Institut für Informatik
der Technischen Universität München
Lehrstuhl für Informatik VIII

Performance Analysis of Intermediate Systems Serving Aggregated ON/OFF Traffic with Long-Range Dependent Properties

Hans-Peter Schwefel

Vollständiger Abdruck der von der Fakultät für Informatik der Technischen Universität München zur Erlangung des akademischen Grades eines

Doktors der Naturwissenschaften (Dr. rer. nat.)

genehmigten Dissertation.

Vorsitzender:	Univ.-Prof. T. Nipkow, Ph.D.
Prüfer der Dissertation:	
1.	Univ.-Prof. Dr. E. Jessen
2.	Prof. L. Lipsky, Ph.D. University of Connecticut, USA

Die Dissertation wurde am 19.9.2000 bei der Technischen Universität München eingereicht und durch die Fakultät für Informatik am 09.11.2000 angenommen.

Performance Analysis of Intermediate Systems Serving Aggregated ON/OFF Traffic with Long-Range Dependent Properties

Hans-Peter Schwefel

Institut für Informatik
Lehrstuhl für Rechnerkommunikation
Technische Universität München, Germany
schwefel@in.tum.de

Keywords

Telecommunication Networks, Network Planning, Quality of Service, ON/OFF-Models, Self-Similarity, Long-Range Dependence, N-Burst Process, Power-Tail Distributions, Truncated Power-Tails, Transient Analysis.

Abstract

Although it is now widely recognized that network traffic frequently exhibits so-called *Long-Range Dependent* (LRD) properties, the impact of such properties on performance is still not well understood. This thesis presents a thorough discussion of a particular class of traffic models, namely the aggregation of ON/OFF traffic sources, called N-Burst. A special family of *Matrix-Exponential* (or Phase-Type) distributions – called *Truncated Power-Tail* distributions – is used for the ON time distribution in order to mimic LRD properties, while still remaining tractable for queueing analysis via *Matrix-Analytic* methods.

In order to make the model applicable to realistic scenarios, adequate procedures for parameter estimation are developed, which are then applied to a set of actual data from measurements of inter-cell times in an IP-over-ATM network.

The steady-state analysis of queueing models witch such LRD arrival processes reveals distinctively different behavior than for models without LRD properties: All performance parameters show several so-called blow-ups at particular, well-defined points in the parameter space of the model, at which performance deteriorates dramatically. The knowledge about the existence and the location of these blow-up points is crucial for any network design task. Applicable engineering rules for the purpose of performance oriented design are derived from the gained insights.

In many scenarios, the steady-state analysis of such LRD models does not provide a satisfactory description of the performance behavior. Therefore, methods for transient analysis are developed and applied. The discussion of the transient results leads to new characterizations of the fluctuations in the performance behavior of such LRD models.

Finally, a variation of the basic analytic queueing model is discussed in which the arrival process is throttled when congestion occurs in the network. The results of that analysis provide insights

into the behavior of elastic traffic, such as assumed by the *Transmission Control Protocol* (TCP) in the Internet.

The contributions of the thesis are twofold: First, contributions to the methodology for performance analysis using Matrix-Analytic models with Long-Range Dependent properties are made. Second, the results from the performance analysis of the rather general N-Burst model class provide new insights for the network design task. Although the thesis is directed at modeling and performance evaluation of telecommunication systems, the methodologies can easily be transfered to related areas, such as storage systems and CPU job-times, both of which are known to be subject to the impact of highly fluctuating loads.

Acknowledgment

First of all I would like to thank my two advisors, Professor Eike Jessen at Technische Universität München, and Professor Lester Lipsky at University of Connecticut. Both of them supported me greatly during the last three years. Special thanks to Lester Lipsky: I learned so much from him academically, and his great personality made working with him an extraordinary pleasant and valuable experience.

Parts of this work were performed while I enjoyed the hospitality of the Computer Science Department at the University of Connecticut and the Department of Network Design and Analysis at AT&T Labs in New Jersey. Thanks to both institutions.

In addition, I want to thank Deutsche Telekom AG who provided the funding for large parts of the work that is presented in this thesis. The work about TCP modeling in Chapter 9 was funded by a fellowship from the German Academic Exchange Agency (DAAD).

Last but not least, I would like to thank the following people for the interesting discussion that contributed to this thesis: Dr. Manfred Jobmann at TU München, Prof. Sören Asmussen at Lund University, Prof. Daryl Daley at Australian National University, Dr. Dan Heyman at AT&T Labs, Dr. V. Ramaswami at AT&T Research, Dr. Udo Krieger at Deutsche Telekom AG, Dr. Volker Fischer at IBM Consulting, Dr. Helmut Gogl at Siemens AG.

German Abstract (condensed form)

Verkehr in modernen, paketvermittelten Telekommunikationsnetzen weist Charakteristika auf, die in klassischen Verkehrsmodellen nicht adäquat berücksichtigt werden: In Messungen wurden sogenannte *Langzeitabhängigkeiten* (auch *Selbstähnlichkeit*) des Paketstromes festgestellt, deren Einfluß auf die Übertragungsqualität aber noch diskutiert wird. Diese Dissertation entwickelt eine realistische Klasse von Verkehrsmodellen, genannt N-Burst, die langzeitabhängigen Verkehr beschreiben kann. Aus der Diskussion einer Vielzahl stationärer und transienter Leistungsparameter in analytischen Warteschlangenmodellen werden detaillierte Schlußfolgerungen über das Leistungsverhalten solcher langzeitabhängiger Verkehre gewonnen, die beispielhaft auf reale Problemstellungen der Netzplanung angewendet werden. Dabei zeigt sich, daß in diesem Fall herkömmliche (stationäre) Leistungsparameter das Verhalten der Warteschlangenmodelle nur ungenügend beschreiben. Neue transiente Methoden werden deswegen eingeführt und angewendet.

Contents

Chapter 1
Overview and Background

1.1 Introduction

With emerging new applications, the requirements on modern telecommunication networks have changed. One of the currently unsolved problems is the demand of the users or applications for a certain transmission quality, often referred to as *Quality of Service* (QoS). Such QoS might be provided by careful analysis and estimation of the traffic in the system together with an effective methodology to design the network and its resources. Another possibility is the development of new transmission protocols, whose effectiveness has to be investigated.

In both cases, modern techniques in stochastic modeling are a necessity. In such models, the description of the traffic is a crucial part: It is now widely accepted that standard Poisson models do not adequately describe network traffic since they do not take into account its inherently bursty nature. This is accentuated by a series of measurements in the last decade (one of the earliest discussed by [Leland et al. 94]) that have shown that network traffic often exhibits so-called *self-similar* or *Long-Range Dependent* properties: It shows burstiness over a wide range of time-scales.

According to several measurements of network traffic, one such series done in [Gogl 00] at the entrance of a national backbone, multiplexed ON/OFF traffic with Long-Range Dependent (LRD) properties appears to be a good match to actual real data traffic. This thesis develops techniques for and provides results from the analysis of queueing models with such traffic as input. Thereby, insights are gained into how QoS problems at single 'bottleneck' switches in the network occur and what remedies are possible.

The remainder of this chapter provides the necessary background and an overview of existing work on traffic modeling. Chapter 2 introduces a very flexible class of ON/OFF models, called N-Burst, whose analytic queueing models are nevertheless tractable for performance analysis using matrix algebraic methods. The traffic of the N-Burst model shows Long-Range Dependent properties if a certain class of distributions, called *Power-Tail* (PT) distributions, are used for the ON-times. Chapter 3 discusses such distributions in detail with particular emphasis on a family of Phase-type distributions with PT characteristics.

For practical applications, it is necessary to estimate the parameters of the traffic model from a given set of measurements. In Chapter 4, a method for parameter estimation is developed and applied to actual measurements of inter-cell times at the entrance point to an IP-over-ATM backbone. The numerical examples that illustrate the performance behavior in the subsequent chapters are all based on the model parameters that are derived from real data in Chapter 4.

The performance analysis at the bottleneck switches starts off in Chapter 5 with an investigation of the steady-state performance parameters *mean Cell Delay* (mCD), *Cell Loss Probability* (CLP), and *Buffer Overflow Probability* (BOP). The analysis shows that QoS does not improve uniformly with, for instance, increasing service rate. Instead, well-defined so-called *blow-up points* are

observed at which the QoS parameters deteriorate dramatically. The location of the blow-up points is derived, and it is shown later in Chapter 7 how that knowledge can be used in network design.

The analysis of the steady-state behavior provides some fundamental insight into the peculiar behavior of aggregated ON/OFF traffic with LRD properties. However, the slowly decaying correlation in the cell-arrival process, together with the finite duration of the 4-8 busy hours in the daily profile, could have the effect that steady-state behavior for the performance-critical time period might never be observed in practice. The results of the transient analysis in Chapter 6 confirm such doubts. Furthermore, for data-transmission, it is generally desirable that the steady-state probability, BOP, of an overflow event be very low, in the order of 10^{-5} or less. However, the BOP does not express the correlated nature of overflow events, which is especially accentuated in systems with large buffers and by traffic with LRD properties. Additional transient parameters are discussed in Chapter 6 in order to provide a better description of the switch's behavior.

The results from the steady-state and transient performance analysis are used in Chapter 7 to discuss different scenarios whose solutions contribute to the network design task. Simple engineering rules are developed based on the asymptotic results of the analysis. Chapter 8 contains a summary of the results and their practical implications. Chapters 9 and 10 point out future directions and open problems. An important part of that is the discussion of models for so-called *elastic* traffic: in that scenario the incoming traffic adjusts to the congestion level in the network, thus it is no longer independent of the network state. An important example is the Internet traffic via the *Transmission Control Protocol* (TCP). Chapter 9 presents a viable approach for capturing the essential features of such a feedback between the network and the incoming traffic.

The body of the thesis concentrates on the model description and the discussion of the performance results. The algorithmic details of how the results are obtained can be found in the appendix. Note that the appendix is not just a compendium of previously known results, but the actual computational techniques are further (or even newly) developed with respect to the specific requirements of the analysis in this thesis.

1.2 Overview of Modeling Approaches

Traffic modeling is not a new problem but it has attracted much attention recently due to the rapid growth of the Internet and the increasing popularity of online access to distributed information. However, it has a long history that started with the early models for telephony traffic: M/G/C//C queueing models were used to obtain the so-called *call-blocking probability*, which is the probability that an incoming call has to be rejected since all C available channels are being used. Since the original use of telephony was for voice-calls, whose statistical properties are rather well-behaved and well-known, even the simpler M/M/C//C models with exponential call-holding times can be appropriate (for the steady-state probabilities of the number of active calls, the distribution of the call-holding times does not even matter in those models, see Appendix D.3).

In traditional telephony traffic, each call is assigned its own channel and that channel is used exclusively during the whole duration of the call. Such a so-called *circuit based* approach was the basis of the early telephony models of type M/G/C//C. Modern telecommunication networks use a more flexible technique which is based on the transmission of small data units, called *packet* or *cells*, instead. The packets of different connections can be transmitted on the same link and therefore the available transmission capacity can be utilized much more efficiently, in particular

when the individual connections show time-intervals with very little or no data-transmission at all. This is known as *multiplex* gain, which is one of the benefits of the packet-based approach.

The simplest packet based performance model for intermediate systems in telecommunication models is an M/M/1 queue: packets arrive at a single server unit according to a Poisson process. If the server is idle, the arriving packet obtains service for an exponentially distributed time before it leaves the system. If the server is busy, the packet is stored in a buffer with infinite capacity. Such simple M/M/1 models are in fact used in network planning tools for capacity planning in order to meet some constraints on the delay, see e.g. [Cahn 98]. The advantage of this model is its simplicity and the availability of explicit closed-form formulas for many of its performance parameters, see Appendix A. The main impact of the use of M/M/1 models in network design is that it predicts major performance problems when its utilization ρ approaches 100%. As a consequence, network components are designed to operate at utilization values between 40% and 60%, as a rule of thumb. However, the assumption that the inter-arrival times of the packets/cells can be adequately described by a Poisson process is questionable in most cases, see [Paxson & Floyd 95]. Consequently, more sophisticated models had to be developed.

One approach to describe the correlated, bursty behavior of packet traffic is the use of ON/OFF models, see also Sect. 2.1: packets or cells are generated at high rate during ON periods, while nothing (or very little) is generated during the OFF periods. If both, ON and OFF times are exponentially distributed, and the arrivals during the ON times can be described by a Poisson process, such models have representations as Markov Modulated Poisson Processes (MMPPs), see Sect. 2.2. [Heffes 80] is one of the earliest papers that discusses such models with application to telecommunication traffic. More general MMPPs with several levels but still exponential state-holding times are discussed in [Pancha et al. 97] and [Heyman 00b]. Other approaches to describe the bursty nature of packet traffic use the notion of bulk arrivals, see [Chaudhry & Templeton 83] and [Lucantoni 93]. Furthermore, certain models were designed with the goal of capturing the empiric correlation structure of measured traffic data for small lags (i.e. the short range correlation). The TES model (Transform and Expand Sample) in [Melamed 91] belongs to that group, however it is almost exclusively used in simulation models.

