follows. Define

$$A_T^{\text{lo}} = \sum_{m=1}^M \sum_{i=1}^{\xi_T^{(m)} - 1} 1_{[x_T, \infty)}(X_i^{(m)}) \quad \text{and} \quad A_T^{\text{up}} = \sum_{m=1}^M \sum_{i=1}^{\xi_T^{(m)}} 1_{[x_T, \infty)}(X_i^{(m)}). \quad (6.3)$$

Observe that for all ω ,

$$A_T^{\text{lo}} \leq A_T \leq A_T^{\text{up}}. \quad (6.4)$$

As in the proof of Theorem 5.5, with Lévy motion in the limit, it is possible to replace all $\xi_T^{(m)}$ by their means μ_T . For large T , the distributions of A_T^{lo} and A_T^{up} are approximately binomial with $M[T/\mu]$ trials and success probability $\bar{F}_{\text{on}}(x_T)$. The Poisson approximation to the binomial distribution then guarantees $A_T^{\text{lo}, \text{up}} \xrightarrow{d} \text{Poi}(1/\mu)$ (and by (6.4) also $A_T \xrightarrow{d} \text{Poi}(1/\mu)$) if

$$EA_T^{\text{lo}, \text{up}} \approx M \frac{T}{\mu} \bar{F}_{\text{on}}(x_T) \rightarrow \frac{1}{\mu} \quad \text{as } T \rightarrow \infty.$$

This implies $x_T \sim b(MT)$. Moreover, since the Slow Growth Condition is satisfied, $x_T = o(T)$.

For $1 < \alpha < 2$, there is a clear connection with Theorem 5.5. Let x_n be such that $n\bar{F}_{\text{on}}(x_n) \sim 1$, as $n \rightarrow \infty$, i.e. $x_n \sim b(n)$. Since \bar{F}_{on} is regularly varying with tail parameter α , as $n \rightarrow \infty$

$$x_n^{-1} \sum_{i=1}^n (X_i - \mu_{\text{on}}) \xrightarrow{d} S_\alpha, \quad (6.5)$$

where S_α is a totally skewed to the right α -stable distribution. Also, as $n \rightarrow \infty$

$$x_n^{-1} \max(X_1, \dots, X_n) \xrightarrow{d} \Phi_\alpha, \quad (6.6)$$

where $P(\Phi_\alpha \leq y) = \exp\{-y^{-\alpha}\}$ is the Fréchet distribution. Finally, due to the Poisson limit theorem, the definition of x_n guarantees that as $n \rightarrow \infty$

$$\sum_{i=1}^n 1_{[x_n, \infty)}(X_i) \xrightarrow{d} \text{Poi}(1). \quad (6.7)$$

For $\alpha \geq 2$, (6.6) and (6.7) still hold, but (6.5) translates into the Central Limit Theorem where the normalization is $\sqrt{n\text{Var}(X)}$ (if $\alpha = 2$ and $EX^2 = \infty$ a different normalization has to be used).

If M satisfies the Fast Growth Condition the situation is more complicated. In order to apply the Poisson limit theorem, we must have $EA_T \sim (\text{const})$. Since $\xi_T^{(m)}$ cannot be replaced by μ_T , there is no straightforward method of

calculating $E\hat{A}_T$. However, notice that $\xi_T^{(m)}$ is a stopping time with respect to the filtration

$$\mathcal{F}_n^{(m)} = \sigma(D^{(m)}, X_1^{(m)}, Y_1^{(m)}, \dots, X_n^{(m)}, Y_n^{(m)}), \quad n \geq 1.$$

Therefore we also include the ON-periods $X_{\xi_T^{(m)}}^{(m)}$, which allows us to use Wald's identity for the expectation of random sums (see Resnick [78], Section 1.8.1 or Gut [43], Theorem I5.3). We consider the following number of exceedances:

$$\hat{A}_T = \sum_{m=1}^M \sum_{i=1}^{\xi_T^{(m)}} 1_{[x_T, T)}(X_i^{(m)}). \quad (6.8)$$

We need the restriction $X_i^{(m)} < T$ to obtain a non-degenerate limit for \hat{A}_T . Using Wald's identity, we have

$$\begin{aligned} E\hat{A}_T &= M \frac{T}{\mu} [\bar{F}_{\text{on}}(x_T) - \bar{F}_{\text{on}}(T)] \\ &= M \frac{T}{\mu} \bar{F}_{\text{on}}(T) \left[\frac{\bar{F}_{\text{on}}(x_T)}{\bar{F}_{\text{on}}(T)} - 1 \right]. \end{aligned} \quad (6.9)$$

Since the Fast Growth Condition holds, $MT\bar{F}_{\text{on}}(T) \rightarrow \infty$. To have $E\hat{A}_T \sim (\text{const})$, we must choose x_T such that $\bar{F}_{\text{on}}(x_T) \sim \bar{F}_{\text{on}}(T)$. Using the monotonicity and regular variation (see Resnick [77], Proposition 0.8 (iii)) of \bar{F}_{on} , we see that $x_T \sim T$ must hold. Moreover, since $x_T < T$, we can write $x_T = T - a_T$ where a_T is a positive sequence satisfying $a_T = o(T)$. We obtain a Poisson limit for \hat{A}_T by balancing M and a_T such that $E\hat{A}_T \sim (\text{const})$.

It appears that, if M satisfies the Fast Growth Condition, the space where exceedances can occur must be chosen very small in order to obtain a non-degenerate limit. Since the ON-periods cannot be larger than T , there will be more and more of them near T as M increases. The faster M grows, the smaller the region $[T - a_T, T)$ has to be to ensure a non-degenerate limiting number of exceedances. A large M must be compensated by a small a_T . Thus, for M very fast, Poisson convergence of \hat{A}_T is due to the number of ON-periods with lengths which are practically indistinguishable from T .

In Section 6.1 we consider the case when M satisfies the Slow Growth Condition. We use the theory of point processes to show that the number of exceedances of the threshold $x_T = b(MT) = o(T)$ converges to a Poisson Random Measure. We apply this result to show that the number of exceedances up to time Tt converges weakly to a homogeneous Poisson process in $(\mathbb{D}[0, \infty), J_1)$ as $T \rightarrow \infty$.

In Section 6.2 we consider the case when M satisfies the Fast Growth Condition, using point processes as in Section 6.1. We obtain convergence to a Poisson Random Measure for a threshold $x_T = T - a_T$ with $a_T = o(T)$ under a balancing condition on M and a_T .

In Section 6.3 we show that

$$\sum_{m=1}^M \sum_{i=1}^{\xi_T^{(m)}} 1_{[x_T, T]}(X_i^{(m)})$$

obeys the Central Limit Theorem under the condition that

$$MT[\bar{F}_{\text{on}}(x_T) - \bar{F}_{\text{on}}(T)] \rightarrow \infty.$$

This ensures an increasing number of exceedances. Here, M does not have to satisfy a particular growth condition.

In Sections 6.A-6.C proofs of the results in Sections 6.1-6.3 can be found. Section 6.D contains theorems on weak convergence of a sequence of point processes to a simple limit. Finally, Section 6.E contains a result needed in the proofs mentioned above. The analysis in this chapter was done in Stegeman [94].

6.1 Slow Growth

Here we consider the case where M is either fixed or $M = M_T \rightarrow \infty$ is a non-decreasing integer-valued function satisfying the Slow Growth Condition, i.e.

$$\text{Slow Growth Condition: } \lim_{T \rightarrow \infty} \frac{b(MT)}{T} = 0 \iff \lim_{T \rightarrow \infty} M T \bar{F}_{\text{on}}(T) = 0,$$

where b is the quantile function of F_{on} . We derive the limit distribution of A_T (see (6.2)), which is the total number of exceedances, up to time T , by completed and running ON-periods of all M sources. The threshold x_T is such that as $T \rightarrow \infty$

$$M T \bar{F}_{\text{on}}(x_T) \sim 1. \quad (6.10)$$

Notice that $x_T \sim b(MT)$. Moreover, since the Slow Growth Condition holds, we have $x_T = o(T)$.

We use the theory of point processes to derive the limit distribution of A_T . Recall the definition of the stationary renewal sequence $(T_n^{(m)})$ from (5.2). For $T > 0$ we define

$$N_T^{\text{lo}} = \sum_{m=1}^M \sum_{i=1}^{\infty} \varepsilon_{T_i^{(m)}/T, X_i^{(m)}/x_T} \quad \text{and} \quad N_T^{\text{up}} = \sum_{m=1}^M \sum_{i=1}^{\infty} \varepsilon_{T_{i-1}^{(m)}/T, X_i^{(m)}/x_T}, \quad (6.11)$$

which are point processes on the state space $E = [0, \infty) \times (0, \infty]$. For $(u, v) \in E$ and $C \times D \subset E$, the Dirac measure ε is defined by

$$\varepsilon_{(u,v)}(C \times D) = \begin{cases} 1 & \text{if } (u, v) \in C \times D, \\ 0 & \text{if } (u, v) \notin C \times D. \end{cases}$$

For fixed T , the point process N_T^{lo} resembles the point process of exceedances (see Embrechts et al. [28], Example 5.1.3). The difference is that N_T^{lo} also takes into account the ‘real’ times, in the form of the renewal sequences $(T_n^{(m)})$, at which the exceedances occur. An additional complication is that the $X_i^{(m)}$ and $T_n^{(m)}$ are heavily dependent. Let A_T^{lo} and A_T^{up} be defined by (6.3). Notice that

$$\begin{aligned} N_T^{\text{lo}}([0, 1) \times [1, \infty)) &= \sum_{m=1}^M \sum_{i=1}^{\xi_T^{(m)} - 1} 1_{[x_T, \infty)}(X_i^{(m)}) = A_T^{\text{lo}}, \\ N_T^{\text{up}}([0, 1) \times [1, \infty)) &= \sum_{m=1}^M \sum_{i=1}^{\xi_T^{(m)}} 1_{[x_T, \infty)}(X_i^{(m)}) = A_T^{\text{up}}. \end{aligned}$$

Let $M_p(E)$ denote the space of all point measures defined on E , equipped with the vague topology (see Resnick [77], Section 3.4). We will show that both N_T^{lo} and N_T^{up} converge in distribution to a Poisson Random Measure (PRM) in $M_p(E)$. Recall that for a PRM N on E with mean measure ν , $N(A)$ has a Poisson distribution with expectation $\nu(A)$ for any Borel set $A \subset E$, and $N(A_1)$ and $N(A_2)$ are independent if A_1 and A_2 are disjoint.

Let \mathbb{L} denote Lebesgue measure. We have the following result.

Theorem 6.1 *Let F_{on} and F_{off} satisfy (6.1). Suppose M satisfies the Slow Growth Condition and x_T satisfies (6.10). Then, as $T \rightarrow \infty$,*

$$N_T^{\text{lo}} \xrightarrow{d} N \quad \text{and} \quad N_T^{\text{up}} \xrightarrow{d} N \quad \text{in } M_p(E),$$

where N is PRM with mean measure $\mu^{-1}\mathbb{L} \times \nu$ and, for $0 < a < b$,

$$\nu[a, b) = \int_a^b \alpha t^{-(\alpha+1)} dt.$$

♡

Notice that, since $\mu^{-1}(\mathbb{L} \times \nu)([0, 1) \times [1, \infty)) = \mu^{-1}$, Theorem 6.1 implies that as $T \rightarrow \infty$

$$A_T^{\text{lo}} \xrightarrow{d} \text{Poi}(\mu^{-1}) \quad \text{and} \quad A_T^{\text{up}} \xrightarrow{d} \text{Poi}(\mu^{-1}).$$

By (6.4) also $A_T \xrightarrow{d} \text{Poi}(\mu^{-1})$. The proof of Theorem 6.1 can be found in Section 6.A. In Figure 6.1 the exceedances of simulated series of ON-periods are depicted for $M_T = \log(T)$.

Next, we use the convergence of N_T^{lo} and N_T^{up} to a PRM to show that the number of exceedances up to time Tt by the double array $X_i^{(m)}$ converges to a homogeneous Poisson process as $T \rightarrow \infty$. Again we assume that there are M iid sources, where M is either fixed or increasing as $T \rightarrow \infty$. In the latter case we also suppose that M satisfies the Slow Growth Condition. The total number

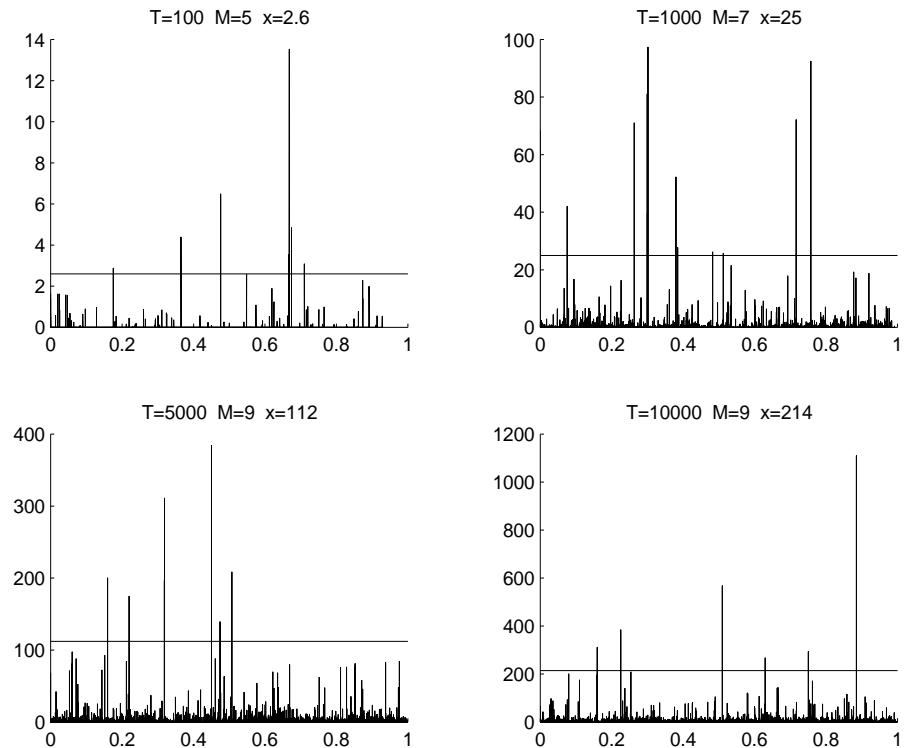


Figure 6.1: Illustration of the exceedances of simulated series of ON-periods under the Slow Growth Condition. Both F_{on} and F_{off} are Pareto with tail parameters $\alpha = 1.2$ and 1.8 and means $\mu_{\text{on}} = 3$ and $\mu_{\text{off}} = 5$, respectively. Here $M_T = \log(T)$ and x_T is chosen such that $M_T \bar{F}_{\text{on}}(x_T)/\mu \sim 10$. It is clear from the plots that the number of exceedances stabilizes as T becomes large. Moreover, the exceedances do not appear in clusters, which is what one might expect in the case of a Poisson limit.

of exceedances up to time Tt is given by

$$A_T(t) = \sum_{m=1}^M \left[\sum_{i=1}^{\xi_{Tt}^{(m)} - 1} 1_{[x_T, \infty)}(X_i^{(m)}) + 1_{[x_T, \infty)}(\min(T - S_{\xi_{Tt}^{(m)} - 1}^{(m)}, X_{\xi_{Tt}^{(m)}}^{(m)})) \right].$$

We assume that the threshold x_T satisfies (6.10). Notice that the processes A_T contain more information on the times when exceedances occur, while N_T^{lo} and N_T^{up} hold more information on their sizes.