A series of measurements of network traffic – one of the earliest discusses Ethernet traffic in [Leland et al. 94] – has indicated that burstiness by itself as it is described by MMPP models might not be sufficient. These measurements have revealed self-similar or Long-Range Dependent properties: they showed burstiness over a wide range of time-scales. Such self-similar behavior is illustrated in the left column of Fig. 1.1, see Chapter 3 for more details. The mathematical theory of self-similar processes is described in-depth by [Samorodnitsky & Taqqu 94].

Those early measurements of network traffic initiated a controversial and still ongoing discussion about the impact of such LRD properties on the performance of telecommunication networks. The detailed analysis in this thesis provides an answer to this discussion for the family of N-Burst models that is developed in this thesis. The impatient reader has to remain impatient at this time, because the answer is: it depends! The impact of the LRD properties on performance can range from practically zero to devastating. Chapters 5 and 6 present the details.

There is evidence (see [Willinger et al. 95] and [Crovella & Bestavros 96]) that such LRD properties are caused by high variance distributions of the size of the individual transfers. Inspired by that fact, the use of so-called sub-exponential distributions has become popular in queueing models of various types, e.g. [Heyman 00] and [Dumas & Simonian 00]. The necessity to include such distributions in traffic models is the second major difference of modern

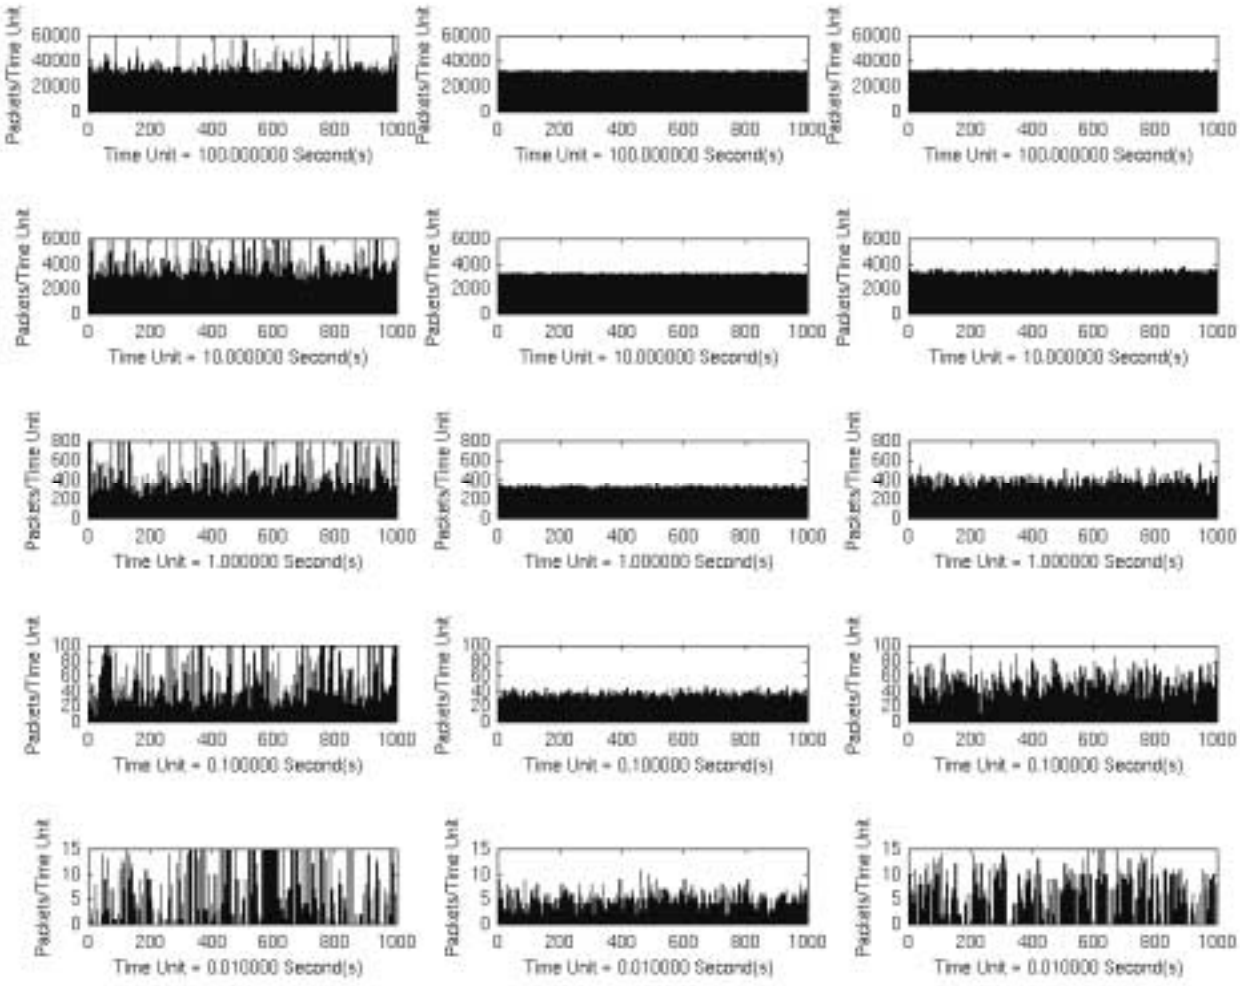

Figure 1.1: **Illustration of Self-Similar Properties in the Counting Process:** Even though the interval size is increased by a factor of 10^4 from bottom to top, the number of cells per interval still shows large fluctuations for the LRD 1-Burst process in the left column. The Poisson process in the middle and the ON/OFF model with exponential state-holding times in the right column become smooth more quickly. The graphs are created in simulation experiments, see [Lipsky et al. 99b]. A plot of the counting process for Ethernet measurements is shown in [Leland et al. 94]: the fluctuations disappeared only gradually as in the left column of this figure.

telecommunication models as compared to classical telephony models. The latter were only concerned about voice-traffic which shows rather well-behaved call-holding times. As a reminder, the first major difference is the packet-switched approach in data networks as opposed to the circuit-switched approach of classical telephony.

Although there are a large number of models for LRD traffic in the literature, the main examples are so-called Fractional Gaussian Noise (FGN) models, see e.g. [Norros 95] and [Krishnan 96], the so-called M/G/∞ model in e.g. [Guerin et al., 00], and fluid-flow ON/OFF models with sub-exponential ON periods in e.g. [Dumas & Simonian 00], it is very difficult to obtain a complete set of performance results for analytic queueing models with such traffic. Frequently, the investigations are only based on simulation experiments, but as the analysis in Chapter 6 shows, simulation results for such LRD traffic can be subject to large fluctuations and therefore have to be interpreted very carefully. For analytic models, mostly only asymptotic bounds or asymptotic approximations (e.g. from Large Deviation theory) have been obtained for a small set of performance parameters.

This thesis provides a very detailed understanding of the performance behavior of queues for a very general class of arrival processes, called N-Burst. The discussion includes the detailed setup of the model, the introduction of LRD properties in the model, methods for parameter estimation from real data, and finally the detailed analysis of large set of performance parameters in the analytic model, both by steady-state analysis and transient analysis. All the individual parts are necessary to make the model applicable to real-world engineering problems.

The major benefit of this thesis is that it provides a very thorough understanding of the performance behavior of such N-Burst traffic. In addition, several methods of analysis and some of the most important results in this thesis were previously unknown. Those are:

- Systematic use of *truncated Power-Tails* in MMPP models:
  This thesis uses a certain family of distributions that was developed by [Greiner et al. 99] to include LRD properties in classical MMPP models, see Chapter 3. This general approach is also taken by [Andersen & Nielsen 98] and [Robert & Boudec 97], but with different methods: [Andersen & Nielsen 98] use the superposition of many two-state MMPPs and develops an algorithm to fit a certain correlation structure of the packet counts with this aggregated model. [Robert & Boudec 97] use an MMPP ON/OFF model with a rudimentary form of the distributions in [Greiner et al. 99], but for the duration of the OFF times. Section 5.7 shows that this has only minimal impact on the performance behavior of queueing models with that arrival process.

  The family of distributions in [Greiner et al. 99] was used in another Semi-Markov model in [Lipsky et al. 97], [Fiorini 97] and [Schwefel 97]: the superposition of two renewal processes. Such a model shows also LRD properties but did not match other important properties of network traffic. Also, rudimentary forms of the 1-Burst model are analyzed to a very limited extent in that earlier work, which can be seen as predecessor of this thesis.

  In order to investigate the impact of the truncation of Power-Tails, which are in our case embedded in more complicated arrival processes, the definition of a *Power-Tail Range* is necessary, see Sect. 3.5. Only with that definition can conclusive results be obtained that are independent of the specific algorithmic representation of the truncated Power-Tail distributions.

- Several results of the performance analysis of N-Burst/M/1 models were previously unknown but turn out to have major impact: In particular the existence of so-called *blow-up points* and the understanding of why those occur in such models is essential. Furthermore, the *impact of the truncation* of Power-Tail distribution can be quantified with the help of the definition of a Power-Tail Range. Finally, the necessity to perform *transient analysis* leads to the definition of new performance parameters (the transient overflow probabilities and the so-called conditional overflow ratio). Those transient parameters also show surprising behavior, which characterizes the kind of fluctuations that can be observed in performance models with such traffic.

Finally, the in-depth analysis settles the question of when and how LRD properties matter, however only for the model class that is discussed here. But it is a very general and – as we believe and measurements show – in certain scenarios a very realistic model class.

1.3 Definitions & Notation

This thesis deals solely with packet- or cell-based data networks (as opposed to circuit switched networks). Although the terminology is chosen in correspondence to ATM networks (cells, switches, etc.), the results can be transfered to any packet based transmission protocol, including IP. However, at first we neglect the fact that packets can have varying size, and that the size of subsequent packets can be correlated.

Hence, we assume that the incoming cell stream is completely described by the inter-arrival process (X_i). The time X_i between arrival of cells i and $i+1$ is also called *inter-cell time*. If the X_i are independent and identically distributed, then the inter-arrival process is a renewal process. However, actual measurements of real network traffic have shown that the inter-cell times are highly correlated, see [Leland et al. 94, Schwefel et al. 97, Gogl 00, Paxson & Floyd 95].

The inter-arrival times X_i are positive random variables. Therefore, the *Cumulative Probability Distribution Function* (CDF) is one-sided, and defined by:[1]

$$F(x) := \mathbb{P}\{X \le x\}$$

with $F(0_-) = 0$, and $\lim_{x\to\infty} F(x) = 1$. The *complementary distribution function*, also called *Reliability Function*, is defined as

$$R(x) := \mathbb{P}\{X > x\} = 1 - F(x).$$

The *probability density function* (pdf) is given by (if it exists):

$$f(x) := \frac{d\,F(x)}{dx} = -\frac{d\,R(x)}{dx}.$$

The ℓ^{th} *moment* of the distribution (or the *expectation of* X^ℓ), is defined by:

$$\mathbb{E}\left\{X^\ell\right\} := \int_0^\infty x^\ell \cdot f(x)\,dx. \tag{1.1}$$

The first two moments define the *variance*:

$$\mathrm{Var}(X) := \sigma^2 := \mathbb{E}\left\{[X - \mathbb{E}\{X\}]^2\right\} = \mathbb{E}\left\{X^2\right\} - [\mathbb{E}\{X\}]^2,$$

[1] '$\mathbb{P}\{U\}$' means 'Probability that U is true'.

and the dimensionless quantity, the *coefficient of variation*, is the ratio of the variance and the square of the mean (first moment):

$$C^2 := \frac{\sigma^2}{[\,\mathbb{E}\,\{X\}]^2} = \frac{\mathbb{E}\left\{X^2\right\}}{[\,\mathbb{E}\,\{X\}]^2} - 1\,. \tag{1.2}$$

Later on, the so-called *tail* of a distribution, that is the behavior of $R(x)$ for large x, will be of special interest. In order to characterize the behavior of a function $R(x)$ for large x, the following notation is useful:

$$f(x) \sim g(x) \quad :\Leftrightarrow \quad \lim_{x\to\infty} \frac{f(x)}{g(x)} = c, \quad \text{where} \quad 0 < c < \infty.$$

The slightly more relaxed definition,

$$f(x) \overset{\bullet}{\sim} g(x) \quad :\Leftrightarrow \quad \liminf_{x\to\infty} \frac{f(x)}{g(x)} = c_1 \quad \text{and} \quad \limsup_{x\to\infty} \frac{f(x)}{g(x)} = c_2, \quad \text{where } 0 < c_1 \le c_2 < \infty,$$

is necessary later on in Chapter 3. If the two limits are close, $c_1 \approx c_2$, $f(x) \overset{\bullet}{\sim} g(x)$ is practically (but not mathematically!) equivalent to $f(x) \sim g(x)$.

One of the peculiar properties of network-traffic is that the inter-cell times are highly correlated. Necessary for the definition of correlation is the *Covariance* of two random variables X and Y:

$$\mathrm{Cov}(X,Y) := \mathbb{E}\{\,[X - \mathbb{E}\,\{X\}]\,[Y - \mathbb{E}\,\{Y\}]\,\} = \mathbb{E}\,\{XY\} - \mathbb{E}\,\{X\}\;\mathbb{E}\,\{Y\}\,.$$

The *Correlation Coefficient* of X and Y is then defined as

$$r(X,Y) := \frac{\mathrm{Cov}(X,Y)}{\sigma_X\,\sigma_Y}\,. \tag{1.3}$$

From its definition, it can be shown that $-1 \le r(X,Y) \le 1$. If $r(X,Y) = 0$ then X and Y are called *uncorrelated*. If X and Y are independent, then it follows that they are uncorrelated. On the other hand, even if two variables are uncorrelated, they still may not be independent.

The *autocorrelation coefficient lag-k* for an inter-arrival process (X_i) is defined as

$$r_i(k) := r(X_i, X_{i+k}).$$

For so-called *covariance stationary* processes, the value of $r_i(k)$ is the same for all indices i, so we just use $r(k)$. A positive $r(k)$ means that if the value of one inter-cell time, X_i, is bigger than the mean inter-cell time, X_{i+k} k steps later is more likely to be bigger than the mean as well.

Finally, the fundamental distribution for continuous time Markov processes is the *exponential* distribution (with rate μ). It has as its reliability function and density function:

$$R(x) = e^{-\mu\,x}, \quad f(x) = \mu\,e^{-\mu\,x}. \tag{1.4}$$

Let X be the exponentially distributed sojourn-time for a particular state with leaving rate μ of a Markov Chain. Then it follows:

$$\mathbb{E}\,\{X\} = \frac{1}{\mu}, \quad \mathrm{Var}(X) = \frac{1}{\mu^2} \quad \Rightarrow \quad C^2 = 1.$$

One significant property of the exponential distribution is that it memoryless. It does not matter at what time in the past the last event happened, the residual time is still exponentially distributed with the same rate:

$$\mathbb{P}\{X > x+h | X > x\} = \frac{\mathbb{P}\{X > x+h\}}{\mathbb{P}\{X > x\}} = \frac{e^{-\mu(x+h)}}{e^{-\mu x}} = e^{-\mu h} = \mathbb{P}\{X > h\} \quad \text{for} \quad h \geq 0. \tag{1.5}$$

If there are two independent, exponentially distributed variables X_1 and X_2 with rates μ_1 and μ_2, then the probability that the first one is smaller is

$$\mathbb{P}\{X_1 < X_2\} = \int_0^\infty R_{X_2}(x)\, f_{X_1}(x)\, dx = \int_0^\infty F_{X_1}(x)\, f_{X_2}(x)\, dx = \frac{\mu_1}{\mu_1 + \mu_2}\,. \tag{1.6}$$

The Counting Process: A less detailed, but in practice frequently used description of the traffic is the *Counting Process* ($N_i(\Delta)$), for a given time-interval $\Delta > 0$. $N_i(\Delta)$ is a random variable whose value corresponds to the number of cell arrivals in the time-interval $[(i-1)\Delta,\, i\Delta[$. Note that even if the inter-cell times X_i stem from a renewal process, the counting process can be correlated.

Queueing Models: Queues will be used in this work for two purposes:

1. As a model for the bottleneck switch in the network: Cells arrive at the switch as described by the inter-cell time process (X_i). The switch needs some time to process the information in the cell header and to transmit the cell over the outgoing link. This is called the service-time. We almost always assume exponentially distributed service times with rate ν in this work. A discussion about the meaningfulness and the impact of that exponential distribution is given in Chapter 2. If cells arrive more quickly than they can be served, they will queue up in a buffer. Both cases, finite and infinite buffers are discussed herein.
2. As a description of the number of active bursts: closed queueing systems with two queues and finite population are used for the modulating process that determines the cell-arrival rate at the switch, see Chapter 2 for details. More general service-time distributions and also multiple servers come into play for that purpose.

A general introduction in queueing theory is given in [Kleinrock 75]. Closed queueing systems with finite population are treated in [Lipsky 92], whose algebraic approach is the basis for the evaluation techniques used in this thesis, see Appendix B.

The performance behavior of a queueing model is determined in large part by the queue-length process, Q_t. In case of the queue that models the network component, the queue length is also referred to as *buffer-occupancy*. In our terminology, the queue length always includes the customers (here cells or bursts) that are currently in service. Frequently, the queue length at the arrival-time of cell j, $Q_j^{(a)}$ is of interest:

$$Q_j^{(a)} := Q_{T_j}, \quad \text{where} \quad T_j = \sum_{i=0}^{j-1} X_i\,.$$

The transient analysis in Chapter 6 is in major parts concerned with the First Passage Time to buffer-level n, which is the following random variable:

$$\tau(n) := \min_{t \geq 0} \{t \,|\, Q_t \geq n\}\,.$$

Furthermore, the concepts of *Busy Periods* and *Busy Cycles* are important for the transient analysis:

- **Busy Period**: A Busy Period of a queueing model starts when a customer enters an empty queue. It ends when the queue is empty (all servers are idle) for the first time again.
- **Busy Cycle**: The definition of a Busy Cycle is limited to queueing models with an underlying Markovian structure of arrival and service processes (or more general: the existence of a regenerative period in the model). The Busy Cycle starts when the queue is empty *and* the arrival process is in a specific state. It ends when that particular combination, empty queue and specific state of arrival process, is reached again for the first time. That definition proves useful for theoretical derivations, since regenerative arguments can be used, i.e. each Busy Cycle is stochastically identical.

  For an M/G/1 queue, the Busy Cycle contains the Busy Period and the preceding idle period, while in the case of a GI/M/1 queue, a Busy Cycle is a Busy Period and its succeeding idle period.

In Chapter 9 the inter-cell times X_i are made dependent on the buffer-occupancy (and other parameters of the server) in order to model the feedback introduced by flow-control mechanisms such as in TCP. However, in the main part of the thesis, the incoming traffic is assumed to be independent of the state of the queue.

Abbreviations

ATM	Asynchronous Transfer Mode	
BOP	Buffer-Overflow Probability (steady-state) (infinite backup buffer, N-Burst/M/1 queueing model)	Sects. 2.6 & 5.2
BOR	Buffer-Overflow Ratio (transient) (infinite backup buffer, x/M/1 queueing model)	Sect. 6.4
B-WiN	'Breitband Wissenschaftsnetz'	Sect. 4.1
CD	Cell Delay	Sect. 5.2
CLP	Cell Loss Probability (steady-state) (finite-buffer N-Burst/M/1/B queueing model)	Sects. 2.6 & 5.5
CLR	Cell Loss Ratio (transient) (finite-buffer N-Burst/M/1/B queueing model)	Sect. 6.4
FGN	Fractional Gaussian Noise	Sect. 2.7.5
FPT	First Passage Time (transient)	Sect. 6.2
iid	independent, identically distributed	
IP	Internet Protocol	
IS model	Independent Source model	Sect. 2.4
LAQT	Linear Algebraic Queueing Theory	App. B & [Lipsky 92]
LRD	Long-Range Dependence	Sect. 3.1
MBS	Maximum Burst Size	Sect. 3.5
mCD	mean Cell Delay (steady-state or transient)	Sects. 5.1 & 6.1
ME	Matrix Exponential	App. B.1 & [Lipsky 92]
mFPT	mean First Passage Time (transient)	Sect. 6.2.2
MMPP	Markov Modulated Poisson Process	Sect. 2.2
PRA	Peak-Rate Allocation	Sect. 5.1
PT	Power-Tail: $R(x) \to c/x^{\alpha}$	Sect. 3.2

QoS	Quality of Service	
SM	Semi Markov	App. C.1
SST	Self-Similar Traffic	Sect. 3.1
	(here often: asymptotic second order self-similarity)	
TCP	Transmission Control Protocol	Chapt. 9
TPT	Truncated Power-Tail	Sect. 3.4

Notation

Bold-faced lower-case letters (e.g. $\boldsymbol{\pi}$, $\mathbf{p}$) represent row vectors while bold-face capital letters represent matrices. $\mathbf{u}'$ expresses the transposed of the row-vector $\mathbf{u}$, analogously for matrices. The unit-matrices are denoted as $\mathbf{I}$ or $\boldsymbol{\mathcal{I}}$ (depending on the dimension). $\mathbf{e_i}$ represents the vector with all components equal to zero except for its ith component which is set to 1. $\boldsymbol{\epsilon}'$ and $\boldsymbol{\varepsilon}'$ are column vectors (of different dimension) with all components set to 1.

N-Burst Parameters		Sect. 2.4
N	Maximum number of simultaneous bursts	
$\overline{n_p}$	Mean number of cells per burst	
λ	Average cell rate of all sources together	
κ	Average cell rate per single source: $\kappa = \lambda/N$	
λ_p	Peak-rate within a single burst	
b	Burstiness: $b = 1 - \kappa/\lambda_p$	
α	Tail exponent of PT burst-length distribution	
$c_{PT}^{(1)}(\alpha)$	Tail-constant of normalized PT distr.	
MBS	Maximum Burst Size	

Queue Parameters		Sect. 2.6
ν	Service rate	
ρ	Utilization: $\rho = \lambda/\nu$	
B	Buffer Size	

Performance Analysis		
i_0	Number of blow-up region, $i_0 \in \{1, ..., N\}$	Eq. (5.2)
i_Δ	location within blow-up region	Eq. (5.3)
β	Tail-exponent of over-saturation periods	Eq. (5.4)
$\mathrm{Rng}\,(R(x))$	Power-Tail Range of reliability function	Sect. 3.5
$\overline{q}_{i_0}$	PT Range of queue-length distribution	Eq. (5.14)
$\overline{q}_N$	largest PT Range of queue-length distribution	Sect. 5.3
c_{xxx}	Tail-constant for Power-Law behavior of 'xxx'	Eq. (5.19)
$\tau(n)$	First Passage Time to level n	Sect. 6.2
$\gamma(t_0, B)$	Probability of overflow in time t_0	Eq. (6.4)
$\mathrm{BOR}_c(t_0, B)$	conditional Buffer Overflow Ratio	Sect. 6.4
$\mathrm{CLR}_c(t_0, B)$	conditional Cell Loss Ratio	Sect. 6.4

Chapter 2

Traffic and Performance Models

In general, network traffic is not smooth. It is very bursty and the inter-cell times are highly correlated. Even worse, the correlation holds up over a large number of cells. That effect is called *Long Range Dependence* (LRD) and can cause devastating performance effects in some scenarios.

Yet, the impact of LRD is still not completely clear. A reasonable model has to have the ability of reproducing that effect, if only to decide on the circumstances when its devastating performance effects show up.

Moreover current and future heterogeneous networks have to serve as media for a wide range of applications. Voice-, Video- and Data-traffic have very different properties. A traffic-model thus must be very flexible and capable of capturing those application mixes.

For practical application, traffic models should have as few parameters as possible, also known as *parsimonious* models. Otherwise the problem of parameter estimation becomes very hard. So it is allowed and even necessary to simplify models by leaving out properties that do not appreciably affect performance. Furthermore, the more simplified the models are, the higher are the chances that the analytic models remain tractable. Of course, the hard part is to determine which properties of real traffic need to be captured and which are superfluous due to their minor impact on performance.

For reasons of extrapolation and forecasting, it is desirable that model parameters are operational, i.e. they have a measurable physical equivalence. Only then, scenarios such as an increase in traffic due to, e.g., a larger number of users can be considered appropriately in the model.

This chapter introduces a very flexible class of models, called N-Burst. They are based on a certain subset of ON/OFF models that are fitted into the tractable framework of MMPPs by the use of Matrix-Exponential (often Phase-Type) distributions. The reader is advised to turn to the appendix in order to obtain the necessary background about *Matrix-Exponential* distributions (Appendix B.1) and LAQT representations of MMPPs (Appendix C.2). Also, closed queueing systems with a finite number of customers as discussed in Appendix D.1 play a major role in the setup of the N-Burst models.

2.1 ON/OFF Models

The use of so-called ON/OFF models has become increasingly popular, since they reproduce the bursty properties of network traffic. In these models, traffic is generated only during ON periods (here also called *active periods* or *bursts*). The duration of the ON periods is described by a random process (A_i). Each ON period is followed by an OFF period during which the source is idle, so no data is transmitted. Let the duration of the OFF periods be described by the stochastic process (B_i).

In a cell/packet based model, the traffic during the ON periods is best described by a point-process model. The most important examples are:

- Poisson process: The inter-cell times during ON periods are iid, and their marginal distribution is exponential.
- Deterministic inter-cell times: While the source is ON, cells are generated with a fixed inter-cell time.