We have the following result.

Proposition 6.2 *Let $(Z(t), t \geq 0)$ be a homogeneous Poisson process with intensity μ^{-1} . If M satisfies the Slow Growth Condition and x_T satisfies (6.10), then as $T \rightarrow \infty$*

$$(A_T(t), t \geq 0) \xrightarrow{d} (Z(t), t \geq 0),$$

where \xrightarrow{d} denotes weak convergence in the space $\mathbb{D}[0, \infty)$, equipped with the J_1 -topology.

Proof. From Theorem 6.1 it follows that the families of point processes $N_T^{\text{lo}}(\cdot \times [1, \infty))$ and $N_T^{\text{up}}(\cdot \times [1, \infty))$ converge weakly in $M_p([0, \infty))$ to $N(\cdot \times [1, \infty))$. Define for $t \geq 0$

$$A_T^{\text{lo}}(t) = \sum_{m=1}^M \sum_{i=1}^{\xi_{Tt}^{(m)} - 1} 1_{[x_T, \infty)}(X_i^{(m)}) \quad \text{and} \quad A_T^{\text{up}}(t) = \sum_{m=1}^M \sum_{i=1}^{\xi_{Tt}^{(m)}} 1_{[x_T, \infty)}(X_i^{(m)}).$$

Since

$$A_T^{\text{lo}}(t) = N_T^{\text{lo}}([0, t) \times [1, \infty)) \quad \text{and} \quad A_T^{\text{up}}(t) = N_T^{\text{up}}([0, t) \times [1, \infty)),$$

A_T^{lo} and A_T^{up} are the corresponding families of cumulative processes. Notice that for all $t \geq 0$ and all ω

$$A_T^{\text{lo}}(t) \leq A_T(t) \leq A_T^{\text{up}}(t).$$

The result now follows from these bounds, Daley and Vere-Jones [22], Lemma 9.1.X, and the fact that $N([0, t) \times [1, \infty))$ is a homogeneous Poisson process with intensity μ^{-1} . \heartsuit

Remark An alternative proof of Proposition 6.2 is as follows. Convergence of the finite-dimensional distributions follows from Theorem 6.1 by considering sets $[0, t_j) \times [1, \infty)$ with $t_j \geq 0$, $j = 1, \dots, n$, $n \geq 1$. For J_1 -tightness in $\mathbb{D}[0, k]$, $k > 0$, the conditions formulated by Gut and Janson [44] can be checked.

6.2 Fast Growth

In this section we derive the limit distribution of the number of exceedances if M satisfies the Fast Growth Condition, i.e.

$$\text{Fast Growth Condition: } \lim_{T \rightarrow \infty} \frac{b(MT)}{T} = \infty \iff \lim_{T \rightarrow \infty} M T \bar{F}_{\text{on}}(T) = \infty.$$

We consider the number of exceedances given by \hat{A}_T in (6.8), where a_T is a positive sequence satisfying $a_T = o(T)$. We have

$$E\hat{A}_T = M \frac{T}{\mu} \bar{F}_{\text{on}}(T) \left[\frac{\bar{F}_{\text{on}}(T - a_T)}{\bar{F}_{\text{on}}(T)} - 1 \right].$$

For any M satisfying the Fast Growth Condition we choose a_T such that $E\hat{A}_T \sim (\text{const})$. An immediate consequence of regular variation of \bar{F}_{on} (see Resnick [77], Proposition 0.8 (iii)) is that $\bar{F}_{\text{on}}(T - a_T)/\bar{F}_{\text{on}}(T) = 1 + o(1)$, but in order to have $E\hat{A}_T \sim (\text{const})$ we need to know the rate at which the $o(1)$ term converges to zero. This requires to impose a second order regular variation condition on \bar{F}_{on} . In practice, however, such a condition often cannot be verified. Therefore, and since we are only interested in a qualitative characterization concerning the limit of \hat{A}_T , we will from now on assume that for some x_0 and some constant $c > 0$,

$$\bar{F}_{\text{on}}(x) = c x^{-\alpha} \quad \text{for } x \geq x_0. \quad (6.12)$$

As before, $\alpha > 1$ and $\bar{F}_{\text{off}}(x) = o(\bar{F}_{\text{on}}(x))$. Then, for T large enough

$$E\hat{A}_T = M \frac{T}{\mu} \bar{F}_{\text{on}}(T) \left[\left(1 - \frac{a_T}{T}\right)^{-\alpha} - 1 \right].$$

Using a first order Taylor expansion, we obtain

$$\begin{aligned} E\hat{A}_T &= M \frac{T}{\mu} \bar{F}_{\text{on}}(T) \left[\alpha \frac{a_T}{T} + o\left(\frac{a_T}{T}\right) \right] \\ &= \frac{\alpha}{\mu} M a_T \bar{F}_{\text{on}}(T) + o(M a_T \bar{F}_{\text{on}}(T)). \end{aligned}$$

Therefore, we choose a_T such that

$$a_T = o(T) \quad \text{and} \quad M a_T \bar{F}_{\text{on}}(T) \sim 1. \quad (6.13)$$

In this way,

$$\lim_{T \rightarrow \infty} E\hat{A}_T = \frac{\alpha}{\mu}.$$

We will use weak convergence of point processes to a PRM to show that $\hat{A}_T \xrightarrow{d} \text{Poi}(\alpha\mu^{-1})$. Define the point processes

$$\hat{N}_T = \sum_{m=1}^M \hat{N}_T^{(m)},$$

where

$$\widehat{N}_T^{(m)} = \sum_{i=1}^{\infty} \varepsilon_{T_{i-1}^{(m)}/T, (T-X_i^{(m)})/a_T}.$$

Let $E = [0, \infty) \times (0, 1]$ be the state space of \widehat{N}_T . Notice that

$$\widehat{N}_T([0, 1) \times (0, 1]) = \sum_{m=1}^M \sum_{i=1}^{\xi_T^{(m)}} 1_{[T-a_T, T)}(X_i^{(m)}) = \widehat{A}_T.$$

In the following theorem we show that \widehat{N}_T weakly converges to a PRM.

Theorem 6.3 *Assume \overline{F}_{on} satisfies (6.12). If M satisfies the Fast Growth Condition and a_T satisfies (6.13), then as $T \rightarrow \infty$*

$$\widehat{N}_T \xrightarrow{d} \widehat{N} \quad \text{in } M_p(E),$$

where \widehat{N} is PRM with mean measure $\alpha\mu^{-1}\mathbb{L} \times \mathbb{L}$. ♡

Since $\alpha\mu^{-1}(\mathbb{L} \times \mathbb{L})([0, 1) \times (0, 1]) = \alpha\mu^{-1}$, Theorem 6.3 implies that as $T \rightarrow \infty$

$$\widehat{A}_T = \sum_{m=1}^M \sum_{i=1}^{\xi_T^{(m)}} 1_{[T-a_T, T)}(X_i^{(m)}) \xrightarrow{d} \text{Poi}(\alpha\mu^{-1}).$$

The proof of Theorem 6.3 is contained in Section 6.B. In Figure 6.2 the exceedances of simulated series of ON-periods are depicted for $M_T = T$.

Remark In Figure 6.2, almost all exceedances in the plots with $T = 5000$ and $T = 10000$ are due to the ON-periods $X_{\xi_T^{(m)}}^{(m)}$. This can be seen as follows. Let $n \in \{1, \dots, \xi_T^{(m)} - 1\}$. By definition $T_{\xi_T^{(m)}}^{(m)} \leq T$ and hence the maximum length of an ON-period $X_n^{(m)}$ is $T - T_{n-1}^{(m)}$. For an ON-period $X_n^{(m)}$ exceeding x_T we must have $T - T_{n-1}^{(m)} \geq x_T$, which is equivalent to $T_{n-1}^{(m)} \leq T - x_T$. In the plots we normalized the time scale to the interval $[0, 1]$. Hence, for any ON-period $X_i^{(m)}$ exceeding x_T and starting at a time later than $(T - x_T)/T$, we have $i = \xi_T^{(m)}$. For the plots with $T = 5000$ and $T = 10000$ we have $(T - x_T)/T = 0.134$ and 0.077 , respectively. So for $T = 10000$ all 7 exceedances are due to the ON-periods $X_{\xi_T^{(m)}}^{(m)}$. For $T = 5000$ this is true for at least 11 of the 13 exceedances.

This raises the question as to whether the same holds for \widehat{A}_T , i.e. as $T \rightarrow \infty$, does

$$\sum_{m=1}^M 1_{[T-a_T, T)}(X_{\xi_T^{(m)}}^{(m)}) \xrightarrow{d} \text{Poi}(\alpha\mu^{-1}) \quad ?$$

We were, however, not able to obtain a conclusive answer.

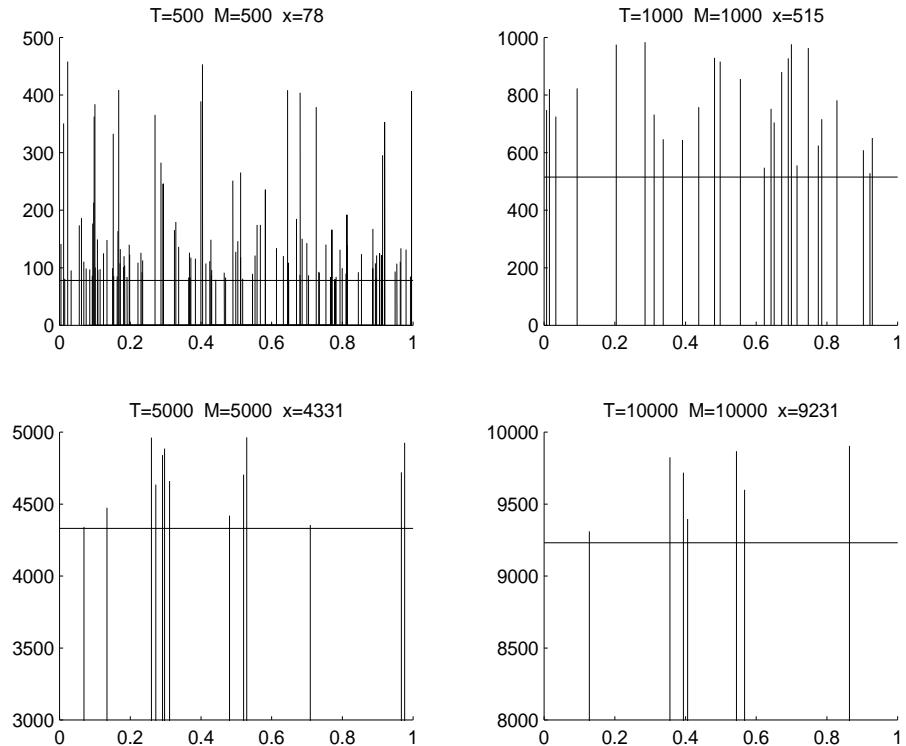


Figure 6.2: Illustration of the exceedances of simulated series of ON-periods under the Fast Growth Condition. Both F_{on} and F_{off} are Pareto with tail parameters $\alpha = 1.2$ and 1.8 and means $\mu_{\text{on}} = 3$ and $\mu_{\text{off}} = 5$, respectively. Here $M_T = T$ and x_T is chosen such that $MT[\bar{F}_{\text{on}}(x_T) - F_{\text{on}}(T)]/\mu \sim 10$. To make the plots more clear, only the ON-period lengths that lie in $[x_T, T]$ are depicted. As in the slow growth situation the number of exceedances stabilizes and there are no clusters, which is consistent with a Poisson limit.

The following result is analogous to Proposition 6.2. Since the proof is similar it is omitted. Define for $t \in [0, 1]$

$$\widehat{A}_T(t) = \sum_{m=1}^M \sum_{i=1}^{\xi_T^{(m)}} 1_{[T-a_T, T)}(X_i^{(m)}).$$

Proposition 6.4 *Assume \overline{F}_{on} satisfies (6.12). Let $(Z(t), t \geq 0)$ be a homogeneous Poisson process with intensity $\alpha\mu^{-1}$. If M satisfies the Fast Growth Condition and a_T satisfies (6.13), then as $T \rightarrow \infty$*

$$(\widehat{A}_T(t), t \geq 0) \xrightarrow{d} (Z(t), t \geq 0),$$

where \xrightarrow{d} denotes weak convergence in the space $\mathbb{D}[0, 1]$, equipped with the J_1 -topology. \heartsuit

6.3 A Central Limit Theorem

Here we show that the number of exceedances

$$\sum_{m=1}^M \sum_{i=1}^{\xi_T^{(m)}} 1_{[x_T, T]}(X_i^{(m)})$$

satisfies the Central Limit Theorem for certain thresholds $x_T \rightarrow \infty$. The number of sources $M = M_T \rightarrow \infty$ is a non-decreasing integer-valued function. Together, M and x_T have to satisfy

$$x_T \leq T \quad \text{and} \quad M T (\overline{F}_{\text{on}}(x_T) - \overline{F}_{\text{on}}(T)) \rightarrow \infty. \quad (6.14)$$

The latter is needed to have an increasing amount of exceedances (see (6.9)). Define

$$Z_T = \sum_{m=1}^M Z_T^{(m)} = \sum_{m=1}^M \left[\frac{\sum_{i=1}^{\xi_T^{(m)}} [1_{[x_T, T]}(X_i^{(m)}) - (\overline{F}_{\text{on}}(x_T) - \overline{F}_{\text{on}}(T))] }{[MT(\overline{F}_{\text{on}}(x_T) - \overline{F}_{\text{on}}(T))]^{1/2}} \right].$$

Our goal is to show that Z_T has a Gaussian limit as $T \rightarrow \infty$. In doing so we will use the central limit theorem. Since each $\xi_T^{(m)}$ is a stopping time, we have $E Z_T = 0$. Using Theorem I5.3 in Gut [43] and the fact that

$$\begin{aligned} \text{Var}(1_{[x_T, T]}(X_1)) &= (\overline{F}_{\text{on}}(x_T) - \overline{F}_{\text{on}}(T))(1 - (\overline{F}_{\text{on}}(x_T) - \overline{F}_{\text{on}}(T))) \\ &\sim \overline{F}_{\text{on}}(x_T) - \overline{F}_{\text{on}}(T), \end{aligned}$$

we obtain

$$\text{Var}(Z_T) = \frac{M E(\xi_T) \text{Var}(1_{[x_T, T]}(X_1))}{M T (\overline{F}_{\text{on}}(x_T) - \overline{F}_{\text{on}}(T))} \sim \mu^{-1}.$$

Notice that, if we replaced each $\xi_T^{(m)}$ by μ_T , Z_T would actually be a centered and normalized sum of M iid binomial random variables with $[T/\mu]$ trials and success probability $\overline{F}_{\text{on}}(x_T) - \overline{F}_{\text{on}}(T)$. Here, we do not need the assumption that \overline{F}_{on} is regularly varying. In fact, it suffices to have

$$\mu = \int_0^\infty (\overline{F}_{\text{on}}(s) + \overline{F}_{\text{off}}(s)) \, ds < \infty.$$

In this way, the existence of a stationary renewal sequence (T_n) , defined by (5.2), is guaranteed.