An alternative to point-process models are the so-called fluid-flow models, in which the traffic is not measured in discrete units (cells) but instead by a real number, which is increased continuously at some rate during ON periods. Fluid-flow models can often be described as a limit process in which traffic is measured in infinitely small units, see Sect. 2.7.3. Fluid-flow models are frequently used for reasons of better tractability.

A second coarse classification of an ON/OFF model can be derived from the ON/OFF process itself, which can either be a discrete-time or a continuous-time process. In the former, a certain time-step Δ is chosen and transitions from ON to OFF periods and vice versa can only occur at time $j\Delta$, $j = 1, 2, ...$

Discrete time models with deterministic inter-cell times during bursts are frequently used for ATM models. In contrast to that, we focus on continuous time ON/OFF models with Poisson cell-arrivals within ON periods. However, we will see in Chapter 5 that the major cause of occurring congestion in our performance models is not the Poisson variability during ON times (which is arguable in ATM models), but instead the burstiness caused by a special ON/OFF structure.

Extensions of the ON/OFF approach: A plain ON/OFF structure as introduced above is the simplest possible case. Two, not necessarily different ways for extensions of that simple ON/OFF approach are listed in the following:

- Aggregation: Plain ON/OFF sources are often used to model the behavior of a single user. A superposition of several ON/OFF sources is necessary to describe the collective behavior of several users together. The superposition can be done for independent and identical ON/OFF models, or arbitrary complex scenarios can be created. Aggregated fluid-flow ON/OFF sources are treated for instance in [Dumas & Simonian 00]. A discrete model with deterministic inter-cell times during ON times is discussed in [Fan & Georganas 96].
- Multiple Levels: The aggregation of cells from N identical ON/OFF models results in a process in which cells are generated at $N + 1$ distinguished rates, also called levels. In the simplest case, those rates are $i\lambda_p$, $i = 0, 1, ..., N$, where λ_p is the rate at which cells are generated during the ON period of a single source. In general, a process is obtained that generates cells at discrete but possibly infinitely many rates. The simple two-level ON/OFF concept is thus extended to a larger number of levels. A common example is the so-called M/G/∞ model, see Sect. 2.7.5 and [Guerin et al., 00], in which the number of busy servers of an M/G/∞ queue represents the level, hence controls the rate at which cells (or fluid) are generated.

Both concepts, the superposition of multiple ON/OFF sources and also more general multi-level processes, are part of the N-Burst model that is developed in this chapter.

2.2 Markov Modulated Poisson Processes

As already mentioned in the previous section, our focus here is on continuous time ON/OFF models with Poisson cell transmissions during ON. Under certain restrictions – in the simplest case, ON and OFF periods are independent and their duration is exponentially distributed – such ON/OFF models can be expressed as *Markov Modulated Poisson Processes* (MMPPs).

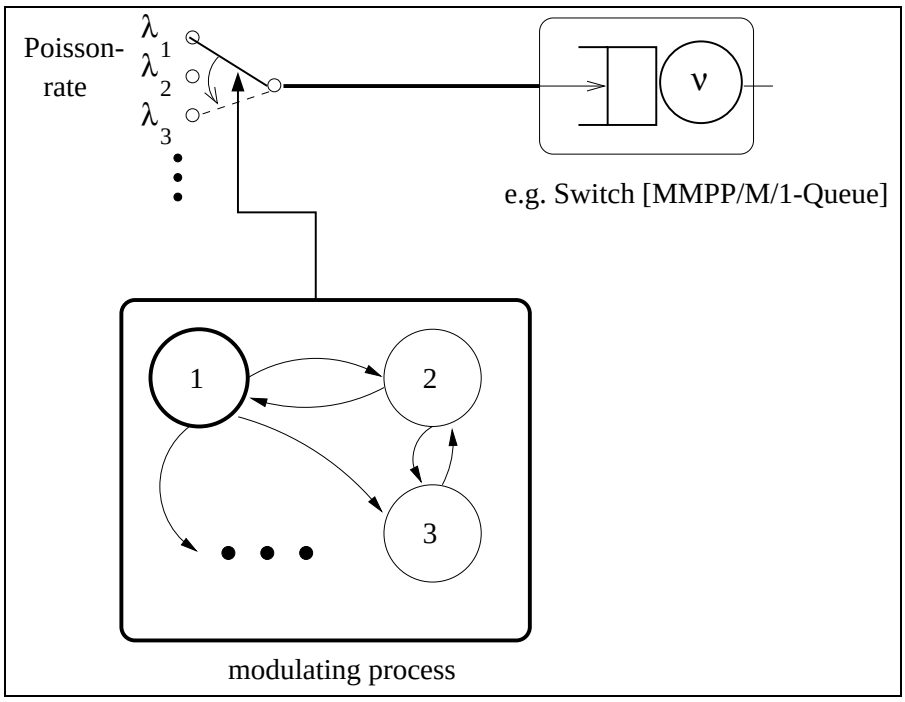

Figure 2.1: **An MMPP as Arrival Process to a Queueing Model:** A Markov Modulated Poisson Process (MMPP) consists of a Markov chain (lower box) which describes the modulating process. A token moves around in the box but never leaves it. The actual arrivals (of e.g. cells to a switch) are generated by a Poisson process, whose rate is determined by the position of the token within the modulating process.

In an MMPP model, cells are always generated by a Poisson process, however its rate λ_i depends on the state i of a Markov chain (called the modulating process). Figure 2.1 illustrates the MMPPs in the context of a queueing model with an MMPP arrival process.

The advantage of the MMPP representation is that queueing models with such arrivals are tractable and matrix-analytic methods can be used for their evaluation. See Appendix C.2 for a description of the matrix representation that is used herein. Matrix analytic methods for MMPP/M/1 queues are described in [Meier-Hellstern & Fischer 92] and in Appendix D.

The simplest case is a two state MMPP which is an ON/OFF model with exponential ON and OFF duration. One can also allow for some positive Poisson cell rate during OFF. The Poisson traffic during OFF can be viewed as background Poisson traffic. Such a model was discussed with respect to network traffic modeling 20 years ago in [Heffes 80]. This simple ON/OFF model is in fact a special case of the 1-Burst model which is introduced in the next Section. Recently, [Andersen & Nielsen 98] used the superposition of several different two-state MMPPs to achieve Long-Range Dependent (LRD) properties. Traffic with such properties can also be generated by the N-Burst models. However the approach is different, since it is achieved here by more complicated ON time distributions, see Chapter 3.

MMPPs belong to the larger class of Markovian Arrival Processes (MAPs), also called Semi-

Markov Processes, see Appendix C.1. If one allows for bulk-arrivals the so-called B-MAP processes are obtained, as discussed in detail in [Lucantoni 93].

The attractive feature of MMPP models for network traffic models is that they allow the variability that is observed in the cell rates to be captured. However, one criticism is that the restriction to exponential state-holding times in the modulating process is not acceptable. We avoid that restriction in our N-Burst models, however for the price of a more complicated but still Markovian modulating process.

Empiric MMPP models are frequently used for traffic modeling: the number of states and the Poisson cell rates λ_i are obtained by some heuristic approach (see e.g. [Heyman 00b]) and the transition rates/probabilities are determined from a set of measurements, see e.g. [Pancha et al. 97]. However, [Rimkus 99] shows that some of the properties that are connected with the LRD properties and were actually observed in actual measurements are not reflected by such empiric MMPP models. In particular, a slow convergence to steady-state is observed in [Rimkus 99] for trace-driven simulations of queueing models with the measured data as arrival process, while the performance parameters in simulations of a fitted empiric MMPP model reached steady-state rather quickly.

Furthermore, such empiric MMPP models have no underlying structure in the modulating process. Extrapolation problems such as how to model an increase in traffic volume are difficult to tackle in such models. Plain scaling of the cell-rates and the transition rates would be a viable but very doubtful approach. The underlying physical structure of e.g. ON/OFF models allows for different physical scenarios of increasing the traffic volume: For instance, the increase of the ON duration or the decrease of the OFF duration are meaningful approaches if single ON/OFF sources are used to model aggregated traffic, as done in Sect. 4.3.

2.3 1-Burst Model

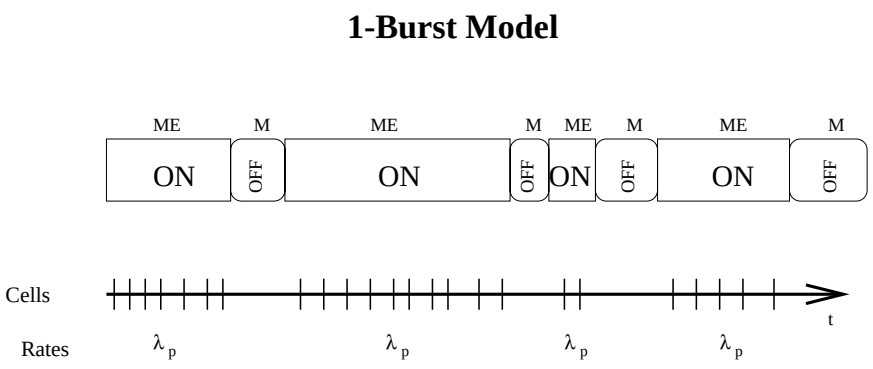

Figure 2.2: **The 1-Burst Model:** traffic is generated with Poisson rate (peak-rate) λ_p only during ON periods, whose duration is described by a general Matrix-Exponential distribution. The duration of the OFF periods is assumed to be exponentially distributed.

As base models we use single source ON/OFF models, called 1-Burst, as illustrated in Figure 2.2. During ON periods, also called bursts, cells are generated with Poisson rate λ_p. The distribution of the duration of the ON periods is not necessarily exponential. It can be any Matrix-Exponential distribution $\langle \mathbf{p}, \mathbf{B} \rangle$, see Appendix B. It is shown in Chapter 3, how LRD properties can be included in the model by a particular choice of the ON time distribution, which is confirmed in many measurements of real traffic. Furthermore, it is shown in Chapters 5 and 6 that the choice of the actual ON time distribution has a great impact on the performance

behavior of queueing models. The OFF time distribution on the other hand is less critical, see Sect. 5.7. We assume here that the OFF time distribution is always exponential with mean Z, although it could easily be replaced by another ME distribution.

The use of ME distributions for the ON times in the 1-Burst model allows for more general ON/OFF models than in [Heffes 80], yet the 1-Burst model still has an MMPP representation, hence the queueing models remain tractable.

In addition to the cell-rate (peak-rate) λ_p during bursts, and the mean OFF time Z, we use the following parameters:

$$\begin{aligned} \text{Mean ON time:} \quad & \overline{x_p} := \mathbf{p}\mathbf{B}^{-1}\boldsymbol{\varepsilon}' \\ \text{Mean number of cells per burst:} \quad & \overline{n_p} = \lambda_p \overline{x_p} \\ \text{Average overall cell-rate:} \quad & \kappa = \frac{\overline{n_p}}{Z + \overline{x_p}} \end{aligned}$$

κ is the average cell-rate of the ON and OFF times together. The ratio κ/λ_p can be used as a measure of the burstiness of the source: Define the *burstiness parameter*

$$b := 1 - \frac{\kappa}{\lambda_p} \cdot \tag{2.1}$$

Using the equation for κ, the burstiness parameter b can also be computed by:

$$b = \frac{Z}{Z + \overline{x_p}},$$

i.e. it is the fraction of time the source is OFF. It is easy to see, that for $b = 0$ and fixed κ, there is no ON/OFF structure any more ($Z = 0$), and the process reduces to pure Poisson. At the other end when $b \to 1$, $\lambda_p \to \infty$, so all cells in a burst are transmitted simultaneously. See Section 2.7 for a more detailed discussion of these limits. For now, it is enough to mention that the higher b, the more bursty the source appears, so that the name *burstiness parameter* is justified.

The ON/OFF structure of the modulating process of the 1-Burst model can also be described by an M/ME/1//1 queueing system with finite (here one customer) population. See [Lipsky 92] or Appendix D for a discussion of such queueing systems with finite population. The busy period of the ME server in the M/ME/1//1 queueing system corresponds to an ON period. It would appear that the description of the modulating process as a finite queueing system with one customer is a rather complicated, theoretical viewpoint of the 1-Burst model. However, its usefulness for generalizations will become obvious when more complicated, aggregated models are introduced.

In addition to the ON/OFF cell traffic, the 1-Burst model also allows for smooth background Poisson traffic with rate λ_0. In the MMPP representation, that background traffic just raises the Poisson rates to λ_0 during OFF periods and to $\lambda_p + \lambda_0$ during ON periods. The overall cell-rate is then

$$\lambda := \lambda_0 + \kappa\,.$$

Physical Equivalences: Three different real scenarios are natural matches for a single source ON/OFF model:

1. Single User or single end-to-end connection.

2. Multiple Sources (Users) but exclusive access: Whenever one user is transferring data, no other bursts can become active. However, those blocked requests are just dropped (if they wait, the result is an [IS/G/1/K] model which is introduced in the next section). If all the sources have exponential OFF periods, the overall observable OFF period is again exponentially distributed, but with smaller mean.

3. Aggregated Traffic on a single link: Instead of assigning a burst to a single source at the network endpoint, it is also a viable approach to model the traffic on a link within the network by an ON/OFF scheme regardless of the origin of the individual cells. Since such a link may be the output of another intermediate system (switch/router), the cell stream on that link can very well show ON/OFF behavior, where the ON periods are the Busy Periods of the switch, see also Sect. 10.4. In the measurements that are discussed in Chapter 4, such single-source ON/OFF behavior is found, although the traffic is an aggregation of cell-streams from a large number (several hundreds) of physical sources.

Thus, a burst can be initiated by a single source-destination pair (which is the viewpoint in Chapter 9), or its definition can be related only to the observed traffic on a single link, i.e. independent of the origin and the destination of incoming cells, a burst just represents a number of subsequent cells that can be adequately described by a Poisson process with rate λ_p. The latter viewpoint is taken in most parts of this thesis except for Chapter 9.

The background Poisson rate λ_0 can be used to describe traffic outside the ON/OFF pattern, e.g. network management information that tends to show a smoother (sometimes periodic) behavior. However, depending on the specific management task, such traffic can also be very bursty with larger chunks of data transmission, e.g. for updates of routing tables. In the latter case, that traffic should not be modeled via the Poisson rate λ_0, but rather by an additional ON/OFF source.

2.4 Aggregated Traffic: N-Burst Model

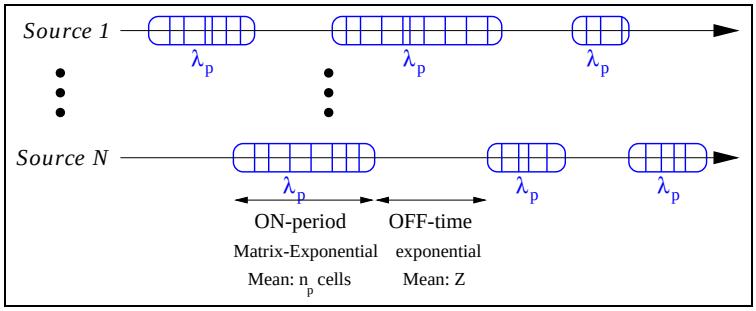

Figure 2.3: **The N-Burst Independent Source [IS] Model:** Cells from N ON/OFF sources are multiplexed together.

A natural way of aggregating the traffic of N ON/OFF sources is shown in Fig. 2.3: Here the sources are assumed to be independent and identical. The cells of up to N simultaneous bursts are simply multiplexed together. The parameters for the description of the individual source remain as in the previous section, with the only difference that the average cell rate λ is now

$$\lambda = N\kappa + \lambda_0 \,.$$

this model is called *N-Burst Independent Source (IS) model* hereafter.

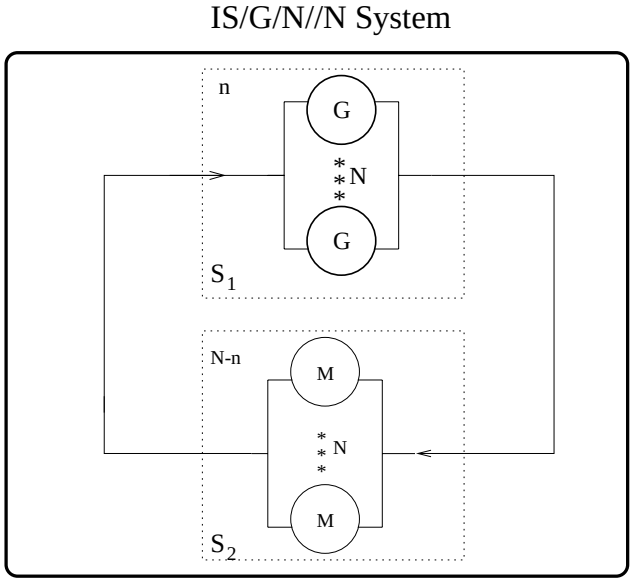

Figure 2.4: **The IS/G/$N//N$ model as modulating process of the N-Burst Independent Source Model:** The number of active servers in S_1 corresponds to the number of active bursts. Subsystem S_2 represents a delay-stage, that creates the exponentially distributed OFF times.

The aggregation of N sources in the N-Burst [IS] model results in a modulated Poisson process with $N+1$ levels, corresponding to the cell-rates $i\lambda_p + \lambda_0$, $i = 0, 1, ..., N$. The modulating process of the N-Burst [IS] model can again be expressed as a closed queueing system of type IS/ME/$N//N$ with N customers, shown in Fig. 2.4. The number of active ME servers in S_1 corresponds to the number of active bursts. The delay at the subsystem S_2 corresponds to the exponentially distributed idle time of the sources. The matrix representation of the resulting MMPP model is described in Appendix E.

Since the OFF periods are assumed to be exponentially distributed, and there is no need to distinguish between the N sources, the N-server delay stage in subsystem S_2 in Fig. 2.4 can be replaced by single-server queue with load-dependent exponential service-rate, as pictured in Fig. 2.5: if n sources are active, i.e. n customers are at S_1, then $N-n$ sources are idle, and the rate of a burst-start is $(N-n)/Z$. This is discussed further in the next paragraph.

Load dependent burst-arrivals:

The formalization of the modulating process of the N-Burst [IS] Model as a closed queueing system with load-dependent subsystem S_2 allows for another generalization: We can allow any load-dependent $M_{ld}/G/N//N$ queueing system as modulating process. We started off with the IS model, in which the burst-start rate $\mu(n)$ for n active sources is given as:

$$\mu(n) = \frac{N-n}{Z}, \quad n = 0, ..., N.$$

Another physically meaningful scenario is obtained when the burst-start rate is constant for $n < N$:

$$\mu(n) \equiv \mu \quad \text{for } n = 0, ..., N-1, \quad \text{and} \quad \mu(N) = 0.$$

The resulting modulating process is an M/G/N//N queueing system with no load-dependence in the arrival process. Thus, the burst-arrival process is a Poisson process with constant rate with the only restriction that the burst-arrivals are blocked when N bursts are active. In reference to that Poisson property, these models are called N-Burst [M] Model.

Physical Scenarios: Both, the [IS] model and the [M] model correspond to practically important physical scenarios:

- N-Burst [IS] Model: The aggregation of the cell-streams of N sources corresponds to the multiplexing of traffic from different links (each described by an ON/OFF source) at a switch entrance. The assumption here is that individual bursts do not react to the presence of others, i.e. no feedback or flow control mechanisms exists. The major parts of this thesis deal with such a base model. However, Section 2.5 provides the framework to include such a feedback, and Chapter 9 elaborates that approach and discusses a potential application to the TCP flow control mechanism in IP traffic.
- N-Burst [M] Model: Whenever a large number of independent sources generate bursts according to some stochastic process, the aggregated burst-start process becomes close to a Poisson process. A central limit theorem for stochastic processes (see [Heyman & Sobel 82]) provides the theoretical justification of that behavior. Consequently, the [M] Model corresponds to the scenario that a large number of independent sources can potentially become active. While being active they do not influence each other, except for the fact that new burst-starts are blocked when N sources are active. Note that in contrast to the [IS] model, N in the [M] model is not the number of sources, but rather the upper bound for the number of simultaneously active sources. The latter is normally a consequence of the finite network capacity.

More general scenarios, such as discouraged arrivals of bursts due to the reaction of users on congestion, can be discussed by selecting any function $\mu(n)$, $n = 0, ..., N-1$ for the load-dependent burst-arrival rate. Results for such models will not be discussed here, although the numerical evaluation is already set-up in full generality. The same statement holds for the following generalization:

Waiting Bursts:

For a final generalization of the process that describes the number of active bursts, we can allow whole bursts to queue up until they can access the line. This leads to an M_{ld}/G/N//K queueing system as modulating process, where any $K \geq N$ is meaningful. In case $K > N$, a queue at Subsystem S_1 is necessary, see Fig. 2.5. Such models will still be called N-Burst models, since at most N bursts can be active simultaneously. The difference now is that up to $K - N$ bursts can be waiting for their turn to transmit cells. For an exact specification of the model, we add the type of the modulating process in brackets, e.g. the IS models are referred to as N-Burst [IS/G/N//K] models. If $K > N$ is not specified explicitly, then the model without waiting bursts ($K = N$) is assumed.

See Appendix E for the matrix representation of the N-Burst models.

We briefly discuss the special case of the 1-Burst [M/G/1//K] Model, which is a variation of the 1-Burst [M] Model with up to $K - 1$ waiting bursts. In this model, the burst-arrival process is a Poisson process with constant rate, unless the upper limit of waiting sources is reached. Since

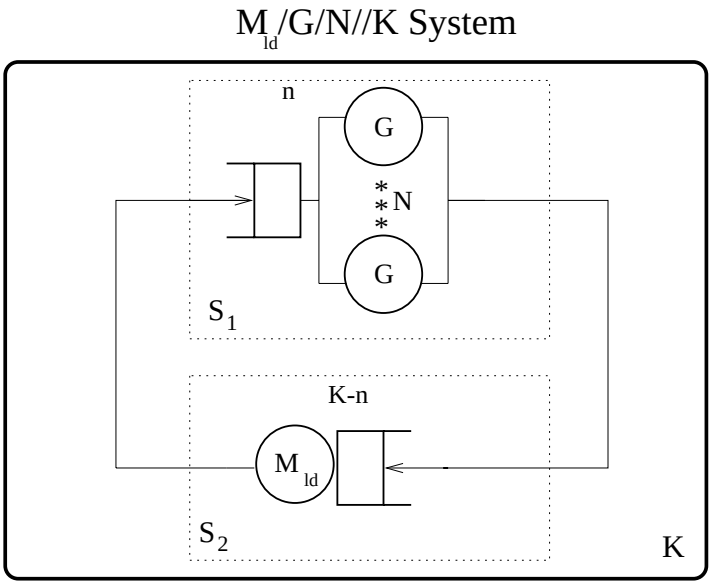

Figure 2.5: **The Modulating Process of a General N-Burst [M_{ld}/G/N//K] Model:** The load-dependent server in S_2 provides a general representation that includes the [M] and the [IS] model as special cases. For $K > N$, queueing in front of the N server unit in S_1 can occur.

$N = 1$, access to the network line is exclusive, i.e. only one source can transmit cells at a time. Sources that become active during the time that the line is used will be waiting until they get access. However, there is an upper limit: only up to $K - 1$ sources will be waiting. If that limit is reached, the burst-arrival process is blocked. As an alternative view to blocking, additional burst-starts can be allowed to happen but they are just discarded. That viewpoint is really an M/G/1/K loss system. However, due to the exponential distribution of the inter-arrival times, the M/G/1/K and the M/G/1//K model are equivalent here – the number of customers in S_1 is exactly the same stochastic process in both models.

Since one of the waiting bursts becomes active immediately when the current burst finishes, two or more ON periods can happen right after each other. The 1-Burst [M/G/1//K] is therefore still a plain ON/OFF model, but the 'new' ON time can be the convolution of several individual burst-durations. Nevertheless, the [M/G/1//K] model on the cell level could be also expressed as [M/G$'$/1//1] model, but with a much more complicated 'new' ON time distribution G$'$, where G$'$ is the Busy Period duration of the M/G/1//K system.

We have introduced a large, very flexible class of N-Burst models in this section. The representation of these models as MMPPs with closed M_{ld}/ME/N//K queueing systems as modulator enables us to apply Matrix-Analytic methods to all these models. However, the performance analysis here will mainly focus on the N-Burst [IS/G/N//N] model, i.e. the model with N independent sources and no waiting bursts. For the rest of the thesis, if we use the term 'N-Burst' without further specification of the modulating process, it refers to the [IS/G/N//N] model. When models with waiting bursts are considered, the parameter $K > N$ is pointed out explicitly.

2.5 Application Mixes & Dynamically Adjustable Cell-Rates

So far, we assumed all bursts in the N-Burst model to be described by an identical parameter set $(\langle \mathbf{p}, \mathbf{B} \rangle, \overline{n_p}, \lambda_p)$ and to be independent of each other: their cell-rates simply add up. In particular, the sources do not react to occurring congestion. For the performance analysis in the subsequent chapters, we concentrate mainly on such a 'base' model. However, the general formulation of the N-Burst model allows for an incorporation of more complex source behavior. Such modifications are described in this section together with their relevance for traffic modeling. Performance results for such more complex models are discussed briefly in Sect. 5.8 for application mixes and in Chapter 9 for more complex transmission protocols for so-called elastic traffic.