The following result shows that Z_T converges weakly to a normal distribution as $T \rightarrow \infty$.

Theorem 6.5 *Suppose $\mu < \infty$. If M and x_T satisfy (6.14), then as $T \rightarrow \infty$*

$$\sum_{m=1}^M Z_T^{(m)} \xrightarrow{d} N(0, \mu^{-1}).$$

♡

Notice that there is no distinction between fast and slow growth of M . The proof of Theorem 6.5 is presented in Section 6.C.

Appendix

6.A Proof of Theorem 6.1

We start with proving convergence of N_T^{lo} . We use Kallenberg [53] (see Theorem 6.D.1) and consider the class \mathcal{A} of sets

$$A = \bigcup_{j=1}^k [a_j, b_j) \times (c_j, d_j],$$

for $k \geq 1$, $[a_j, b_j) \times (c_j, d_j] \subset [0, \infty) \times (0, \infty]$, $j = 1, \dots, k$. We may assume that the sets $[a_j, b_j) \times (c_j, d_j]$ are mutually disjoint. We first show

$$P(N_T^{\text{lo}}(A) = 0) \rightarrow P(N(A) = 0).$$

Notice that

$$N_T^{\text{lo}}(A) = \sum_{j=1}^k \sum_{m=1}^M \sum_{i=\xi_{T a_j}^{(m)}}^{\xi_{T b_j}^{(m)}-1} 1_{(c_j x_T, d_j x_T]}(X_i^{(m)}).$$

Let $\epsilon_T \rightarrow 0$ satisfy

$$b(MT) = o(\epsilon_T T) \quad \text{and} \quad 1/\log(T) = o(\epsilon_T) \quad \text{as } T \rightarrow \infty. \quad (\text{A.1})$$

Define the event

$$B_T = \{|\xi_{T t}^{(m)} - \mu_{T t}| \leq \epsilon_T \mu_{T t}, t = a_j, b_j, j = 1, \dots, k, m = 1, \dots, M\}.$$

From Lemma 5.A.3 it follows that $P(B_T^c) = o(1)$. (Lemma 5.A.3 assumes $\alpha \in (1, 2)$. However, a close inspection of the proof shows that the result holds for any $\alpha > 1$.) Define

$$N_{\mu_T^{\pm}}^{\text{lo}}(A) = \sum_{j=1}^k \sum_{m=1}^M \sum_{i=[(1 \mp \epsilon_T) \mu_{T a_j}]}^{[(1 \pm \epsilon_T) \mu_{T b_j}] - 1} 1_{(c_j x_T, d_j x_T]}(X_i^{(m)}).$$

We have the inequalities

$$P(N_T^{\text{lo}}(A) = 0, B_T) \leq P(N_{\mu_T^-}^{\text{lo}}(A) = 0),$$

and

$$P(N_T^{\text{lo}}(A) = 0, B_T) \geq P(N_{\mu_T^+}^{\text{lo}}(A) = 0, B_T) \geq P(N_{\mu_T^+}^{\text{lo}}(A) = 0) - P(B_T^c).$$

We will show that

$$P(N_{\mu_T^+}^{\text{lo}}(A) = 0) \rightarrow P(N(A) = 0).$$

The proof for $P(N_{\mu_T}^{\text{lo}}(A) = 0)$ is analogous. First, we assume that the sets $[a_j, b_j]$ are mutually disjoint. This assumption is relaxed later on. Moreover, assume that

$$0 \leq a_1 < b_1 \leq a_2 < b_2 \leq \cdots \leq a_k < b_k < \infty.$$

Then, for large T , $[(1 + \epsilon_T)\mu_{Tb_j}] - 1 < [(1 - \epsilon_T)\mu_{Ta_{j+1}}]$, for $j = 1, \dots, k - 1$. Using (6.10), we have

$$\begin{aligned} & P(N_{\mu_T}^{\text{lo}}(A) = 0) \\ &= \prod_{j=1}^k \left[P(X_i/x_T \notin (c_j, d_j), i = [(1 - \epsilon_T)\mu_{Ta_j}], \dots, [(1 + \epsilon_T)\mu_{Tb_j}] - 1) \right]^M \\ &\sim \prod_{j=1}^k \left[1 - \bar{F}_{\text{on}}(c_j x_T) + \bar{F}_{\text{on}}(d_j x_T) \right]^{MT(b_j - a_j)\mu^{-1}} \\ &\sim \prod_{j=1}^k \left[1 - \frac{c_j^{-\alpha}}{MT} + \frac{d_j^{-\alpha}}{MT} \right]^{MT(b_j - a_j)\mu^{-1}} \\ &\sim \prod_{j=1}^k \left[1 - \frac{(b_j - a_j)(c_j^{-\alpha} - d_j^{-\alpha})}{\mu MT} \right]^{MT} \\ &\sim \exp \left\{ -\mu^{-1} \sum_{j=1}^k (b_j - a_j)(c_j^{-\alpha} - d_j^{-\alpha}) \right\} \\ &= P(N(A) = 0). \end{aligned}$$

The final step consists of dropping the assumption that the $[a_j, b_j]$ are mutually disjoint. We only consider the case $k = 2$. The general case is completely analogous. Let

$$A = \bigcup_{j=1}^2 [a_j, b_j) \times (c_j, d_j),$$

where $a_1 < a_2 \leq b_1 < b_2$ and, for example, $c_1 < d_1 \leq c_2 < d_2$. Then $P(N_{\mu_T}^{\text{lo}}(A) = 0)$ equals

$$\begin{aligned} & \left[P(X_i/x_T \notin (c_1, d_1), i = [(1 - \epsilon_T)\mu_{Ta_1}], \dots, [(1 - \epsilon_T)\mu_{Ta_2}] - 1, \right. \\ & \quad X_i/x_T \notin (c_1, d_1) \cup (c_2, d_2), i = [(1 - \epsilon_T)\mu_{Ta_2}], \dots, [(1 + \epsilon_T)\mu_{Tb_1}] - 1, \\ & \quad \left. X_i/x_T \notin (c_2, d_2), i = [(1 + \epsilon_T)\mu_{Tb_1}], \dots, [(1 + \epsilon_T)\mu_{Tb_2}] - 1 \right]^M \\ &\sim \left[(1 - \bar{F}_{\text{on}}(c_1 x_T) + \bar{F}_{\text{on}}(d_1 x_T))^{T(a_2 - a_1)\mu^{-1}} \right. \\ & \quad \left. (1 - \bar{F}_{\text{on}}(c_2 x_T) + \bar{F}_{\text{on}}(d_2 x_T))^{T(a_2 - a_1)\mu^{-1}} \right]^M \end{aligned}$$

$$\begin{aligned}
& (1 - \bar{F}_{\text{on}}(c_1 x_T) + \bar{F}_{\text{on}}(d_1 x_T))(1 - \bar{F}_{\text{on}}(c_2 x_T) + \bar{F}_{\text{on}}(d_2 x_T))^{T(b_1 - a_2)\mu^{-1}} \\
& (1 - \bar{F}_{\text{on}}(c_2 x_T) + \bar{F}_{\text{on}}(d_2 x_T))^{T(b_2 - b_1)\mu^{-1}} \Big]^M \\
= & \prod_{j=1}^2 \left[1 - \bar{F}_{\text{on}}(c_j x_T) + \bar{F}_{\text{on}}(d_j x_T) \right]^{MT(b_j - a_j)\mu^{-1}},
\end{aligned}$$

which converges to $P(N(A) = 0)$ as before.

It remains to show that for $I = [a, b] \times (c, d] \subset E$, we have

$$\limsup_{T \rightarrow \infty} P(N_T^{\text{lo}}(I) > 1) \leq P(N(I) > 1).$$

Since $I \in \mathcal{A}$, by the first part of the proof,

$$\begin{aligned}
\limsup_{T \rightarrow \infty} P(N_T^{\text{lo}}(I) > 1) & \leq 1 - \liminf_{T \rightarrow \infty} P(N_T^{\text{lo}}(I) = 0) - \liminf_{T \rightarrow \infty} P(N_T^{\text{lo}}(I) = 1) \\
& = 1 - P(N(I) = 0) - \liminf_{T \rightarrow \infty} P(N_T^{\text{lo}}(I) = 1).
\end{aligned}$$

Thus it suffices to show that

$$\liminf_{T \rightarrow \infty} P(N_T^{\text{lo}}(I) = 1) \geq P(N(I) = 1).$$

As before, let $\epsilon_T \rightarrow 0$ satisfy (A.1) and define

$$B_T = \{|\xi_{Tt}^{(m)} - \mu_{Tt}| \leq \epsilon_T \mu_{Tt}, t = a, b, m = 1, \dots, M\}.$$

Write (slightly abusing notation)

$$\begin{aligned}
N_{\mu_T}^{\text{lo}}(I) & = \sum_{m=1}^M \sum_{i=\lfloor (1+\epsilon_T) \mu_{Ta} \rfloor}^{\lfloor (1-\epsilon_T) \mu_{Tb} \rfloor - 1} 1_{(c x_T, d x_T]}(X_i^{(m)}) , \\
N_{\epsilon_T}^a & = \sum_{m=1}^M \sum_{i=\lfloor (1-\epsilon_T) \mu_{Ta} \rfloor}^{\lfloor (1+\epsilon_T) \mu_{Ta} \rfloor - 1} 1_{(c x_T, d x_T]}(X_i^{(m)}) , \\
N_{\epsilon_T}^b & = \sum_{m=1}^M \sum_{i=\lfloor (1-\epsilon_T) \mu_{Tb} \rfloor}^{\lfloor (1+\epsilon_T) \mu_{Tb} \rfloor - 1} 1_{(c x_T, d x_T]}(X_i^{(m)}).
\end{aligned}$$

Since $N_{\mu_T}^{\text{lo}}(I)$, $N_{\epsilon_T}^a$ and $N_{\epsilon_T}^b$ are mutually independent, we have

$$\begin{aligned}
P(N_T^{\text{lo}}(I) = 1) & \geq P(N_T^{\text{lo}}(I) = 1, B_T) \\
& \geq P(N_{\mu_T}^{\text{lo}}(I) = 1, N_{\epsilon_T}^a = 0, N_{\epsilon_T}^b = 0, B_T) \\
& \geq P(N_{\mu_T}^{\text{lo}}(I) = 1, N_{\epsilon_T}^a = 0, N_{\epsilon_T}^b = 0) - P(B_T^c) \\
& = P(N_{\mu_T}^{\text{lo}}(I) = 1) P(N_{\epsilon_T}^a = 0) P(N_{\epsilon_T}^b = 0) - P(B_T^c).
\end{aligned}$$

By Lemma 5.A.3 (as above), $P(B_T^c) \rightarrow 0$. Using (6.10), we observe that as $T \rightarrow \infty$

$$\begin{aligned} P(N_{\epsilon_T}^a = 0) &= \left[1 - \overline{F}_{\text{on}}(cx_T) + \overline{F}_{\text{on}}(dx_T) \right]^{M(2\epsilon_T \mu_{T^a} - 1)} \\ &\sim \left[1 - \frac{c^{-\alpha}}{MT} + \frac{d^{-\alpha}}{MT} \right]^{M(2\epsilon_T T a \mu^{-1} - 1)} \\ &\sim \left[1 - \frac{2a(c^{-\alpha} - d^{-\alpha})}{\mu MT} \right]^{\epsilon_T MT} \rightarrow 1. \end{aligned}$$

Analogously, $P(N_{\epsilon_T}^b = 0) \rightarrow 1$. Set

$$r_T = M([(1 - \epsilon_T)\mu_{T^b}] - 2 - [(1 + \epsilon_T)\mu_{T^a}]).$$

Finally, by (6.10) and the first part of the proof, $P(N_{\mu_T}^{\text{lo}}(I) = 1)$ equals

$$\begin{aligned} r_T \left[\overline{F}_{\text{on}}(cx_T) - \overline{F}_{\text{on}}(dx_T) \right] \left[1 - \overline{F}_{\text{on}}(cx_T) + \overline{F}_{\text{on}}(dx_T) \right]^{r_T - 1} \\ \sim \frac{MT(b - a)}{\mu} \frac{(c^{-\alpha} - d^{-\alpha})}{MT} P(N(I) = 0) \\ = \frac{(b - a)(c^{-\alpha} - d^{-\alpha})}{\mu} \exp\{\mu^{-1}(b - a)(c^{-\alpha} - d^{-\alpha})\} \\ = P(N(I) = 1). \end{aligned}$$

This completes the proof for N_T^{lo} . Since the proof is based upon the fact that the counting process ξ_T can be replaced by its mean μ_T , the same techniques can be used to show convergence of N_T^{up} . \heartsuit

Remark To prove convergence of N_T^{up} we can also apply Kallenberg [52], Theorem 4.7 (see also Resnick [78], Proposition 3.22). This involves showing that $EN_T^{\text{up}}(A) \rightarrow EN(A)$, which follows from Wald's identity and (6.10). Another alternative way to prove Theorem 6.1 is by starting with convergence of either N_T^{lo} or N_T^{up} and then showing that the distance between N_T^{lo} and N_T^{up} in the vague topology on $M_p(E)$ (see Resnick [77], Section 3.4) becomes arbitrarily small as $T \rightarrow \infty$. By virtue of Kallenberg [52], Theorem 4.2, this means showing that for each continuous function $f : E \rightarrow \mathbb{R}$ with compact support, one has

$$P(|N_T^{\text{lo}}(f) - N_T^{\text{up}}(f)| > \epsilon) \rightarrow 0,$$

for every $\epsilon > 0$.

6.B Proof of Theorem 6.3

We need the following lemma.