2.5.1 Model of Interaction between Sources

In many scenarios, e.g. real-time or streaming applications, and also the transmission of IP packets via UDP, the assumption of independently transmitting sources that are not affected by any occurring congestion holds. In contrast to such real-time applications, many data-transmissions are not necessarily time-critical – except for the impatience of the human users. Such traffic is usually called *elastic traffic*, and frequently there are so-called *flow-control mechanisms* implemented in the transmission protocols that allow the sending rate of the elastic traffic to adjust in correspondence to the congestion level of the network. The most important example is the TCP flow-control mechanism, which is discussed in more detail and modeled by variations of the N-Burst model in Chapter 9. This section describes the fundamental general modifications that are necessary to incorporate such feedback in the N-Burst models. In the end, we obtain a very general, but still physically meaningful MMPP model where the modulating process is a closed queueing system of type $\mathrm{M}_{ld}/\mathrm{G}_{ld}/N//K$. The load dependence of the general, Matrix-Exponential server can model a 'back-off behavior' of the sources as described below. Arbitrary complex notions of load dependence are possible, some of which are described in subsequent sections.

The essential property of sources of elastic traffic is that they try to share the available bandwidth on a transmission path between all the active sources. Thus the rate λ_p of cell transmission during an ON period has to depend on the number n of active bursts. That dependence is captured by load-dependence factors $\beta(n)$, $n = 1, ..., N$ for the N-Burst model (or equivalently $n = 1, ..., K$ for the $[\mathrm{M}_{ld}/\mathrm{G}/N//K]$ model). Therefore, the individual cell-rate λ_p is replaced by $[\beta(n) \cdot \lambda_p]$, with the scaling factors $\beta(n)$ usually decreasing with higher n in most physically relevant scenarios (it is not very likely that additional traffic speeds up the active transfers).

However, the amount of data in a burst (i.e. the distribution of the number of cells per burst, with mean $\overline{n_p} = \lambda_p \overline{x_p}$) should remain unchanged by such a scaling of the intra-burst cell rate. Therefore, a scaling of the cell-rate requires an adequate adjustment of the duration of the ON time. In case of Phase-type distributed ON periods, an appropriate scaling of the Phase transition rates (in $\mathbf{B}$) with the same factor $\beta(n)$ is possible, see Chapter 9 for more details and for a discussion of the theoretical questions that arise in that situation. In terms of the representation of the modulating process as a closed queueing system, such a scaling of the Poisson rates requires the use of load-dependent servers in S_1 in order to keep the number of cells per burst the same. Thus, the result is a $[\mathrm{M}_{ld}/\mathrm{G}_{ld}/N//K]$ model with load-dependent burst-arrival rates (in S_2) and load-dependent multiple servers in S_1 for the ON time duration.

2.5.2 Traffic Mixes

So far, we assumed identical sources with ON/OFF characteristic. However, the traffic on data-networks is usually generated by a large number of different applications (such as FTP, HTTP, Video, Audio). Traffic from different applications also has different properties: for instance, file sizes are known to be Power-Tail distributed (e.g. [GARG ET AL. 92]) thus FTP-data can be expected to show the same distribution type for the ON periods. On the other hand, the frame-sizes of a video transmission are closer to an Erlangian distribution, see e.g. [GOGL 00] for a detailed investigation. It should be possible to characterize such application mixes and include them in the traffic model. This section demonstrates how the N-Burst model can be extended in such a way. Section 5.8 then discusses the performance impact.

There are two fundamentally different possibilities of how traffic mixes can come up:

1. All N sources are identical, yet the individual source can produce different types of bursts. Let us assume that we want to model t different traffic classes corresponding to e.g. HTTP-requests, FTP-transfers, voice-traffic, etc. . However, the traffic classes are not assigned to individual sources, but every source generates the same mix of all the different classes. Each of the t classes is described by a separate burst-length distribution, $\langle \mathbf{p}^{(i)}, \mathbf{B}^{(i)} \rangle$, and its peculiar cell-rate during bursts, $\lambda_p^{(i)}$.

   The model in this case is the usual N-Burst [IS] model but with a composite burst-length distribution and/or different peak cell-rates, see Appendix B.2.

2. The second possibility is a mixture of different sources. Having n_i sources that produce traffic of class i only $(i = 1, ..., t)$, each of the classes has to be modeled separately by an n_i-Burst [IS] model with representation $\langle \boldsymbol{\mathcal{Q}}^{(i)}, \boldsymbol{\mathcal{L}}^{(i)} \rangle$ and the individual classes have to be merged explicitly which implies using the product space as a new state space:

$$\boldsymbol{\mathcal{Q}} = \boldsymbol{\mathcal{Q}}^{(1)} \oplus \boldsymbol{\mathcal{Q}}^{(2)} \oplus ... \oplus \boldsymbol{\mathcal{Q}}^{(t)}, \quad \text{and} \quad \boldsymbol{\mathcal{L}} = \boldsymbol{\mathcal{L}}^{(1)} \oplus \boldsymbol{\mathcal{L}}^{(2)} \oplus ... \oplus \boldsymbol{\mathcal{L}}^{(t)}.$$

Apart from the obviously different physical scenarios, the main feature with different implications on performance is: When using identical sources as in Case (1), every source can generate cells with any of the possible rates $\lambda_p^{(i)}$, so there is a wider choice of possibilities for the resulting Poisson rates than in Case (2). For example, the peak-rate in Case (1) is $\left(t \max(\lambda_p^{(i)})\right)$, while it is $\sum_{i=1}^{t} \lambda_p^{(i)}$ and thus smaller in Case (2) (assuming $n_i = 1$). Section 5.8 describes how this property affects performance, namely the number of so-called *blow-up points* is different in the two scenarios.

When using traffic mixes in the way described in Case (1), the different types of bursts are caused by different applications which might belong to different traffic classes which normally also show different back-off behavior in overload situations.. ATM for example offers real-time VBR traffic, that must not back off to the same extent, as e.g. ABR-traffic. Other protocols offer prioritized connections that do not have any or only moderate back-off behavior.

The scaling-factors $\beta(n)$ for the back-off of the cell-rates during a burst as introduced in the previous section do not allow for different back-off behavior yet. They have to be extended to be different for each type of traffic, as shown in Fig. 2.6. When using an ME representation of the ON times as described in Appendix B.2, a different series of back-off factors could be assigned even to every phase of the ME representation. Since this modification can lead to an enormously large parameter-space for the model, the discussion of the performance behavior does not investigate such models of traffic mixes with different back-off factors.

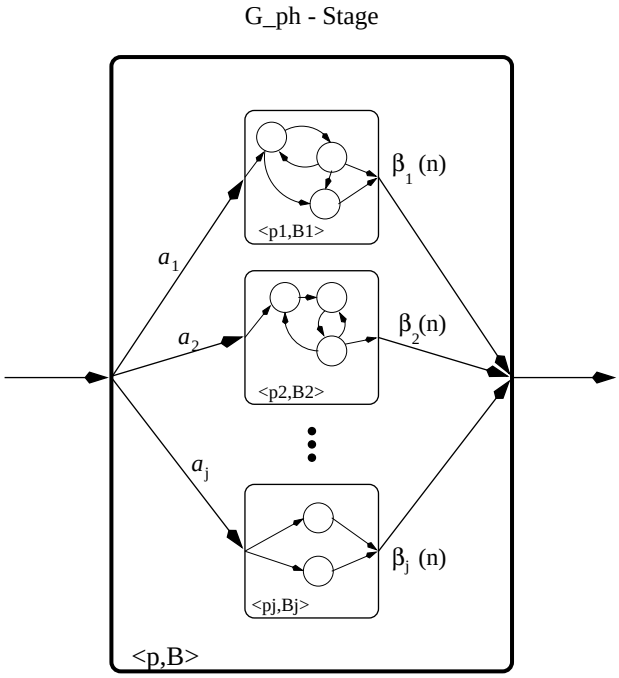

Figure 2.6: **Burst-Length Distribution in the Scenario of Traffic Mixes:** different traffic classes (e.g. rt-VBR, ABR) can have different back-off behavior. Therefore the scaling factors that determine the load-dependent behavior of the individual bursts, have to be extended to allow for different factors for different phases of the burst-length distribution.

2.5.3 Back-Off According to Life-Time of Burst

Another potentially relevant variant of the back-off behavior is introduced in this section: different back-off behavior according to the current life-time of the bursts. Again, the discussion of the performance behavior of such models is beyond the scope of this thesis.

When the active bursts are distinguished according to the order that they started, i.e. the longest active burst is assigned to Server 1 in S_1, the second longest to Server 2, etc., it is also possible to assign different back-off factors. For instance, in case of Power-Tailed burst-lengths it would be a better strategy in terms of throughput of bursts to decrease the cell-rate of the first few (older) bursts to spare more bandwidth for the newer bursts. The reason why that is more effective for burst-length distributions with high variance is that if the bursts have been active for a long time, the residual life-time will be very likely very long as well. If those long bursts are throttled, the rare huge events do not block the line as much and let the frequent small bursts pass by quicker.

The scaling factors $\beta^{(i)}(n)$ then have to be different for each load-dependent server i, $i = 1, ..., N$, in S_1 (see Figure 2.7). Consequently, the servers are not identical any more and have to be distinguished in the state-space representation of the $M/G_{ldS}/N//K$-System. A reduced state space representation as mentioned in Section E.2.1 cannot be used in this case.

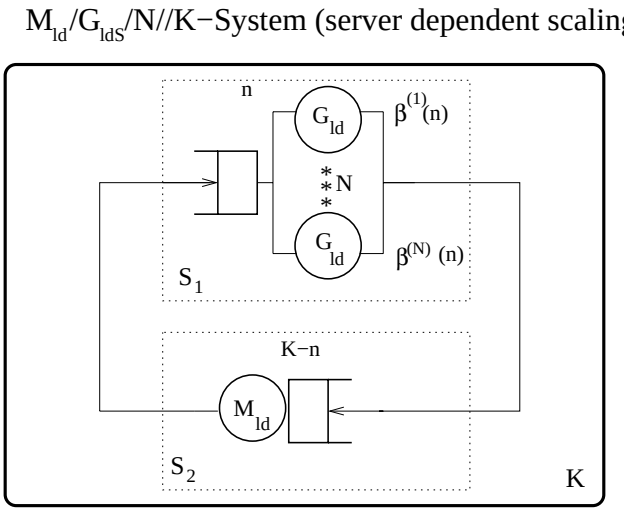

Figure 2.7: **Separate Load-Dependence Factors for Each Server in S_1:** if the order of the bursts matters for their load-dependent behavior, then different scaling factors $\beta^{(j)}(n)$ for the servers $j = 1, ..., N$ are necessary.

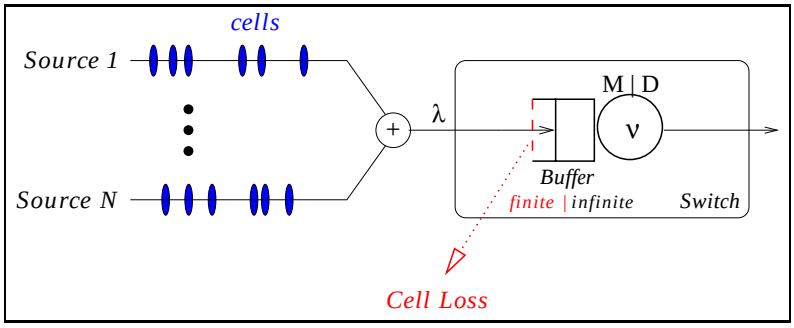

Figure 2.8: **N-Burst/M/1 Queue as a Model of an Output Port of an ATM-Switch**

2.6 Performance Models

So far, we introduced the N-Burst model which describes in its various variants the network traffic at the entrance of a network component, here assumed to be an ATM switch. For the model of the switch itself, we pick the simplest possible model: the output port of the ATM-switch is modeled by an exponential server of rate ν and some amount of FIFO buffer space, i.e. a queueing model of type N-Burst/M/1 as shown in Figure 2.8. Since ATM-cells have a fixed size of 53 bytes and thus the service times in reality are close to deterministic, the assumption of an exponential distribution for the service times represents an approximation. However, buffer sizes in ATM-switches are rather large (for instance 10^3 to 10^6 cells) and the time to drain a buffer of n cells is Erlangian-n distributed, which approximates the deterministic distribution very well for large n. Therefore, the assumption of exponential service times should have little impact on buffer-overflow in large-buffer systems. See [Lipsky & Schwefel 00] for a more detailed discussion of the impact of the service-time distribution.

In reality, physical buffers are always of finite, but potentially large, size. However, it also makes sense to investigate the behavior of models with infinite buffers, as explained below. The two different buffer-models are compared in the following:

- **Finite Buffer Models (Loss Systems):** The ATM-switch can store a large, but finite number of at most B cells in its buffer. Cells that arrive at a full buffer are discarded (queueing model of type N-Burst/M/1/B). The probability of such an event is called the *Cell Loss Probability* (CLP). A discussion of performance results is given in Section 5.5.