Lemma 6.B.1 *Assume \bar{F}_{on} satisfies (6.12). Suppose M satisfies the Fast Growth Condition and a_T satisfies (6.13). Let $0 \leq a < b \leq 1$ and define*

$$B_i = \{(T - X_i)/a_T \in (a, b]\}, \quad i \geq 1.$$

As $T \rightarrow \infty$

$$M E \left(\sum_{i=2}^{\xi_T} \sum_{l=1}^{i-1} 1_{B_i \cap B_l} \right) \rightarrow 0.$$

Proof. Using a first order Taylor expansion, we have for T large enough

$$\begin{aligned} P(B_1) &= \bar{F}_{\text{on}}(T - ba_T) - \bar{F}_{\text{on}}(T - aa_T) \\ &= \bar{F}_{\text{on}}(T) \left(\frac{\bar{F}_{\text{on}}(T - ba_T)}{\bar{F}_{\text{on}}(T)} - \frac{\bar{F}_{\text{on}}(T - aa_T)}{\bar{F}_{\text{on}}(T)} \right) \\ &= \bar{F}_{\text{on}}(T) \left(\left(1 - b \frac{a_T}{T}\right)^{-\alpha} - \left(1 - a \frac{a_T}{T}\right)^{-\alpha} \right) \\ &= \bar{F}_{\text{on}}(T) \left(\alpha(b-a) \frac{a_T}{T} + o\left(\frac{a_T}{T}\right) \right). \end{aligned}$$

Hence, by (6.12) and (6.13)

$$M T P(B_1) \sim \alpha(b-a) M a_T \bar{F}_{\text{on}}(T) \sim \alpha(b-a). \quad (\text{B.1})$$

We have

$$\begin{aligned} 2 \sum_{i=2}^{\xi_T} \sum_{l=1}^{i-1} 1_{B_i \cap B_l} &= \left(\sum_{i=1}^{\xi_T} 1_{B_i} \right)^2 - \sum_{i=1}^{\xi_T} 1_{B_i} \\ &= \left(\sum_{i=1}^{\xi_T} (1_{B_i} - P(B_1)) \right)^2 + 2 P(B_1) \xi_T \sum_{i=1}^{\xi_T} 1_{B_i} \\ &\quad - [P(B_1)]^2 \xi_T^2 - \sum_{i=1}^{\xi_T} 1_{B_i} \\ &= I_T + 2 II_T - III_T - IV_T. \end{aligned}$$

From Gut [43], Theorem II5.1, we have that for $r > 0$

$$E \xi_T^r \sim \frac{T^r}{\mu^r} \quad \text{as } T \rightarrow \infty. \quad (\text{B.2})$$

Using this fact and (B.1), we obtain

$$M EIII_T \sim (\text{const}) M^{-1} \rightarrow 0.$$

Gut [43], Theorem I5.3, gives

$$E \left(\sum_{i=1}^{\xi_T} (1_{B_i} - P(B_1)) \right)^2 = \text{Var}(1_{B_1}) E \xi_T. \quad (\text{B.3})$$

Using this, we find that

$$\begin{aligned} M (EI_T - EIV_T) &= M \text{Var}(1_{B_1}) E \xi_T - M P(B_1) E \xi_T \\ &= M (P(B_1) - [P(B_1)]^2) E \xi_T - M P(B_1) E \xi_T \\ &= -M [P(B_1)]^2 E \xi_T, \end{aligned}$$

which is $o(1)$ by (B.1) and (B.2). For II_T we have

$$M EII_T = M P(B_1) E \left(\xi_T \sum_{i=1}^{\xi_T} (1_{B_i} - P(B_1)) \right) + M [P(B_1)]^2 E \xi_T^2.$$

The second term on the right-hand side is $o(1)$ by (B.1) and (B.2). Using the Cauchy-Schwarz inequality and (B.3), we obtain for the first term the upper bound

$$M P(B_1) (E \xi_T^2)^{1/2} (P(B_1) [1 - P(B_1)] E \xi_T)^{1/2},$$

which is $o(1)$ by (B.1) and (B.2). This completes the proof. \heartsuit

Proof of Theorem 6.3 We consider sets of the form

$$A = \bigcup_{j=1}^k [r_j, s_j) \times (a_j, b_j),$$

for $k \geq 1$, $[r_j, s_j) \times (a_j, b_j) \subset [0, \infty) \times (0, 1]$, $j = 1, \dots, k$. We may assume that the sets $[r_j, s_j) \times (a_j, b_j)$ are mutually disjoint. Notice that \widehat{N}_T are the row sums of a triangular array of point processes. According to Theorem 6.D.2, it suffices to show

$$M P(\widehat{N}_T^{(1)}(A) \geq 1) \rightarrow E \widehat{N}(A) \quad \text{and} \quad M P(\widehat{N}_T^{(1)}(A) \geq 2) \rightarrow 0.$$

Since

$$E \widehat{N}_T(A) = \sum_{n=1}^{\infty} M P(\widehat{N}_T^{(1)}(A) \geq n),$$

it also suffices to show

$$E \widehat{N}_T(A) \rightarrow E \widehat{N}(A) \quad \text{and} \quad M P(\widehat{N}_T^{(1)}(A) \geq 1) \rightarrow E \widehat{N}(A).$$

Define

$$B_{ij} = \{(T - X_i)/a_T \in (a_j, b_j]\} \quad i \geq 1, \quad j = 1, \dots, k.$$

Notice that

$$\widehat{N}_T(A) = \sum_{m=1}^M \sum_{i=\xi_{T r_j}^{(m)} + 1}^{\xi_{T s_j}^{(m)}} 1_{[T - b_j a_T, T - a_j a_T]}(X_i^{(m)}) .$$

Using the fact that ξ_T is a stopping time, Wald's identity and (B.1), we obtain

$$\begin{aligned} E\widehat{N}_T(A) &= \mu^{-1} \sum_{j=1}^k (s_j - r_j) M T P(B_{1j}) \\ &\sim \alpha \mu^{-1} \sum_{j=1}^k (s_j - r_j)(b_j - a_j) \\ &= E\widehat{N}(A). \end{aligned}$$

It remains to show that $MP(\widehat{N}_T^{(1)}(A) \geq 1) \rightarrow E\widehat{N}(A)$. By Markov's inequality we have

$$\limsup_{T \rightarrow \infty} M P(\widehat{N}_T^{(1)}(A) \geq 1) \leq E\widehat{N}(A).$$

For nonnegative vectors $\mathbf{n}_r = (n_r^{(1)}, \dots, n_r^{(k)})$ and $\mathbf{n}_s = (n_s^{(1)}, \dots, n_s^{(k)})$, let

$$A_{\mathbf{n}_r}^{\mathbf{n}_s} = \{\xi_{T r_j} = n_r^{(j)}, \xi_{T s_j} = n_s^{(j)}, j = 1, \dots, k\}.$$

Notice that for $\mathbf{n}_r \leq \mathbf{n}_s$, which means that $n_r^{(j)} \leq n_s^{(j)}$ for $j = 1, \dots, k$, the $A_{\mathbf{n}_r}^{\mathbf{n}_s}$ constitute a partition of Ω . We write

$$\begin{aligned} P(\widehat{N}_T^{(1)}(A) \geq 1) &= P\left(\sum_{j=1}^k \sum_{i=\xi_{T r_j} + 1}^{\xi_{T s_j}} 1_{B_{ij}} \geq 1\right) \\ &= P\left(\bigcup_{j=1}^k \bigcup_{i=\xi_{T r_j} + 1}^{\xi_{T s_j}} B_{ij}\right) \\ &= \sum_{\mathbf{n}_r \leq \mathbf{n}_s} P\left(\bigcup_{j=1}^k \bigcup_{i=n_r^{(j)} + 1}^{n_s^{(j)}} (B_{ij} \cap A_{\mathbf{n}_r}^{\mathbf{n}_s})\right). \end{aligned}$$

For fixed \mathbf{n}_r and \mathbf{n}_s , we apply the inclusion-exclusion formula on the two unions as a whole, which yields the lower bound

$$\sum_{j=1}^k \sum_{i=n_r^{(j)} + 1}^{n_s^{(j)}} P(B_{ij} \cap A_{\mathbf{n}_r}^{\mathbf{n}_s}) - \sum_{j=2}^k \sum_{n=1}^{j-1} \sum_{i=n_r^{(j)} + 1}^{n_s^{(j)}} \sum_{l=n_r^{(n)} + 1}^{n_s^{(n)}} P(B_{ij} \cap B_{ln} \cap A_{\mathbf{n}_r}^{\mathbf{n}_s})$$

$$- \sum_{j=1}^k \sum_{i=n_r^{(j)}+2}^{n_s^{(j)}} \sum_{l=n_r^{(j)}+1}^{i-1} P(B_{ij} \cap B_{lj} \cap A_{\mathbf{n_r}}^{\mathbf{n_s}}).$$

Using $P(B) = E(1_B)$ and the linearity of the expectation, this equals

$$\begin{aligned} \sum_{j=1}^k E \left(\sum_{i=n_r^{(j)}+1}^{n_s^{(j)}} 1_{B_{ij} \cap A_{\mathbf{n_r}}^{\mathbf{n_s}}} \right) & - \sum_{j=2}^k \sum_{n=1}^{j-1} E \left(\sum_{i=n_r^{(j)}+1}^{n_s^{(j)}} \sum_{l=n_r^{(n)}+1}^{n_s^{(n)}} 1_{B_{ij} \cap B_{ln} \cap A_{\mathbf{n_r}}^{\mathbf{n_s}}} \right) \\ & - \sum_{j=1}^k E \left(\sum_{i=n_r^{(j)}+2}^{n_s^{(j)}} \sum_{l=n_r^{(j)}+1}^{i-1} 1_{B_{ij} \cap B_{lj} \cap A_{\mathbf{n_r}}^{\mathbf{n_s}}} \right). \end{aligned}$$

Since the $A_{\mathbf{n_r}}^{\mathbf{n_s}}$ are a partition of Ω , we have shown that

$$\begin{aligned} M P(\widehat{N}_T^{(1)}(A) \geq 1) & \geq M \sum_{j=1}^k E \left(\sum_{i=\xi_{Tr_j}+1}^{\xi_{Ts_j}} 1_{B_{ij}} \right) - M \sum_{j=2}^k \sum_{n=1}^{j-1} E \left(\sum_{i=\xi_{Tr_j}+1}^{\xi_{Ts_j}} \sum_{l=\xi_{Tr_n}+1}^{\xi_{Ts_n}} 1_{B_{ij} \cap B_{ln}} \right) \\ & - M \sum_{j=1}^k E \left(\sum_{i=\xi_{Tr_j}+2}^{\xi_{Ts_j}} \sum_{l=\xi_{Tr_j}+1}^{i-1} 1_{B_{ij} \cap B_{lj}} \right) \\ & = I_T - II_T - III_T. \end{aligned}$$

Notice that $I_T = E\widehat{N}_T(A)$, which converges to $E\widehat{N}(A)$. It remains to show that II_T and III_T converge to zero. Since, for fixed j

$$\sum_{i=\xi_{Tr_j}+2}^{\xi_{Ts_j}} \sum_{l=\xi_{Tr_j}+1}^{i-1} 1_{B_{ij} \cap B_{lj}} \leq \sum_{i=2}^k \sum_{l=1}^{i-1} 1_{B_{ij} \cap B_{lj}},$$

$III_T \rightarrow 0$ follows from Lemma 6.B.1. Next we deal with II_T . Fix j and n . We have $n < j$. Assume that the sets $[r_j, s_j)$ are ordered such that

$$n < j \Rightarrow r_n \leq r_j.$$

Denote

$$D_T^{(j,n)} = \sum_{i=\xi_{Tr_j}+1}^{\xi_{Ts_j}} \sum_{l=\xi_{Tr_n}+1}^{\xi_{Ts_n}} 1_{B_{ij} \cap B_{ln}},$$

and define the following disjoint partition of Ω :

$$\begin{aligned} C_1(j, n) &= \{\xi_{Ts_n} \leq \xi_{Tr_j}\}, \\ C_2(j, n) &= \{\xi_{Ts_j} \geq \xi_{Ts_n} > \xi_{Tr_j}\}, \\ C_3(j, n) &= \{\xi_{Ts_n} > \xi_{Ts_j}\}. \end{aligned}$$

We will show that

$$D_T^{(j,n)} \leq 3 \sum_{i=2}^{\xi_{T s_j} - 1} \sum_{l=1}^{\xi_{T s_n} - 1} 1_{\tilde{B}_i \cap \tilde{B}_l} + 2 \sum_{l=2}^{\xi_{T s_n} - 1} \sum_{i=1}^{\xi_{T s_j} - 1} 1_{\tilde{B}_i \cap \tilde{B}_l}, \quad (\text{B.4})$$

where

$$\tilde{B}_i = \{(T - X_i)/a_T \in (\min(a_j, a_n), \max(b_j, b_n)]\}, \quad i \geq 1.$$

Then $II_T \rightarrow 0$ follows from Lemma 6.B.1. It is not difficult to see that

$$D_T^{(j,n)} 1_{C_1(j,n)} \leq \sum_{i=\xi_{T r_j} + 1}^{\xi_{T s_j}} \sum_{l=\xi_{T r_n} + 1}^{\xi_{T s_n}} 1_{\tilde{B}_i \cap \tilde{B}_l} \leq \sum_{i=2}^{\xi_{T s_j} - 1} \sum_{l=1}^{\xi_{T s_j} - 1} 1_{\tilde{B}_i \cap \tilde{B}_l}. \quad (\text{B.5})$$

Next consider $C_2(j,n)$. From the definition of ξ_T in (5.3) it follows that if $C_2(j,n) \neq \emptyset$, then $s_n > r_j$. The ordering of the sets $[r_j, s_j)$ now implies that $[r_j, s_j)$ and $[r_n, s_n)$ are not disjoint.

Since $[r_j, s_j) \times (a_j, b_j]$ are mutually disjoint, it follows that $(a_j, b_j) \cap (a_n, b_n) = \emptyset$. Hence, $B_{ij} \cap B_{ln} = \emptyset$ if $i = l$. The reasoning above implies that

$$\begin{aligned} D_T^{(j,n)} 1_{C_2(j,n)} &\leq \sum_{i=\xi_{T r_n} + 1}^{\xi_{T s_j}} \sum_{\substack{l=\xi_{T r_n} + 1 \\ l \neq i}}^{\xi_{T s_j}} 1_{B_{ij} \cap B_{ln}} \\ &\leq \sum_{i=\xi_{T r_n} + 1}^{\xi_{T s_j}} \sum_{\substack{l=\xi_{T r_n} + 1 \\ l \neq i}}^{\xi_{T s_j}} 1_{\tilde{B}_i \cap \tilde{B}_l} \\ &\leq 2 \sum_{i=2}^{\xi_{T s_j} - 1} \sum_{l=1}^{\xi_{T s_j} - 1} 1_{\tilde{B}_i \cap \tilde{B}_l}. \end{aligned} \quad (\text{B.6})$$

Since $C_3(j,n) \neq \emptyset$ also implies that $[r_j, s_j)$ and $[r_n, s_n)$ are not disjoint, we have analogously

$$D_T^{(j,n)} 1_{C_3(j,n)} \leq \sum_{i=\xi_{T r_n} + 1}^{\xi_{T s_n}} \sum_{\substack{l=\xi_{T r_n} + 1 \\ l \neq i}}^{\xi_{T s_n}} 1_{B_{ij} \cap B_{ln}} \leq 2 \sum_{i=2}^{\xi_{T s_n} - 1} \sum_{l=1}^{\xi_{T s_n} - 1} 1_{\tilde{B}_i \cap \tilde{B}_l}. \quad (\text{B.7})$$

The proof now follows by observing that (B.5)–(B.7) together imply (B.4). \heartsuit

Remark From the steps in the proof of Theorem 6.3 it is clear that similar results hold if one imposes a second order regular variation condition on \bar{F}_{on} , instead of assuming (6.12). Proposition 6.4 will also follow in this case. Such a proof would, however, be more technical.

6.C Proof of Theorem 6.5

For ease of presentation, we introduce some notation. Set

$$\begin{aligned}\tilde{X}_i &= 1_{[x_T, T]}(X_i), \quad \tilde{\mu} = E\tilde{X}, \quad \tilde{\sigma}^2 = \text{Var}(\tilde{X}), \\ \tilde{\mu}_3 &= E(\tilde{X} - \tilde{\mu})^3, \quad \tilde{\mu}_4 = E(\tilde{X} - \tilde{\mu})^4.\end{aligned}$$

For $a \leq b$, write

$$\overline{F}_{\text{on}}(a, b) = \overline{F}_{\text{on}}(a) - \overline{F}_{\text{on}}(b).$$

Notice that

$$\tilde{\sigma}^2 \sim \tilde{\mu}_3 \sim \tilde{\mu}_4 \sim \overline{F}_{\text{on}}(x_T, T). \quad (\text{C.1})$$

For $n \geq 1$, let

$$\tilde{S}_0 = 0, \quad \tilde{S}_n = \tilde{X}_1 + \cdots + \tilde{X}_n \quad \text{and} \quad \mathcal{F}_n = \sigma(D, X_1, Y_1, \dots, X_n, Y_n).$$

Recall that ξ_T is a stopping time with respect to the filtration (\mathcal{F}_n) .

We will derive Lyapunov's condition (with $\delta = 2$) for the central limit theorem, i.e. as $T \rightarrow \infty$

$$\frac{M E(Z_T^{(1)})^4}{[M E(Z_T^{(1)})^2]^2} \sim \mu^2 M E(Z_T^{(1)})^4 \rightarrow 0.$$

We have

$$M E(Z_T^{(1)})^4 = \frac{E(\tilde{S}_{\xi_T} - \xi_T \tilde{\mu})^4}{M T^2 [\overline{F}_{\text{on}}(x_T, T)]^2}. \quad (\text{C.2})$$

The approach is to calculate $E(\tilde{S}_{\xi_T} - \xi_T \tilde{\mu})^4$ by first constructing a martingale and then using the optional stopping theorem. It can be seen that

$$\begin{aligned}E[(\tilde{S}_n - n\tilde{\mu})^4 | \mathcal{F}_{n-1}] &= (\tilde{S}_{n-1} - (n-1)\tilde{\mu})^4 + 6\tilde{\sigma}^2 (\tilde{S}_{n-1} - (n-1)\tilde{\mu})^2 \\ &\quad + 4\tilde{\mu}_3 (\tilde{S}_{n-1} - (n-1)\tilde{\mu}) + \tilde{\mu}_4.\end{aligned}$$

Define

$$A_N = \sum_{n=1}^N \left[(\tilde{S}_n - n\tilde{\mu})^4 - E[(\tilde{S}_n - n\tilde{\mu})^4 | \mathcal{F}_{n-1}] \right].$$

By definition A_N is a martingale with respect to the filtration (\mathcal{F}_N) and from the optional stopping theorem (see Gut [43], Theorem A2.4) it follows that $EA_{\xi_T} = 0$. This implies

$$\begin{aligned}E(\tilde{S}_{\xi_T} - \xi_T \tilde{\mu})^4 &= 6\tilde{\sigma}^2 E \left[\sum_{n=1}^{\xi_T} (\tilde{S}_{n-1} - (n-1)\tilde{\mu})^2 \right] \\ &\quad + 4\tilde{\mu}_3 E \left[\sum_{n=1}^{\xi_T} (\tilde{S}_{n-1} - (n-1)\tilde{\mu}) \right] + \tilde{\mu}_4 E \xi_T.\end{aligned} \quad (\text{C.3})$$

From (E.1) in Lemma 6.E.1 it follows that

$$\begin{aligned}
E \left[\sum_{n=1}^{\xi_T} (\tilde{S}_{n-1} - (n-1)\tilde{\mu}) \right] &= E \left[\sum_{i=1}^{\xi_T-1} (\tilde{X}_i - \tilde{\mu})(\xi_T - i) \right] \\
&= E \left[\sum_{i=1}^{\xi_T-1} (\tilde{X}_i - \tilde{\mu})i \right] \\
&= E \left[\sum_{i=1}^{\xi_T} (\tilde{X}_i - \tilde{\mu})i \right] - E[\xi_T(\tilde{X}_{\xi_T} - \tilde{\mu})] \\
&= -E[\xi_T(\tilde{X}_{\xi_T} - \tilde{\mu})]. \tag{C.4}
\end{aligned}$$

The final step follows from the fact that $\sum_{i=1}^n (\tilde{X}_i - \tilde{\mu})i$ is a martingale with respect to (\mathcal{F}_n) and the optional stopping theorem.

Continuing the analysis, we write

$$\begin{aligned}
E \left[\sum_{n=1}^{\xi_T} (\tilde{S}_{n-1} - (n-1)\tilde{\mu})^2 \right] &= E \left[\sum_{n=2}^{\xi_T} \sum_{i=1}^{n-1} (\tilde{X}_i - \tilde{\mu})^2 \right] \tag{C.5} \\
&\quad + 2 E \left[\sum_{n=3}^{\xi_T} \sum_{i=2}^{n-1} \sum_{j=1}^{i-1} (\tilde{X}_i - \tilde{\mu})(\tilde{X}_j - \tilde{\mu}) \right].
\end{aligned}$$

By (E.1) in Lemma 6.E.1, the first term equals

$$\begin{aligned}
E \left[\sum_{i=1}^{\xi_T-1} (\tilde{X}_i - \tilde{\mu})^2 (\xi_T - i) \right] &= E \left[\sum_{i=1}^{\xi_T} (\tilde{X}_i - \tilde{\mu})^2 i \right] - E[\xi_T(\tilde{X}_{\xi_T} - \tilde{\mu})^2] \\
&= \frac{\tilde{\sigma}^2}{2} [E\xi_T^2 + E\xi_T] \\
&\quad - E[\xi_T(\tilde{X}_{\xi_T} - \tilde{\mu})^2]. \tag{C.6}
\end{aligned}$$

In the last step we used the fact that $\sum_{i=1}^n (\tilde{X}_i - \tilde{\mu})^2 i - \tilde{\sigma}^2 n(n+1)/2$ is a martingale with respect to (\mathcal{F}_n) and the optional stopping theorem.

By (E.2) in Lemma 6.E.1, the second term in (C.5) equals

$$\begin{aligned}
&2 E \left[\sum_{i=2}^{\xi_T-1} (\tilde{X}_i - \tilde{\mu})(\xi_T - i) \sum_{j=1}^{i-1} (\tilde{X}_j - \tilde{\mu}) \right] \\
&= 2 E \left[\sum_{i=2}^{\xi_T} (\tilde{X}_i - \tilde{\mu}) \sum_{j=1}^{i-1} (\tilde{X}_j - \tilde{\mu}) j \right] - 2 E \left[(\tilde{X}_{\xi_T} - \tilde{\mu}) \sum_{j=1}^{\xi_T-1} (\tilde{X}_j - \tilde{\mu}) j \right]
\end{aligned}$$

$$= -2 E \left[(\tilde{X}_{\xi_T} - \tilde{\mu}) \sum_{j=1}^{\xi_T} (\tilde{X}_j - \tilde{\mu}) j \right] + 2 E[\xi_T (\tilde{X}_{\xi_T} - \tilde{\mu})^2]. \quad (\text{C.7})$$

In the final step we used the fact that $\sum_{i=1}^n (\tilde{X}_i - \tilde{\mu}) \sum_{j=1}^{i-1} (\tilde{X}_j - \tilde{\mu}) j$ is a martingale with respect to (\mathcal{F}_n) and the optional stopping theorem.

Combining (C.3)–(C.7) yields

$$\begin{aligned} E(\tilde{S}_{\xi_T} - \xi_T \tilde{\mu})^4 &= 3 \tilde{\sigma}^4 E\xi_T^2 + (3 \tilde{\sigma}^4 + \tilde{\mu}_4) E\xi_T \\ &\quad - 4 \tilde{\mu}_3 E[\xi_T (\tilde{X}_{\xi_T} - \tilde{\mu})] \\ &\quad + 6 \tilde{\sigma}^2 E[\xi_T (\tilde{X}_{\xi_T} - \tilde{\mu})^2] \\ &\quad - 12 \tilde{\sigma}^2 E \left[(\tilde{X}_{\xi_T} - \tilde{\mu}) \sum_{j=1}^{\xi_T} (\tilde{X}_j - \tilde{\mu}) j \right]. \end{aligned}$$

We will now show that

$$E(\tilde{S}_{\xi_T} - \xi_T \tilde{\mu})^4 = o(M T^2 [\bar{F}_{\text{on}}(x_T, T)]^2).$$

The proof of the Lyapunov condition then follows from (C.2).

By virtue of (B.2), (6.14) and (C.1) we obtain

$$3 \tilde{\sigma}^4 E\xi_T^2 + (3 \tilde{\sigma}^4 + \tilde{\mu}_4) E\xi_T = o(M T^2 [\bar{F}_{\text{on}}(x_T, T)]^2).$$

Furthermore,

$$\begin{aligned} |\tilde{\mu}_3 E[\xi_T (\tilde{X}_{\xi_T} - \tilde{\mu})]| &\leq |\tilde{\mu}_3| E\xi_T \sim \mu^{-1} T \bar{F}_{\text{on}}(x_T, T) \\ &= o(M T^2 [\bar{F}_{\text{on}}(x_T, T)]^2), \end{aligned}$$

and

$$\begin{aligned} \tilde{\sigma}^2 E[\xi_T (\tilde{X}_{\xi_T} - \tilde{\mu})^2] &\leq \tilde{\sigma}^2 E\xi_T \sim \mu^{-1} T \bar{F}_{\text{on}}(x_T, T) \\ &= o(M T^2 [\bar{F}_{\text{on}}(x_T, T)]^2). \end{aligned}$$

Finally,

$$\begin{aligned} \tilde{\sigma}^2 \left| E \left[(\tilde{X}_{\xi_T} - \tilde{\mu}) \sum_{j=1}^{\xi_T} (\tilde{X}_j - \tilde{\mu}) j \right] \right| &\leq \tilde{\sigma}^2 E \left[\sum_{j=1}^{\xi_T} |\tilde{X}_j - \tilde{\mu}| j \right] \\ &= \tilde{\sigma}^2 E \left[\sum_{j=1}^{\xi_T} (|\tilde{X}_j - \tilde{\mu}| j - E|\tilde{X}_j - \tilde{\mu}| j) \right] \end{aligned}$$

$$\begin{aligned}
& + \tilde{\sigma}^2 E \left[\sum_{j=1}^{\xi_T} j \right] E |\tilde{X} - \tilde{\mu}| \\
& \leq 2 \tilde{\sigma}^2 E \xi_T^2 \bar{F}_{\text{on}}(x_T, T) \sim 2 \mu^{-2} T^2 [\bar{F}_{\text{on}}(x_T, T)]^2 \\
& = o(M T^2 [\bar{F}_{\text{on}}(x_T, T)]^2).
\end{aligned}$$

In the third step we again used the fact that ξ_T is a stopping time. This completes the proof. \heartsuit

6.D Convergence to a simple point process

Let $M_p(E)$ denote the space of all point measures defined on the state space $E = [0, \infty) \times (0, \infty]$. The space $M_p(E)$ is equipped with the vague topology (see Resnick [77], Section 3.4). A point process N on E is called *simple* if

$$P(N(\{x\}) \leq 1 \text{ for all } x \in E) = 1,$$

i.e. if any point x counts at most once, with probability 1. Notice that a Poisson Random Measure (PRM) N with mean measure ν is simple if ν is atomless.

Let \mathcal{I} be the class of rectangles $[a, b) \times (c, d] \subset E$ and let \mathcal{A} be the collection of sets

$$A = \bigcup_{j=1}^k I_j, \quad I_j \in \mathcal{I}, \quad j = 1, \dots, k, \quad k \geq 1.$$

Observe that we may assume that the sets $[a_j, b_j) \times (c_j, d_j]$ are mutually disjoint. If two such sets intersect, they are the union of at most three mutually disjoint sets of the form $[a, b) \times (c, d]$. Notice that \mathcal{A} is closed under finite unions and intersections.

It can be seen that for any compact $K \subset E$ and open $G \subset E$ with $K \subset G$, there exists an $A \in \mathcal{A}$ such that $K \subset A \subset G$. In Kallenberg [53] a class \mathcal{A} with this property is called a *separating class*. Evidently, finite unions of elements in \mathcal{I} constitute a separating class. Any such class \mathcal{I} is called a *pre-separating class*.

The following result is due to Kallenberg [53] and is an improved version of Kallenberg [52], Theorem 4.7.

Theorem 6.D.1 *Suppose N and N_T , $T > 0$, are point processes on E and N is simple. If for all $A \in \mathcal{A}$ and $I \in \mathcal{I}$*

$$\lim_{T \rightarrow \infty} P(N_T(A) = 0) = P(N(A) = 0),$$

$$\limsup_{T \rightarrow \infty} P(N_T(I) > 1) \leq P(N(I) > 1),$$

then $N_T \xrightarrow{d} N$ in $M_p(E)$. \heartsuit

Next we state a result on the convergence of a triangular array of point processes to a PRM. Let M_T be a non-decreasing integer-valued function such that $M_T \rightarrow \infty$ as $T \rightarrow \infty$. Let

$$(N_T^{(m)}, m = 1, \dots, M_T, T > 0)$$

be a triangular array of point processes on E , such that for each T the processes $(N_T^{(m)}, m = 1, \dots, M_T)$ are mutually independent. The array is *uniformly asymptotically negligible* (u.a.n.) if

$$\lim_{T \rightarrow \infty} \sup_{m=1, \dots, M_T} P(N_T^{(m)}(A) > 0) = 0,$$

for all sets $A \in \mathcal{A}$. Define the row sums

$$N_T = \sum_{m=1}^{M_T} N_T^{(m)}, \quad T > 0.$$

The following result can be found in Daley and Vere-Jones [22], Theorem 9.2.V.