- **Infinite Buffer Models (Secondary Buffer Model):** In most telecommunication networks, dropped cells at the ATM layer are retransmitted by higher layers of the network protocols. So they still add to the load at the input link of the switch. Performance evaluation in that scenario can be approximated by using infinite buffer models: the finite primary buffer with a capacity of B cells is complemented by an infinitely large secondary buffer. Available buffer space in the primary buffer is instantaneously filled up with cells from the secondary buffer (queueing model of type N-Burst/M/1). The probability of a cell arriving at a full primary buffer, and therefore having to be stored in the secondary buffer, is called *Buffer Overflow Probability* (BOP).

Performance Parameters: The above description of the two different buffer models already listed two performance parameters, the BOP and the CLP, which describe essentially the same steady-state probability for a buffer-overflow event, but for different buffer models. Another performance parameter of interest here is the distribution of the per-cell delay, which is closely related to the buffer-occupancy or queue-length distribution. Chapter 5 discusses the behavior of those performance parameters in steady-state, i.e. if the observation interval of the system is long enough, such that an equilibrium value is observed. However, it turns out that steady-state analysis might not be sufficient, in particular for highly bursty traffic models. Chapter 6 therefore also discusses the transient counterparts of those performance parameters for finite observation intervals.

Another transient performance parameter of interest is the time until the first overflow (or loss) occurs, called *First Passage Time* $\tau(B)$. Its expected values, called *mean First Passage Time* (mFPT), is of concern in major parts of Chapter 6. Note that in contrast to cell delay or buffer-occupancy, the First Passage Time is independent of the buffer-model.

All these performance parameters can be computed numerically for the analytic N-Burst/M/1[B] model with the Matrix-Analytic methods that are summarized and also partially developed in Appendix F.

2.7 Approximating Performance Models

A large number of different traffic models are found in the literature of network performance analysis (see Section 1.2). Quite a few of these models can be seen as approximations that are obtained as a limit model of the N-Burst/M/1 queue when some parameter approaches a (physically unrealistic) limit. The advantage of the limiting models is that performance results are often easier to obtain, in best case, closed-form analytic solutions exist (e.g. for most steady-state performance parameters of the M/M/1 queue). The full N-Burst model then describes the transition region between the different limit models. First, that transition region is of interest since it contains the most realistic model description. Second, the description of the approximation models as limits of N-Burst/M/1 queues provides an understanding of the comparability or the diversity of the different models. In the following, we list several well known performance models and describe how they can be obtained as limits of the N-Burst model.

2.7.1 M/M/1 Queue

A simple M/M/1 queue with Poisson arrival rate $\lambda = N\kappa$ and service-rate ν is obtained if for fixed parameters N, κ, $\overline{n_p}$, and ν of the N-Burst/M/1 queue, the burstiness parameter b and with it the mean OFF time Z of the [IS] model is reduced to 0, which implies that λ_p reduces and $\overline{x_p}$ grows. In that limit, the ON periods are not separated by OFF periods any more, thus the generated traffic is purely Poisson. The actual distribution of the ON periods does not even matter in that limit. See [Lipsky & Schwefel 00] for a more elaborate discussion.

Performance results for the M/M/1 queue are well known, and in many cases, even closed form analytic expressions exist, see Appendix A. Although it is now widely accepted that the M/M/1 model is not adequate for traffic modeling in telecommunication systems ([Paxson & Floyd 95]), it is still widely used in practical design approaches.

2.7.2 Bulk Arrival Models

The limit $b \to 0$ of the N-Burst [IS] model for fixed N, κ, and $\overline{n_p}$ leads to pure Poisson traffic. At the other end, the limit $b \to 1$ for the same set of fixed parameters, the intra-burst cell rate $\lambda_p \to \infty$, and the duration of the ON periods becomes infinitesimally small, so all cells in a burst arrive simultaneously at the switch. Such a model is called a *bulk arrival* queue. The resulting mean per-cell delay can be computed by the following closed-form expression (see [Chaudhry & Templeton 83]):

$$\mathbb{E}\left\{\mathrm{CD}_{Bulk}\right\} = \frac{g_1 + g_2}{2\, g_1\, \nu\, (1-\rho)}\,,$$

where g_1 and g_2 are the first and second moments of the bulk-size distribution, i.e. the distribution of the number of cells per burst in the N-Burst model, see Appendix C.4. Consequently, in the limit $\lambda_p \to \infty$, the mean Cell Delay only depends on the first two moments of the burst-length distribution.

Since the time between two cells in a burst of the N-Burst model allows for some draining of the queue, the CD and also the BOP in the N-Burst model are lower than in the corresponding bulk-arrival system. Therefore, the performance parameters of the N-Burst model approach the ones of the bulk-arrival system from below for $\lambda_p \to \infty$. As such, the bulk-arrival model can be used to derive worst case upper bounds of the performance parameters.

2.7.3 Fluid-Flow Models

A fluid-flow limit of the N-Burst model is obtained if the number of cells per burst becomes infinitely large while at the same time, the server-speed and also the size of the buffer grow at the same speed. Then, the bursts consist of infinitely small units, which is exactly the fluid-flow approach. To be able to express the limit formally (see also [Lipsky & Schwefel 00]), let

$k :=$ Cell-size divisor (each cell is divided into k parts);
$\overline{n_p}(k) := k \cdot \overline{n_p} =$ New average number of cells during a burst;
$\lambda_p(k) := k \cdot \lambda_p =$ New peak transmission rate during a burst;
$\lambda(k) := k \cdot \lambda =$ New overall arrival rate;
$\nu(k) := k \cdot \nu =$ New mean cell service rate of switch ;
$B(k) := k \cdot B =$ New buffer-size of switch ;
$\rho(k) := \lambda(k)/\nu(k) = \lambda/\nu$ (remains unchanged).

For $k \to \infty$, the fluid-flow version of the N-Burst model is obtained. See [Lipsky & Schwefel 00] for the performance behavior in that case. Also, a discussion of the practical implications with respect to cell-sizes can be found at that reference.

2.7.4 Performance Models on Burst-Level: M/G/1 Queues

The N-Burst model works on cell-level, i.e. the arrival of each individual cell is modeled and its delay and other performance measures are investigated. A less detailed but therefore much easier computable model results, if cell-level is avoided completely and only whole bursts are considered. The service time of a whole burst (with first two moments b_1 and b_2) is then closely related to the burst-size distribution. The so-called *PK formula*, see [Kleinrock 75], allows to compute the expected value of the burst-delay, BD:

$$\mathbb{E}\left\{\mathrm{BD}_{M/G/1}\right\} = \frac{2\,{b_1}^2 - 2\,{b_1}^2\,\rho + b_2\,\rho}{2\,b_1 - 2\,b_1\,\rho} = b_1 + \frac{1}{2}\,\frac{b_2}{b_1}\,\frac{\rho}{1-\rho}\,.$$

The PK formula holds for all non-preemptive service strategies that do not make use of the actual service-times. However, *Processor Sharing* as it is proposed in Sect. 10.6 does not belong to that group. Except for single sections in Chapter 10, we restrict ourselves to queues with FIFO (First In First Out) service strategy. Note that the M/G/1 FIFO queue allows single large bursts to block the server. On the other hand, the N-Burst/M/1 queue intermixes the cells of simultaneously active bursts which results in shorter queues in most scenarios.

The BD of the M/G/1 model can be seen as a user-oriented performance measure in the limit $b \to 1$ (the bulk-arrival model). BD measures the time it takes to complete a single transfer. Again, [Lipsky & Schwefel 00] provides a more detailed discussion of that model.

2.7.5 Limits for $N \to \infty$: M/G/∞ Model, FGN, α-stable Levy Motion

Several other popular classes of traffic models can be obtained from the N-Burst model with the limit $N \to \infty$. Three main limits are of particular interest:

- M/G/∞ Model: The easiest case results when using the N-Burst [M] model. The limit $N \to \infty$ can be taken without any additional scaling of the process. It converges to a model in which bursts arrive according to a Poisson process with no upper limit on the number of active bursts (as in the case of finite N). Such models are used as a modulating process for the arrival rate to a queueing model for instance in [Guerin et al., 00].

- Fractional Gaussian Noise Models: If we start off with the N-Burst [IS] model, it is necessary to do some form of scaling when increasing the number of sources N, otherwise the total cell-arrival rate grows unboundedly with N. The details are omitted here, see [Lipsky & Schwefel 00] for reference. One particular way of scaling for the counting process[1] of N-Burst models results in the limit $N \to \infty$ in a (Fractional) Gaussian Noise Model, i.e. the counting process shows a purely Gaussian marginal distribution, but retains the asymptotic correlation structure of the single-source model. In case of Long-Range Dependent ON/OFF models (see Chapter 3), a so-called Fractional Gaussian Noise Model results. See [Norros 95] and [Krishnan 96] for performance results for that model class.

[1] See Sect. 1.3 for its definition.

- α-stable Levy Motion: A different way of scaling leads to a different limit model for $N \to \infty$ of the N-Burst [IS] Model, called α-stable Levy Motion. Again, see [Lipsky & Schwefel 00] for more information.

2.8 Summary

A class of performance models for telecommunications was introduced in this chapter. The traffic model, called N-Burst, is based on the superposition of ON/OFF models with Poisson cell-traffic during ON. Very general Matrix-Exponential distributions are thereby allowed for the duration of the ON periods. An upper limit N is imposed on the number of simultaneously active ON periods (active bursts). The formulation of this model in the framework of modulated Poisson processes with closed queueing system as modulating process has several advantages:

- Generality: Load dependent stages in the modulator allow different physical scenarios to be represented. Two important ones in terms of the burst-arrival process are the Independent Source Model [IS], and the [M] model with Poisson arrivals of bursts. By introducing load-dependent intra-burst cell rates and corresponding load-dependence in the ON-time distribution, back-off behavior, as for example seen in protocols with flow-control mechanisms, can be captured. An important issue for modeling of source level in modern telecommunication networks is also the potential to allow for mixes of traffic with different properties. It was shown how such traffic mixes can be incorporated in the model in two different ways.

- Tractability: The use of Matrix-Exponential distributions for the ON periods results in an MMPP representation of the N-Burst model. Analytic MMPP/M/1 models can be evaluated numerically, see [Meier-Hellstern & Fischer 92] and Appendix F.

- Operationality: The advantage of using the N-Burst model as opposed to a general (empiric) MMPP model is the underlying structure of the model. Its basic ON/OFF building blocks can be adjusted to real physical scenarios via operational (i.e. measurable) parameters. Also, that underlying structure allows for reasonable extrapolation of the model parameters, for instance in order to investigate the performance behavior for increased traffic volume.

This chapter introduced the model family without confusing the reader with the details about the MMPP representation. Those can be found in Appendix E. Furthermore, algorithms to compute performance parameters of the models are described and partially developed further in Appendix F. One of the problems with the numerical computation is the growing size of the state-space for more complicated model. Specially designed algorithms are discussed in the appendix to be able to deal with that problem, at least to some extent. A comparison of the efficiency of various algorithms is also given in Appendix F.

Chapter 3

Self-Similarity, Long-Range Dependence, and Truncated Power-Tail Distributions

One major motivation of this thesis is to investigate the significance of Long-Range Dependent properties of network traffic for performance. We introduced a class of very flexible traffic models, N-Burst, in the previous chapter without mentioning such properties so far. The N-Burst models belong to the broader class of MMPP models and it is known that due to the inherent exponential properties of MMPPs, no true Long-Range Dependence as in its mathematical definition can be reproduced[1]. However, on finite time-scales – and relevant time-scales for telecommunication networks are always finite, see the reasoning in Sect. 3.4 – Long-Range Dependent properties can very well be reproduced by MMPP models.

One such approach to mimick LRD properties by MMPP models is used by [Andersen & Nielsen 98], who use the superposition of several exponential ON/OFF models with different mean ON times to achieve such LRD properties. Our approach is different: an individual ON/OFF source can show LRD properties when the ON time distribution is appropriately chosen. A family of distributions, called *Truncated Power-Tail* (TPT) distributions, is utilized for that purpose. Furthermore, there is a considerable physical evidence (e.g., [Crovella & Bestavros 96], [Willinger et al. 95]) that such distributions should be included in models for network traffic, see also Sect. 3.2.

Power-Tail distributions are also occasionally proposed for the duration of the OFF periods. Later on in Section 5.7, it will be shown that the shape of the OFF time distribution is less critical for the performance behavior than the ON time distribution, so that the assumption of exponential OFF-times is acceptable, regardless of what the actual distribution turns out to be.

3.1 Self-Similarity and Long-Range Dependence

The notion of self-similarity has already been illustrated in Fig. 1.1 in the Introduction. Intuitively, self-similar network traffic means that regardless of the time-scale (e.g. the interval size Δ of the counting process, see Sect. 1.3), the traffic looks the 'same' when properly scaled. Of course, there is a more precise mathematical definition of the concept of self-similarity:

When self-similarity is used in the context of network traffic, it usually refers to the counting process $(N_i(\Delta))$, i.e. the number of cell arrivals in the i-th interval of length Δ:

[1] The same is true for the broader class of Semi-Markov processes, see Appendix C.1.