Theorem 6.D.2 *Let N be a simple PRM with mean measure ν . If the triangular array $(N_T^{(m)})$ is u.a.n., then $N_T \xrightarrow{d} N$ in $M_p(E)$ if and only if for all sets $A \in \mathcal{A}$*

$$\lim_{T \rightarrow \infty} \sum_{m=1}^{M_T} P(N_T^{(m)}(A) \geq 1) = \nu(A) \quad \text{and} \quad \lim_{T \rightarrow \infty} \sum_{m=1}^{M_T} P(N_T^{(m)}(A) \geq 2) = 0.$$

♡

6.E An identity in law for stopped random sums

Here we use the notation introduced at the beginning of Section 6.C.

Lemma 6.E.1 *Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a measurable function. Then*

$$\sum_{i=1}^{\xi_T-1} f(\tilde{X}_i)(\xi_T - i) \stackrel{d}{=} \sum_{i=1}^{\xi_T-1} f(\tilde{X}_i)i, \quad (\text{E.1})$$

and

$$\sum_{i=1}^{\xi_T-1} f(\tilde{X}_i)(\xi_T - i) \sum_{j=1}^{i-1} f(\tilde{X}_j) \stackrel{d}{=} \sum_{i=1}^{\xi_T-1} f(\tilde{X}_i) \sum_{j=1}^{i-1} f(\tilde{X}_j)j. \quad (\text{E.2})$$

Proof. First we prove (E.1). We have

$$p_T = P \left(\sum_{i=1}^{\xi_T-1} f(\tilde{X}_i)(\xi_T - i) > x \right)$$

$$\begin{aligned}
&= \sum_{n=0}^{\infty} P \left(\sum_{i=1}^{\xi_T-1} f(\tilde{X}_i)(\xi_T - i) > x, \xi_T = n \right) \\
&= \sum_{n=0}^{\infty} P \left(\sum_{i=1}^{n-1} f(\tilde{X}_i)(n - i) > x, S_{n-1} \leq T, S_n > T \right) \\
&= \sum_{n=0}^{\infty} P \left(\sum_{i=1}^{n-1} f(\tilde{X}_{n-i})i > x, S_{n-1} \leq T, S_n > T \right).
\end{aligned}$$

Notice that the probabilities above are invariant under a permutation on X_1, \dots, X_{n-1} . We apply the permutation $X_i \rightarrow X_{n-i}$, $i = 1, \dots, n-1$. This yields

$$\begin{aligned}
p_T &= \sum_{n=0}^{\infty} P \left(\sum_{i=1}^{n-1} f(\tilde{X}_i)i > x, D + \sum_{i=1}^{n-1} (X_{n-i} + Y_i) \leq T, \right. \\
&\quad \left. D + \sum_{i=1}^{n-1} (X_{n-i} + Y_i) + (X_n + Y_n) > T \right) \\
&= \sum_{n=0}^{\infty} P \left(\sum_{i=1}^{\xi_T-1} f(\tilde{X}_i)i > x, \xi_T = n \right) \\
&= P \left(\sum_{i=1}^{\xi_T-1} f(\tilde{X}_i)i > x \right).
\end{aligned}$$

This completes the proof of (E.1). The proof of (E.2) is along the same lines and therefore omitted. \heartsuit

Bibliography

- [1] ABRY, P. AND VEITCH, D. (1998) Wavelet analysis of long-range dependent traffic. *IEEE Transactions on Information Theory* **44**, 2–15.
- [2] ASMUSSEN, S. (1987) *Applied Probability and Queues*. Wiley, Chichester.
- [3] BARTLETT, M.S. (1955) *An Introduction to Stochastic Processes with Special Reference to Methods and Applications, First Edition*. Cambridge University Press, Cambridge (U.K.).
- [4] BERAN, J. (1994) *Statistics for Long-Memory Processes*. Chapman & Hall, New York.
- [5] BERAN, J., SHERMAN, R., TAQQU, M.S. AND WILLINGER, W. (1995) Long-range dependence in variable-bit-rate video traffic. *IEEE Transactions on Communications* **43**, 1566–1579.
- [6] BHATTACHARYA, R.N., GUPTA, V.K. AND WAYMIRE, E. (1983) The Hurst effect under trends. *Journal of Applied Probability* **20**, 649–662.
- [7] BILLINGSLEY, P. (1968) *Convergence of Probability Measures*. Wiley, New York.
- [8] BILLINGSLEY, P. (1995) *Probability and Measure, Third Edition*. Wiley, New York.
- [9] BINGHAM, N.H., GOLDIE, C.M. AND TEUGELS, J.L. (1987) *Regular Variation*. Cambridge University Press, Cambridge (U.K.).
- [10] BLAKE, W. (1994) *The Marriage of Heaven and Hell*. Dover Publications, Inc., New York.
- [11] BOES, D.C. AND SALAS, J.D. (1978) Nonstationarity of the mean and the Hurst phenomenon. *Water Resources Research* **14**, 135–143.
- [12] BOXMA, O.J. (1996) Fluid queues and regular variation. *Performance Evaluation* **27-28**, 699–712.
- [13] BOXMA, O.J. AND DUMAS, V. (1998) Fluid queues with long-tailed activity period distributions. *Computer Communications* **21**, 1509–1529.

- [14] BRAKMO, L.S., O'MALLEY, S.W. AND PETERSON, L.L. (1994) TCP Vegas: New techniques for congestion detection and avoidance. *Proceedings of the SIGCOMM '94 Conference*, University College London, London, 24–35.
- [15] BROCKWELL, P.J. AND DAVIS, R.A. (1991) *Time Series: Theory and Methods, Second Edition*. Springer-Verlag, New York.
- [16] CAO, J., CLEVELAND, W.S., LIN, D. AND SUN, D.X. (2001) On the nonstationarity of Internet traffic. *Proceedings of the ACM SIGMETRICS 2001 Conference*, MIT, Cambridge (U.S.A.), 102–112.
- [17] CAO, J., CLEVELAND, W.S., LIN, D. AND SUN, D.X. (2001) The effect of statistical multiplexing on Internet packet traffic: theory and empirical study. Technical Report, Bell Labs, Murray Hill, New Jersey.
- [18] CLINE, D.B.H. AND HSING, T. (1991) Large deviation probabilities for sums and maxima of random variables with heavy or subexponential tails. Preprint, Texas A & M University.
- [19] COX, D.R. (1984) Long-range dependence: a review. In: DAVID, H.A. AND DAVID, H.T. (editors). *Statistics: An Appraisal*. Iowa State University Press, 55–74.
- [20] CROVELLA, M. AND BESTAVROS, A. (1995) Explaining world wide web traffic self-similarity (revised). Technical Report TR-95-015, Computer Science Department, Boston University.
- [21] CROVELLA, M. AND BESTAVROS, A. (1997) Self-similarity in world wide web traffic, evidence and possible causes. *IEEE/ACM Transactions on Networking* **5**, 835–846.
- [22] DALEY, D.J. AND VERE-JONES, D. (1988) *An Introduction to the Theory of Point Processes*. Springer-Verlag, New York.
- [23] DIJKSTRA, E.W. (1959) A note on two problems in connection with graphs. *Numerical Mathematics* **1**, 269–271.
- [24] DREES, H., DE HAAN, L. AND RESNICK, S. (2000) How to make a Hill plot. *Annals of Statistics* **28**, 254–274.
- [25] DUFFIELD, N.G., LEWIS, J.T., O'CONNELL, N., RUSSELL, R. AND TOOMEY, F. (1994) Statistical issues raised by the Bellcore data. *Proceedings of the 11th IEE UK Teletraffic Symposium*, Cambridge.
- [26] DUMAS, V. AND SIMONIAN, A. (2000) Asymptotic bounds for the fluid queue fed by subexponential ON/OFF sources. *Advances in Applied Probability* **32**, 244–255.

[27] EMBRECHTS, P. AND VERAVERBEKE, N. (1982) Estimates for the probability of ruin with special emphasis on the possibility of large claims. *Insurance: Mathematics and Economics* **1**, 55–72.

[28] EMBRECHTS P., KLÜPPELBERG, C. AND MIKOSCH T. (1997) *Modelling Extremal Events for Insurance and Finance*. Springer-Verlag, Berlin.

[29] ERRAMILLI, A., NARAYAN, O. AND WILLINGER W. (1996) Experimental queueing analysis with long-range dependent packet traffic. *IEEE/ACM Transactions on Networking* **4**, 209–223.

[30] FELDMANN, A., GILBERT, A.C., WILLINGER, W. AND KURTZ, T.G. (1998) The changing nature of network traffic: scaling phenomena. *ACM Computer Communication Review* **28**, 5–29.

[31] FELDMANN, A., GILBERT, A.C. AND WILLINGER, W. (1998) Data networks as cascades: investigating the multifractal nature of Internet WAN traffic. *Proceedings of the ACM SIGCOMM '98 Conference*. Vancouver, 25–38.

[32] FELLER, W. (1951) The asymptotic distribution of the range of sums of independent random variables. *Annals of Mathematical Statistics* **22**, 427–432.

[33] FELLER, W. (1971) *An Introduction to Probability and Its Applications, Volume II, Second Edition*. Wiley, New York.

[34] FOWLER, H.J. AND LELAND, W.E. (1991) Local area network traffic characteristics, with implications for broadband network congestion management. *IEEE Journal on Selected Areas in Communications* **9**, 1139–1149.

[35] GEWEKE, J. AND PORTER-HUDAK, S. (1983) The estimation and application of long memory time series models. *Journal of Time Series Analysis* **4**, 221–238.

[36] GILBERT, A.C., WILLINGER, W. AND FELDMANN, A. (1999) Scaling analysis of conservative cascades, with applications to network traffic. *IEEE Transactions on Information Theory* **45**, 971–991.

[37] GNEDENKO, B.V. AND KOLMOGOROV, A.N. (1954) *Limit Theorems for Sums of Independent Random Variables*. Addison-Wesley, Cambridge (U.S.A.).

[38] GOGL, H. (2000) *Measurement and Characterization of Traffic Streams in High-Speed Wide Area Networks*. Ph.D. thesis, Institut für Informatik, Technische Universität München.

[39] GRASSE, M., FRATER, M.R. AND ARNOLD, J.F. (2000) Testing VBR video traffic for stationarity. *IEEE Transactions on Circuits and Systems for Video Technology* **10**, 448–459.

[40] GROSSGLAUSER, M. AND BOLOT, J-C. (1996) On the relevance of long-range dependence in network traffic. *Proceedings of the ACM SIGCOMM '96 Conference*. Stanford, California, 15–24.

[41] GUERIN, C.A., NYBERG, H., PERRIN, O., RESNICK, S., ROOTZÉN, H. AND STĂRICĂ, C. (2000) Empirical testing of the infinite source Poisson data traffic model. Technical Report TR1257, Cornell University, ORIE.

[42] GUO, L., CROVELLA, M. AND MATTÀ, I. (2000) TCP congestion control and heavy tails. Technical Report TR-2000-017, Boston University, Computer Science Department.

[43] GUT, A. (1988) *Stopped Random Walks: Limit Theorems and Applications*. Springer-Verlag, New York.

[44] GUT, A. AND JANSON, S. (1999) Tightness and weak convergence for jump processes. Technical Report 1999:9, Uppsala University, Department of Mathematics.

[45] HEATH, D., RESNICK, S.I. AND SAMORODNITSKY, G. (1998) Heavy tails and long range dependence in ON/OFF processes and associated fluid models. *Mathematics of Operations Research* **23**, 145–165.

[46] HEYMAN, D.P. (1998) Some issues in performance modeling of data teletraffic. *Performance Evaluation* **34**, 227–247.

[47] HURST, H.E. (1951) Long-term storage capacity of reservoirs. *Transactions of the American Society of Civil Engineers* **116**, 770–799.

[48] IBRAGIMOV, I.A. AND LINNIK, Y.V. (1971) *Independent and Stationary Sequences of Random Variables*. Wolters-Noordhoff, Groningen.

[49] INTERNET TRAFFIC ARCHIVE. <http://www.acm.org/sigcomm/ITA>.

[50] JACOBSON, V. (1988) Congestion avoidance and control. *Proceedings of the SIGCOMM '88 Conference*, Stanford, California, 314–329.

[51] JELENKOVIĆ, P. AND LAZAR, A. (1999) Asymptotic results for multiplexing subexponential on/off sources. *Advances in Applied Probability* **31**, 394–421.

[52] KALLENBERG, O. (1986) *Random Measures*. 4th edition, Akademie-Verlag, Berlin.

[53] KALLENBERG, O. (1996) Improved criteria for distributional convergence of point processes, *Stochastic Processes and their Applications* **64**, 93–102.

[54] KLEMEŠ, V. (1974) The Hurst phenomenon: a puzzle? *Water Resources Research* **10**, 675–688.

[55] KONSTANTOPOULOS, T. AND LIN, S. (1998) Macroscopic models for long-range dependent network traffic. *Queueing Systems. Theory and Applications* **28**, 215–243.

[56] KÜNSCH, H. (1986) Discrimination between monotonic trends and long-range dependence. *Journal of Applied Probability* **23**, 1025–1030.

[57] KURTZ, T.G. (1996) Limit theorems for workload input models. In: KELLY, F.P., ZACHARY, S. AND ZIEDINS, I. (editors). *Stochastic Networks: Theory and Applications*. Clarendon Press, Oxford, U.K., 339–366.

[58] LAMPERTI, J. (1962) Semi-stable stochastic processes. *Transactions of the American Mathematical Society* **104**, 62–78.

[59] LELAND, W. AND WILSON, D. (1991) High time-resolution measurement and analysis of LAN traffic: implications for LAN interconnection. *Proceedings of the IEEE Infocom '91*, Bal Harbour, Florida, 1360–1366.

[60] LELAND, W., TAQQU, M.S., WILLINGER, W. AND WILSON, D. (1994) On the self-similar nature of Ethernet traffic (extended version). *IEEE/ACM Transactions on Networking* **2**, 1–15.

[61] LEVY, J.B. AND TAQQU, M.S. (2000) Renewal reward processes with heavy-tailed interrenewal times and heavy-tailed rewards. *Bernoulli* **6**, 23–44.

[62] MANDELBROT, B.B. AND VAN NESS, J.W. (1968) Fractional Brownian motions, fractional noises and applications. *SIAM Review* **10**, 422–437.

[63] MANDELBROT, B.B. (1975) Limit theorems on the self-normalized range for weakly and strongly dependent processes. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete* **31**, 271–285.

[64] MANDELBROT, B.B. (1982) *The Fractal Geometry of Nature*. W.H. Freeman, New York.

[65] MIKOSCH, T. AND STEGEMAN, A. (1999) The interplay between heavy tails and rates in self-similar network traffic. IWI-preprint 99-5-07, University of Groningen, Department of Mathematics.

[66] MIKOSCH, T., RESNICK, S.I., ROOTZÉN, H. AND STEGEMAN, A.W. (2001) Is network traffic approximated by stable Lévy motion or fractional Brownian motion? *Annals of Applied Probability*, to appear.

[67] MIKOSCH, T. AND STĂRICĂ, C. (1999) Change of structure in financial time series, long range dependence and the GARCH model. IWI-preprint 99-5-06, University of Groningen, Department of Mathematics.

[68] PARK, K., KIM, G. AND CROVELLA, M. (1996) On the relationship between file sizes, transport protocols and self-similar network traffic. *Proceedings of the IEEE International Conference on Network Protocols*, 171–180.

[69] PARK, K., KIM, G. AND CROVELLA, M. (1997) On the effect of traffic self-similarity on network performance. *Proceedings of the SPIE International Conference on Performance and Control of Network Systems*, 296–310.

[70] PAXSON, V. AND FLOYD, S. (1995) Wide area traffic: the failure of Poisson modeling. *IEEE/ACM Transactions on Networking* **3**, 226–244.

[71] PETROV, V.V. (1975) *Sums of Independent Random Variables*. Springer-Verlag, Berlin.

[72] PETROV, V.V. (1995) *Limit Theorems of Probability Theory*. Oxford University Press, Oxford.

[73] PIPIRAS, V. AND TAQQU, M.S. (2000) The limit of a renewal-reward process with heavy-tailed rewards is not a linear fractional stable motion. *Bernoulli* **6**, 607–614.

[74] PIPIRAS, V., TAQQU, M.S. AND LEVY, J.B. (2001) Slow, fast and arbitrary growth conditions for renewal reward processes when the renewals and the rewards are heavy-tailed. Preprint, Boston University, Department of Mathematics.

[75] POLLARD, D. (1984) *Convergence of Stochastic Processes*. Springer-Verlag, Berlin.

[76] PRIESTLEY, M.B. (1981) *Spectral Analysis and Time Series, Volume 1*. Academic Press, London.

[77] RESNICK, S.I. (1987) *Extreme Values, Regular Variation and Point Processes*. Springer-Verlag, New York.

[78] RESNICK, S.I. (1992) *Adventures in Stochastic Processes*. Birkhauser, Boston.

[79] RESNICK, S.I. (1997) Heavy tail modeling and teletraffic data. *Annals of Statistics* **25**, 1805–1869.

[80] RESNICK, S. AND VAN DEN BERG, E. (2000) Weak convergence of high-speed network traffic models. *Journal of Applied Probability* **37**, 575–597.

[81] RESNICK, S., SAMORODNITSKY, G., GILBERT, A. AND WILLINGER W. (1999) Wavelet analysis of conservative cascades. Technical Report TR1243, Cornell University, ORIE.

[82] RESNICK, S. AND ROOTZÉN, H. (2000) Self-similar communication models and very heavy tails. *Annals of Applied Probability* **10**, 753–778.

[83] RESNICK, S. AND SAMORODNITSKY, G. (2001) Limits of on/off hierarchical product models for data transmission. Technical Report TR1281, Cornell University, ORIE.

[84] RIEDI, R.H. AND LÉVY-VÉHEL, J. (1997) Multifractal properties of TCP traffic: a numerical study. Technical Report 3129, INRIA Rocquencourt, France.

[85] RIEDI, R.H., CROUSE, M.S., RIBEIRO, V.J. AND BARANIUK, R.G. (1999) A multifractal wavelet model with application to network traffic. *IEEE Transactions on Information Theory* **45**, 992–1019.

[86] RIEDI, R.H. AND WILLINGER, W. (2000) Toward an improved understanding of network traffic dynamics. In: PARK, K. AND WILLINGER W. (editors). *Self-Similar Network Traffic and Performance Evaluation*. Wiley, New York, 507–530.

[87] ROBINSON, P.M. (1995) Log-periodogram regression of time series with long range dependence. *Annals of Statistics* **23**, 1048–1072.

[88] RYU, B. AND ELWALID, A. (1996) The importance of long-range dependence of VBR video traffic in ATM traffic engineering: myths and realities. *Proceedings of the ACM SIGCOMM '96 Conference*. Stanford, California, 3–14.

[89] SAMORODNITSKY, G. AND TAQQU, M.S. (1994) *Stable Non-Gaussian Random Processes. Stochastic Models with Infinite Variance*. Chapman & Hall, New York.

[90] SHORACK, G.R. AND WELLNER, J.A. (1986) *Empirical Processes with Applications to Statistics*. Wiley, New York.

[91] SKOROHOD, A.V. (1956) Limit theorems for stochastic processes. *Theory of Probability and Its Applications* **1**, 261–290.

[92] STEGEMAN, A. (1998) *Modeling Traffic in High-Speed Networks by ON/OFF Models*. Masters Thesis, University of Groningen, Department of Mathematics.

[93] STEGEMAN, A. (2000) Heavy tails versus long-range dependence in self-similar network traffic. *Statistica Neerlandica* **54**, 293–314.

[94] STEGEMAN, A. (2002) Extremal behavior of heavy-tailed ON-periods in a superposition of ON/OFF processes. *Advances in Applied Probability*, to appear.

[95] STEGEMAN, A. (2001) Non-stationarity versus long-range dependence in computer network traffic measurements. IWI-preprint 2001-5-01, University of Groningen, Department of Mathematics.

[96] TANENBAUM, A.S. (1996) *Computer Networks, 3rd Edition*. Prentice-Hall, London.

[97] TAQQU, M.S. (1975) Weak convergence to fractional Brownian motion and to the Rosenblatt process. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete* **31**, 287–302.

[98] TAQQU, M.S. AND LEVY, J.B. (1986) Using renewal processes to generate long-range dependence and high variability. In: EBERLEIN, E. AND TAQQU, M.S. (editors). *Dependence in Probability and Statistics*. Birkhauser, Boston, 73–89.

[99] TAQQU, M.S. AND TEVEROVSKY, V. (1995) Estimators for long-range dependence: an empirical study. *Fractals* **3**, 785–798.

[100] TAQQU, M.S. AND TEVEROVSKY, V. (1996) On estimating the intensity of long-range dependence in finite and infinite variance time series. In: ADLER, R., FELDMAN, R. AND TAQQU, M.S. (editors). *A Practical Guide To Heavy Tails: Statistical Techniques and Applications*. Birkhauser, Boston, 177–217.

[101] TAQQU, M.S., WILLINGER, W. AND SHERMAN, R. (1997) Proof of a fundamental result in self-similar traffic modeling. *Computer Communications Review* **27**, 5–23.

[102] TEVEROVSKY, V. AND TAQQU, M.S. (1997) Testing for long-range dependence in the presence of shifting means or a slowly decaying trend, using a variance-type estimator. *Journal of Time Series Analysis* **18**, 279–304.

[103] TUAN, T. AND PARK, K. (1999) Multiple time scale congestion control for self-similar network traffic. Preprint, Purdue University, Department of Computer Sciences, Network Systems Lab.

[104] VERES, A. AND BODA, M. (2000) The chaotic nature of TCP congestion control. *Best paper of the IEEE INFOCOM 2000 Conference*. Dan Panorama Hotel, Tel Aviv.

[105] VERES, A., KENESI, Zs., MOLNÁR, S. AND VATTAY, G. (2000) On the propagation of long-range dependence in the Internet. *Proceedings of the ACM SIGCOMM 2000 Conference*. Grand Hotel, Stockholm, 243–254.

[106] WILLINGER, W., TAQQU, M.S., SHERMAN, R. AND WILSON, D.V. (1995) Self-similarity through high-variability: statistical analysis of Ethernet LAN traffic at the source level (extended version). *Computer Communications Review* **25**, 100–113.

[107] WILLINGER W. AND PAXSON, V. (1998) Where mathematics meets the Internet. *Notices of the American Mathematical Society* **45**, 961–970.

[108] WHITT, W. (1970) Weak convergence of probability measures on the function space $C[0, \infty)$. *Annals of Mathematical Statistics* **41**, 939–944.

- [109] ZWART, A.P., BORST, S.C. AND MANDJES, M.R.H. (2000) Exact asymptotics for fluid queues fed by multiple heavy-tailed on-off flows. SPOR-Report 2000-14, Eindhoven University of Technology.
- [110] ZWART, A.P. (2001) *Queueing Systems with Heavy Tails*. Ph.D. thesis, Eindhoven University of Technology.
- [111] ZYGMUND, A. (1968) *Trigonometric Series, Volumes I and II*. Cambridge University Press, Cambridge (U.K.).

Summary

Since the beginning of the 1990s a growing number of computer networks from all over the world has been linked together. Nowadays, they constitute the global network called the *Internet*. The Internet has hundreds of millions of users, from private individuals and corporations to government officials and scientists, and offers a broad range of applications: *e-mail*, *newsgroups*, *remote login*, *file transfer*, *audio and video streams*, *chatrooms* and the popular *World Wide Web*. Although the Internet and computer networks in general have many benefits, some technological difficulties have been encountered. As surely everyone has experienced, the operation of transferring data between two computers in a network can be quite troublesome. At busy hours, when there are a lot of active users, the network can be congested, resulting in long transmission delays or difficulties in establishing a connection between the source and destination computer.

In order to get a better understanding of the dynamics of the data traffic in computer networks, a number of empirical studies of network traffic measurements has been conducted. As a benchmark for comparison voice traffic in the traditional telephone system has been used. In the telephone system call arrivals can be modeled by a Poisson process, i.e. with exponential inter-arrival times. Moreover, the distribution of call lengths has an exponentially bounded tail. From an engineering perspective these are very convenient properties, since the long-term arrival rate of the Poisson process, together with the mean call length, roughly determines the capacity of the network that guarantees reliable telephone communication.

Measurements of computer network traffic

The situation in computer networks, however, has been found to be very different. Computer network traffic has been studied at two levels: the *application level* and the *packet level*. At the application level file sizes, connection durations and transmission times are the main subjects of analysis. Instead of exponential their distributions appear to be *heavy-tailed*, i.e. $P(X > x) \sim c x^{-\alpha}$, $x \rightarrow \infty$, with $c > 0$ and $\alpha \in (1, 2)$. Consequently, $\text{Var}(X) = \infty$ and extremely large values of X occur with non-negligible probability. At the packet level the so-called *workload* of the network is measured. When a file is sent from a source to a destination computer, it is decomposed into small *packets* which are sent through

the network cables. After arriving at the destination computer, the packets are put together again and the original file is reconstructed. The workload on a cable or link in the network is measured by counting the number of packets or bytes passing the measurement point in a small time interval, e.g. 1 second. In this way, the workload per second is determined. Even for large time intervals, the dependence in the series of workload measurements appears to be rather strong: at large lags the sample autocorrelations still seem significantly different from zero. This phenomenon is often referred to as *long-range dependence*. This would indicate that random cycles of arbitrary length are present in the workload data. Another striking feature is that when the time interval used for measuring the workload is increased, the relative variability of the workload remains roughly the same, or, in other words, the workload shows a similar burstiness across a wide range of time scales. This property has been observed for time intervals ranging from 0.01 up to 100 seconds and resembles, in some sense, the theoretical notion of distributional *self-similarity*. Unlike voice traffic, computer network traffic does not smooth out when viewed at increasingly larger time scales.

Heavy tails, long-range dependence and self-similarity are believed to be present in traffic measurements on networks of different scales providing different applications, from the late 1980s till the present day. Therefore, these three features are regarded as *traffic invariants*. On the whole, this implies that computer network traffic behaves rather erratic compared to voice traffic in the telephone system, and, hence, that computer networks are a great challenge to the engineer and the scientist.

Non-stationarity versus long-range dependence

By definition a stochastic process exhibiting long-range dependence is stationary, i.e. its underlying distribution does not change during the time the process is observed. However, no general test for the stationarity of an observed time series is available. Also, the graphical methods that are often used to detect long-range dependence in a time series are not very reliable. It is well-known that these graphical methods can interpret non-stationarities like shifts in the mean or a slowly decaying trend as the presence of long-range dependence. In Chapter 3 of this thesis we show that a realization from a non-stationary ARIMA($p,1,q$) process, with appropriate parameter values, can indeed exhibit long-range dependence in this ‘graphical’ sense. In Chapter 4 we analyze series of workload measurements in various computer networks and find that most of them can be modeled by an ARIMA($p,1,q$) process, with small p and q . This shows that, when using graphical methods, it is virtually impossible to distinguish between non-stationarity and long-range dependence in an observed time series. Here however, given the complicated nature of a computer network, with applications and connections being activated and terminated during the measurement period, the option of non-stationarity is probably the most reasonable one.

Modeling computer network traffic

An attempt to give a ‘physical’ explanation for the observed traffic characteristics has been made by using a mathematical modeling approach. Two simple models have been proposed which both use the assumption of heavy-tailed transmission times to explain the long-range dependence in the workload of the network. Also, it is shown that the centered and properly normalized cumulative workload can be approximated, in some sense, by a self-similar process. One of these two models, the *ON/OFF model* introduced by Willinger et al. [106], is the subject of Chapters 5 and 6 of this thesis. In this model, traffic is generated by M independent and identically distributed *ON/OFF sources*. If a source is ON it transmits data at unit rate, e.g. 1 byte per time unit. If it is OFF it remains silent. In this way, every individual ON/OFF source generates a binary *ON/OFF process*. The lengths of periods in which a source is ON, the *ON-periods*, are independently drawn from a heavy-tailed distribution. Analogously, the *OFF-periods* are also heavy-tailed. The sequences of ON- and OFF-periods are assumed independent. It has been shown by Heath et al. [45] that the stationary version of the ON/OFF process of an individual source exhibits long-range dependence. Using independence, the same is true for the total workload, i.e. the superposition of the M ON/OFF processes.

In Willinger et al. [106] it is shown that the centered cumulative workload up to time T , when properly normalized, converges in finite dimensional distributions to fractional Brownian motion if first $M \rightarrow \infty$ and then $T \rightarrow \infty$. In Taqqu et al. [101] the limits are reversed and a different normalization is used to obtain stable Lévy motion as limit process. Both fractional Brownian motion and stable Lévy motion are self-similar, but their dependence structures are totally different. The increment sequence, at equidistant instants of time, of fractional Brownian motion is stationary and exhibits long-range dependence (thus preserving the long-range dependence in the pre-limit workload), while the increments of stable Lévy motion are independent. In Chapter 5 of this thesis we consider simultaneous limit regimes in which M is a non-decreasing function of T , converging to infinity as $T \rightarrow \infty$. We show that when M grows faster than some ‘critical rate’ fractional Brownian motion is obtained in the limit. On the other hand, if M grows slower than this ‘critical rate’ stable Lévy motion appears as limit process.

In Chapter 6 we use the framework of the ON/OFF model to study the number of ON-periods up to time T exceeding a high threshold. Again, we consider simultaneous limits of M and T . Moreover, also the threshold depends on T . We distinguish between the ‘slow’ and ‘fast’ growth conditions on M . Although different approaches are needed, in both cases we are able to show that the number of exceedances converges to a Poisson random variable if the threshold satisfies a balancing condition guaranteeing a constant average number of exceedances in the limit. We also show that if the threshold grows slower than this ‘balancing rate’, the number of exceedances satisfies the Central Limit Theorem.

Samenvatting

Vanaf het begin van de negentiger jaren van de vorige eeuw zijn steeds meer computernetwerken over de hele wereld aan elkaar gekoppeld. Tegenwoordig vormen zij het wereldwijde netwerk dat bekend staat als het *Internet*. Het Internet heeft honderden miljoenen gebruikers, van individuele burgers en bedrijven, tot overheidsinstanties en wetenschappers. Deze gebruikers staan een breed scala aan toepassingen ter beschikking: *e-mail*, *nieuwsgroepen*, *remote login*, *file transfer*, *audio en video stromen*, *chatboxen* en het populaire *World Wide Web*. Hoewel de voordelen van het Internet en computernetwerken in het algemeen legio zijn, moeten er nog wel wat technische problemen overwonnen worden. Zoals iedereen waarschijnlijk wel eens heeft ervaren verloopt het sturen van data tussen twee computers in een netwerk lang niet altijd even soepel. Wanneer er veel gebruikers op het netwerk aanwezig zijn, kan het netwerk verstoppt raken en is het mogelijk dat een gevraagde verbinding tussen de zendende en ontvangende computer niet tot stand wordt gebracht.

Om een beter begrip te krijgen van de dynamiek van het dataverkeer in computernetwerken, zijn metingen aan dit dataverkeer in een aantal statistische studies geanalyseerd. De resultaten zijn vergeleken met de eigenschappen van gespreksverkeer in het traditionele telefoonnetwerk. In het telefoonverkeer is het redelijk om aankomsttijden van gesprekken te modelleren als een Poisson proces, d.w.z. met exponentiële tussenaankomsttijden. Bovendien duidt de empirische verdeling van gesprekslengtes op het bestaan van een exponentieel begrenste staart. Gezien vanuit het standpunt van de ingenieur zijn deze eigenschappen bijzonder aantrekkelijk, want het gemiddelde van het Poisson proces en de gemiddelde gespreksduur bepalen dan samen in grote lijnen de capaciteit van het netwerk die nodig is om de binnenkomende telefoongesprekken te verwerken.

Metingen van dataverkeer in computernetwerken

Echter, bij computernetwerken ziet het er totaal anders uit. Het dataverkeer in deze netwerken is op twee niveaus bestudeerd: het *applicationiveau* en het *pakketniveau*. Op het applicationiveau zijn de lengtes van bestanden en de duur van connecties en transmissies de meest geanalyseerde variabelen. In plaats van exponentieel lijken hun verdelingen *zwaarstaartig*, d.w.z. $P(X > x) \sim cx^{-\alpha}$, $x \rightarrow \infty$, met $c > 0$ en $\alpha \in (1, 2)$. Dit heeft tot gevolg dat $\text{Var}(X) = \infty$ en dat extreem grote waarden van X optreden met een niet te verwaarlozen kans.

Op het pakketniveau wordt de zogenaamde *werklast* van het netwerk gemeten. Wanneer een bestand van een zender naar een ontvanger wordt gestuurd, wordt deze opgesplitst in kleine *pakketjes* die door de kabels van het netwerk worden verzonden. Na aankomst bij de ontvangende computer worden de pakketjes weer samengevoegd tot het oorspronkelijke bestand. De werklast op een kabel of verbinding in het netwerk wordt gemeten door het aantal pakketjes of bytes te tellen dat het meetpunt passeert in een klein tijdsinterval, bijvoorbeeld 1 seconde. Op deze manier wordt per seconde de werklast bepaald. De afhankelijkheid in de reeks van werklastmetingen schijnt, ook op afstand, nogal sterk te zijn: bij grote tussenperioden lijken de autocorrelaties nog steeds significant verschillend van nul. Dit fenomeen wordt *long-range dependence* genoemd. Dit zou kunnen wijzen op het voorkomen van periodiciteiten met willekeurig lange perioden in de werklastdata. Een andere opzienbarende eigenschap is dat de relatieve variabiliteit van de werklast ruwweg onveranderd blijft wanneer het tijdsinterval, dat voor de werklastmetingen wordt gebruikt, vergroot wordt. Anders gezegd: de werklast vertoont eenzelfde soort grilligheid over een groot aantal tijdsschalen. Deze eigenschap is waargenomen voor tijdsintervallen van 0.01 tot 100 seconden en vertoont, op een bepaalde manier, gelijkenis met het theoretische concept van *self-similarity* van delingen. In tegenstelling tot verkeer in het telefoonnetwerk, vlakt de werklast in computernetwerken niet uit wanneer deze op steeds grotere tijdsschalen wordt bekeken.

Er wordt algemeen aangenomen dat zware staarten, long-range dependence en self-similarity voorkomen in metingen van dataverkeer bij netwerken van verschillende grootte met verschillende applicaties, vanaf 1989 tot nu. Daarom worden deze drie eigenschappen ook wel beschouwd als *invarianten van dataverkeer*. Samenvattend betekent dit dat dataverkeer in computernetwerken zich een stuk grilliger en onregelmatiger gedraagt dan gespreksverkeer in het telefoonnetwerk. Dus vormen computernetwerken een grote uitdaging voor de ingenieur en de wetenschapper.

Niet-stationariteit of long-range dependence

Per definitie moet een stochastisch proces met long-range dependence stationair zijn, d.w.z. dat de onderliggende verdeling onveranderlijk is gedurende de tijd dat het proces geobserveerd wordt. Echter, er bestaat geen algemene toets voor de stationariteit van een waargenomen tijdsreeks. Ook zijn de grafische methoden, die vaak worden gebruikt om long-range dependence in een tijdsreeks op te sporen, niet erg betrouwbaar. Het is bekend dat deze grafische methoden niet-stationariteiten als sprongen in het gemiddelde of een langzaam damende trend kunnen interpreteren als long-range dependence. In Hoofdstuk 3 van dit proefschrift laten we zien dat een realisatie van een niet-stationair ARIMA($p,1,q$) proces, met geschikte parameterwaarden, inderdaad long-range dependence vertoont in deze ‘grafische’ betekenis. In Hoofdstuk 4 analyseren we reeksen van werklastmetingen bij verschillende computernetwerken en komen we tot de conclusie dat de meeste reeksen gemodelleerd kunnen worden als een ARIMA($p,1,q$) proces, met kleine waarden p en q . Dit toont aan dat het schier onmogelijk is om,

op basis van grafische methoden, onderscheid te maken tussen niet-stationariteit en long-range dependence in een geobserveerde tijdreeks. Echter, gegeven het complexe karakter van een computernetwerk, met applicaties en connecties die geactiveerd en beëindigd worden gedurende de meetperiode, is de optie van niet-stationariteit hier waarschijnlijk de meest redelijke.

Het modelleren van dataverkeer in computernetwerken

Door gebruik te maken van wiskundige modellen wordt een poging gedaan om een ‘fysische’ verklaring te geven voor de vermelde eigenschappen van het dataverkeer. Dit heeft twee eenvoudige modellen opgeleverd die beide de aanname van zwaarstaartige transmissieduren gebruiken om de long-range dependence in de werklast van het netwerk te verklaren. Ook is aangetoond dat de gecentreerde cumulatieve werklast kan worden benaderd door een self-similar proces, mits de juiste normalisatie wordt gebruikt. Eén van deze twee modellen, het *AAN/UIT model* van Willinger et al. [106], is het onderwerp van de Hoofdstukken 5 en 6 van dit proefschrift. In dit model wordt dataverkeer gegenereerd door M onafhankelijke en identiek verdeelde *AAN/UIT bronnen*. Wanneer een bron AAN is, stuurt deze 1 eenheid data (bijvoorbeeld 1 byte) per tijdseenheid het netwerk op. Een bron stuurt niks wanneer deze UIT is. Op deze manier genereert elke AAN/UIT bron een binair *AAN/UIT proces*. De lengtes van de perioden dat een bron AAN is, de *AAN-perioden*, worden onafhankelijk van elkaar getrokken uit een zwaarstaartige verdeling. Hetzelfde geldt voor de *UIT-perioden*. De reeksen van AAN- en UIT-perioden worden onafhankelijk verondersteld. Door Heath et al. [45] is bewezen dat de stationaire versie van het AAN/UIT proces van een individuele bron long-range dependence vertoont. Door de onafhankelijkheid van de bronnen geldt dat ook voor de totale werklast, d.w.z. de superpositie van de M AAN/UIT processen.

In Willinger et al. [106] wordt aangetoond dat de eindig-dimensionale verdelingen van de gecentreerde cumulatieve werklast tot tijd T , mits juist genormaliseerd, naar die van *fractional Brownian motion* convergeren als eerst $M \rightarrow \infty$ en dan $T \rightarrow \infty$. Taqqu et al. [101] laten zien dat wanneer eerst $T \rightarrow \infty$ en dan $M \rightarrow \infty$, en een andere normalisatie wordt gebruikt, het limietproces *stable Lévy motion* is. Zowel fractional Brownian motion als stable Lévy motion zijn self-similar, maar hun afhankelijkheidsstructuren zijn totaal verschillend. De reeks van aanwassen, beschouwd over gelijke tijdsintervallen, van fractional Brownian motion is stationair en vertoont long-range dependence (dus de long-range dependence van de werklast blijft ongeschonden in de limiet), terwijl de aanwassen van stable Lévy motion onafhankelijk zijn. In Hoofdstuk 5 van dit proefschrift beschouwen we simultane limieten in M en T , d.w.z. we veronderstellen dat M een niet-dalende functie van T is die naar oneindig convergeert als $T \rightarrow \infty$. We tonen aan dat wanneer M sneller groeit dan een bepaalde ‘kritieke snelheid’, fractional Brownian motion als limietproces wordt verkregen. In het geval dat M langzamer groeit dan deze ‘kritieke snelheid’ is stable Lévy motion het limietproces.

In Hoofdstuk 6 gebruiken we het raamwerk van het AAN/UIT model om

het aantal AAN-perioden tot tijd T , dat een hoge drempelwaarde overschrijdt, te bestuderen. We beschouwen weer simultane limieten in M en T . Ook de drempelwaarde is een functie van T . We maken onderscheid tussen de gevallen waarin M ‘snel’ en ‘langzaam’ groeit. Hoewel er verschillende manieren van aanpak nodig zijn, kunnen we in beide gevallen aantonen dat het aantal overschrijdingen naar een Poisson stochast convergeert. Hierbij is het noodzakelijk dat de drempelwaarde aan een evenwichtsconditie voldoet die ervoor zorgt dat het aantal overschrijdingen in de limiet constante is. Ook laten we zien dat wanneer de drempelwaarde langzamer groeit dan deze ‘evenwichtssnelheid’, het aantal overschrijdingen aan de Centrale Limiet Stelling voldoet.

Dankwoord

Gedurende de vier jaar waarin ik aan het onderzoek in dit proefschrift heb gewerkt, heb ik mogen profiteren van een aantal stimulerende kontakten, zowel in de professionele als in de persoonlijke sfeer. Ik wil graag enkele van deze mensen hier bedanken.

Te beginnen met mijn begeleider en promotor Thomas Mikosch. Door onze prettige samenwerking tijdens het schrijven van mijn scriptie aarzelde ik geen moment toen je voorstelde om er nog vier jaar aan vast te plakken. Je was altijd bereid mij te helpen met wiskundige probleempjes. Ook voor advies over het presenteren van de onderzoeksresultaten, zowel mondelijk als schriftelijk, kon ik bij jou terecht. Zelfs nadat je je carrière voortzette in Kopenhagen heb je nog een paar ‘tussenstops’ in Groningen gemaakt. Bedankt voor alles!

Mijn promotor Herold Dehling wil ik graag bedanken voor een goede samenwerking wat betreft de afhandeling van de formele aspecten rond de promotie, zijn uitnodiging om in Bochum een voordracht te komen houden en de zinvolle discussie over de ‘symbolic embedding’ die ik voor dit boekje bedacht heb. Voor dit laatste ben ik ook dank verschuldigd aan Willem Schaafsma en Marius van der Put. Verder bedank ik de beoordelingscommissie voor hun waardevolle commentaar. Voor een prettig verblijf in Gothenburg gaat mijn dank uit naar de organisatie van de Semstat conferentie 2001.

De financiële middelen die dit onderzoek mogelijk maakten, zijn opgebracht door de Nederlandse Organisatie voor Wetenschappelijk Onderzoek (NWO) en de Stichting Wiskunde-onderzoek Nederland (SWON). Hiervoor bedank ik het NWO en SWON hartelijk. Het NWO wordt bovendien bedankt voor de efficiënte samenwerking. Ook wil ik het Instituut voor Wiskunde en Informatica (IWI) aan de Rijksuniversiteit Groningen bedanken voor de gelegenheid die mij is geboden om er mijn proefschrift te schrijven.

Voor het ter beschikking stellen van een gedeelte van zijn metingen van ATM-dataverkeer wil ik Helmut Gogl bedanken. Marius de Vink en Titus van ’t Veer ben ik dank verschuldigd voor het downloaden en wegschrijven van de enorme datasets van verkeer in computernetwerken. Cedric Thieulot wil ik bedanken voor het installeren van alle software op mijn thuis-PC.

Voor de gezellige sfeer in de groep van statistiek en stochastiek wil ik graag mijn (al dan niet voormalige) kamergenoten en collega's Bojan, Casper, Daniël, Diemer, Evgeny, Hui en Mook bedanken.

Tenslotte gaat mijn dank uit naar enkele personen aan wie ik veel steun heb ervaren in de afgelopen jaren. Hierbij wil ik met name Eddy, Johan, Mart, Miranda, Rob en Sietse noemen. Thanks for being who you are!

Alwin, februari 2002.



Reminiscensce