Definition 3.1.1 *The counting process* $(N_i(\Delta))_{i=1,2,...}$ *is called* self-similar *with Hurst parameter* H*, if for all* $s > 0$*, all finite dimensional distributions of the rescaled process* $\left(s^{-H} N_j(s\Delta)\right)_{j=1,2,...}$ *are the same as for the original process* $(N_i(\Delta))_{i=1,2,...}$.

A related concept is *second order self-similarity* whose definition relies on a similar property, but now only for the correlation structure. Let (X_i) be a *covariance stationary*[2] stochastic process.

Definition 3.1.2 *The process* (X_i) *is called* second order self-similar*, if the autocorrelation structure* $r^{(m)}(k)$ *of the (averaged) aggregated process,*

$$X_j^{(m)} := \frac{1}{m} \sum_{k=1}^{m} X_{jm+k},$$

is indistinguishable from the autocorrelation structure of the original process (X_i)*:*

$$r^{(m)}(k) \equiv r(k), \quad \textit{for all } m, k = 1, 2, \ldots .$$

Frequently (see e.g. [Leland et al. 94]) the above notion of second order self-similarity is also called *exact* second order self-similarity to make a clear distinction to *asymptotically* second order self-similar processes:

Definition 3.1.3 *The process* (X_i) *is called* asymptotically second order self-similar*, if*

$$\lim_{k \to \infty} \frac{r^{(m)}(k)}{r(k)} = 1 \quad \textit{for all } m,$$

i.e. the autocorrelation structure of the aggregated process $\left(X_j^{(m)}\right)$ *is identical to the autocorrelation structure of the original process* (X_i) in the limit of large k.

For network traffic, X_i could be the inter-cell times or alternatively the counting process for some interval Δ. Frequently in the literature, (asymptotic) second order self-similarity is used with respect to the counting process. Later in this chapter, we introduce a version of the N-Burst model that shows asymptotic second order self-similar properties[3] in both, the counting process and the inter-cell times.

For the models that are discussed in this thesis, it is sufficient to restrict ourselves to processes with non-negative autocorrelation function, $r(k) \geq 0$. The following definitions can also be modified for general $r(k)$ functions, but the technical details become more laborious in that case.

With the restriction to non-negative autocorrelation functions, the class of asymptotically second order self-similar processes is a subclass of the so-called *Long-Range Dependent* (LRD) processes ([Daley 00]):

Definition 3.1.4 *A covariance stationary process* (X_i) *is called* Long-Range Dependent *when its autocorrelation function* $r(k)$ *decays so slowly, that its sum diverges:*

$$\lim_{n \to \infty} \sum_{k=1}^{n} r(k) = \infty.$$

[2] i.e., the autocorrelation function $r(k) = r(X_i, X_{i+k})$ exists and depends only on k.

[3] we show later what we mean by that: we never work with processes that satisfy the definitions on all timescales. But as we discuss later, the actual mathematical limits are not meaningful for the practical application towards network traffic modeling.

A sufficient condition for LRD to occur is that the autocorrelation function drops off with a Power-Law with small enough exponent:

$$r(k) \sim \frac{1}{k^{\alpha-1}}, \quad 1 < \alpha < 2 \tag{3.1}$$

The smaller the value of α, the more slowly does the autocorrelation decay, thus α can be used as a measure of the degree of Long-Range Dependence for such processes. Whenever we talk about LRD processes in the following, we refer to processes with the property (3.1). Such LRD processes, whose autocorrelation function decays as a Power-Law, can be shown to be asymptotically second order self-similar ([Daley 00]).

Other manifestations of second-order self-similarity are described in Sect. 3.7, in which they are used to estimate the exponent α.

There is one particular self-similar process that is widely used. It is called *Fractional Brownian Motion*, and its increment process is called *Fractional Gaussian Noise* (FGN). The marginal distribution of the increments is always Gaussian – hence the name. Its Hurst parameter can be in the range $0.5 < H < 1$. FGN is also exactly second order self-similar and LRD with an exponent $\alpha = 3 - 2H$. Although the Hurst parameter is strictly speaking reserved for self-similar processes according to Definition 3.1.1, LRD processes show similar scaling properties with an equivalent 'Hurst-Parameter' $H = (3 - \alpha)/2$. In the R/S statistics in Section 3.7.2, that 'Hurst-Parameter' for LRD processes shows up.

3.2 Power-Tail Distributions

Long-Range Dependent properties have been found in a large number of measurements of network traffic, one of the earliest is described in [Leland et al. 94]. However, the discussion about the impact of the LRD property on the performance behavior is not yet settled at this time. One major motivation for this thesis is to include LRD properties in the N-Burst model in order to find out about its impact on performance behavior.

Traditional ON/OFF models with exponentially distributed ON and OFF durations (as in [Heffes 80]) show a correlation structure – both for inter-cell times and counts – that drops off exponentially. LRD effects can be achieved when so-called *Power-Tail (PT) distributions* are introduced either for the duration of the ON or the OFF periods, or both. Such PT distributions are characterized by a complementary distribution function (reliability function) $R(x)$ that drops off very slowly, namely by a Power-Law with exponent $\alpha > 0$:

$$R(x) \sim \frac{1}{x^{\alpha}}\,. \tag{3.2}$$

Due to the slow decay, the probability that huge events will occur is never negligible. Also, all moments $\mathbb{E}\left\{X^{\ell}\right\}$ for $\ell > \alpha$ are infinite. Thus, if $\alpha \leq 1$ the distribution has infinite mean, and if $\alpha \leq 2$ then the variance is infinite. Such infinite variance distributions can be the cause of LRD properties in ON/OFF models, see Sect. 3.6.

The constant

$$c_{\mathrm{PT}} := \lim_{x\to\infty} x^{\alpha} R(x)\,, \tag{3.3}$$

is called the *tail-constant* of the PT distribution $R(x)$. The tail-exponent α together with the tail-constant fully characterizes the tail of a PT distribution.

There is also further encouragement for the use of such PT distributions from measurements of distributions of file-sizes in both standard file-systems and WWW-servers. [Garg et al. 92] found Power-Tail behavior in the sizes of all individual files of a UNIX file-system. This is confirmed in [Gribble et al. 98] who measured the distribution of the size of requested files and found a Power-Tail with $\alpha = 1.48$. [Crovella & Bestavros 96] found Power-Tail distributions for the file-sizes at a WWW-Server. Therefore, at least for FTP-traffic and WWW-traffic, PT distributions need to be included in the traffic-model. Also, the analysis of Ethernet traffic between specific source-destination pairs in [Willinger et al. 95] showed ON/OFF behavior with Power-Tailed properties.

One type of a PT distribution is the following *Pareto Distribution*:

$$R(x) = \frac{1}{\left(\frac{x}{(\alpha-1)\,\bar{x}} + 1\right)^{\alpha}} \quad \text{for} \quad x > 0,\ \alpha > 1 \cdot \tag{3.4}$$

When $\alpha > 1$, this distribution has the expected value $\mathbb{E}\{X\} = \bar{x}$. See Sect. I.1 for a table with the tail-constants of the Pareto distribution for different values of α.

Although the Pareto function (3.4) has a simple analytic closed-form expression, it is hard to come up with tractable analytic models that make use of this distribution. Consequently, simulation models are often used instead of analytic models. The generation of random variables from Power-Tail distributions is simple, however the analysis of the simulation results can be very tricky, see [Lipsky et al. 99b] and also parts of Chapter 6.

The interesting question now is: can we integrate PT distributions in the N-Burst models? In Chapter 2, the N-Burst models were introduced as the aggregation (in various ways) of several ON/OFF models with ME distributed ON times. Unfortunately, Power-Tail distributions do not have a finite-dimensional ME representation. As a consequence, another step is possible, which is the introduction of truncated Power-Tails in Sect. 3.4. There, further arguments, other than tractability, are given as to why truncated tails are also more reasonable for the practical application in network traffic modeling.

3.3 Properties of Power-Tail Distributions

Before we truncate the Power-Tails that we just introduced, we discuss a few properties of untruncated Power-Tail distributions. Later on, PT distributions will be used for the ON periods of the N-Burst model. During the performance analysis, three particular random variables that are derived from the PT distributed burst duration, A, with exponent α will be of particular interest:

1. Distribution of Overshoot for threshold x_0:
   Knowing the fact that the burst duration has already lasted for time x_0, the probability that it will last for another time period of at least length x_1 behaves asymptotically for large x_0 as:
   $$Pr\left(A > x_0 + x_1 \,|\, A > x_0\right) \sim \left(\frac{x_0}{x_0 + x_1}\right)^{\alpha}$$
   Note that the probability grows with increasing threshold x_0, thus the longer the burst has already been active, the larger the probability that it will remain active for another period of at least length x_1. This is a peculiar property of PT distributions. For an exponentially distributed burst duration, the probability would be independent of x_0 due to its memoryless property.

The expected value of the 'overshoot', $\tilde{A}(x_0) = (A - x_0 \mid A > x_0)$, in excess of the threshold x_0 turns out to be

$$\mathbb{E}\left\{\tilde{A}(x_0)\right\} = \int_{x_1=0}^{\infty} x_1 \, \mathbb{P}\left\{A > x_0 + x_1 \mid A > x_0\right\} dx_1 \sim \frac{x_0}{\alpha - 1} \quad \text{for} \quad \alpha > 1\,.$$

The expected value of the overshoot also grows with the threshold x_0. The growth is asymptotically linear for large x_0. When the PT exponent α is close to 1, the expected overshoot grows rapidly.

2. Distribution of the Residual Time A_{res}:
If it is not known how long the currently active burst has already been active (i.e. in the viewpoint of a random observer), the density function of the residual time A_{res} can be shown to be the renormalized reliability function of the original random variable A, see [Kleinrock 75]:

$$f_{\text{res}}(x) = \frac{1}{\bar{x}}\, R(x)\,, \tag{3.5}$$

Eq. (3.5) is only meaningful if $\bar{x} = \mathbb{E}\{A\}$ is finite.

If the original burst duration is Power-Tailed with tail-exponent $\alpha > 1$ and tail-constant c_{PT}, then it follows from Eq. (3.5) that the residual time A_{res} is also Power-Tailed, but with smaller exponent $\alpha - 1$. Thus, the reliability function of the residual time decays even more slowly than for the original burst duration. The tail-constant of the Power-Tailed residual time follows as

$$c_{\text{PT}}(A_{res}) = \frac{c_{\text{PT}}}{(\alpha - 1)\,\bar{x}} \quad \text{for} \quad \alpha > 1\cdot$$

3. Simultaneous Power-Tail distributions:
In models with an aggregation of several simultaneous bursts, it turns out in the analysis of Chapter 5 that the distribution of the duration of i simultaneously active bursts is of critical importance (for some special i, see Chapter 5). If we assume that a burst with Power-Tailed duration, $R(x) \sim 1/x^{\alpha}$, starts while another independent burst of the same kind is already active, the time-period that both of them are active has a duration that is distributed as

$$R_2(x) = R_{\text{res}}(x)\, R(x) \sim \frac{1}{x^{\alpha-1}} \cdot \frac{1}{x^{\alpha}} = \frac{1}{x^{2\alpha-1}}\ \cdot \tag{3.6}$$

That time period ends if either the first burst ends, as described by the residual time $R_{\text{res}}(x)$, or independently, the second burst ends, whose reliability function is $R(x)$.

If another, third independent burst starts while both of the other sources are active, the same argument leads to

$$R_3(x) = R_{2,res}(x)\, R(x) \sim \frac{1}{x^{3\alpha-2}},$$

and by induction it follows for the duration of a time period with the i long-term active bursts:

$$R_i(x) \sim \frac{1}{x^{i\alpha-i+1}}\ \cdot \tag{3.7}$$

In summary, the distribution of the time-period with i independent bursts simultaneously active is also Power-Tail distributed, but with smaller exponent $i\alpha - (i-1)$.

An analog derivation can be made for the tail-constants.

3.4 Matrix Exponential Representations and Truncated Tails

We mentioned in Sect. 3.2 that Power-Tail distributions do not have a finite-dimensional ME representation. However, for the purpose of network traffic modeling, the question is, how 'real' are infinite Power-Tails? There are physical limits on e.g. file sizes, even though these limits tend to grow with developing technology. Therefore, we would expect that the PT behavior $R(x) \sim x^{-\alpha}$ of the ON time distributions is cut off at some point. Furthermore, the utilization of computer networks shows a very regular fluctuation during the course of a day: normally, there exist the so-called busy hours, during which the network is heavily utilized. Those 4–8 busy hours are of particular interest for performance modeling, because that is the time when performance problems could be expected. If we want to avoid the use of time-varying model parameters – which we do here – there is only hope to describe the network traffic realistically during those busy hours, after which no problems are expected in any case. Therefore, any traffic model would not have to worry about Power-Tailed distributed ON periods with a duration longer than those busy hours. Thus the use of truncated tails is more reasonable, if not to say necessary, for a realistic traffic model.

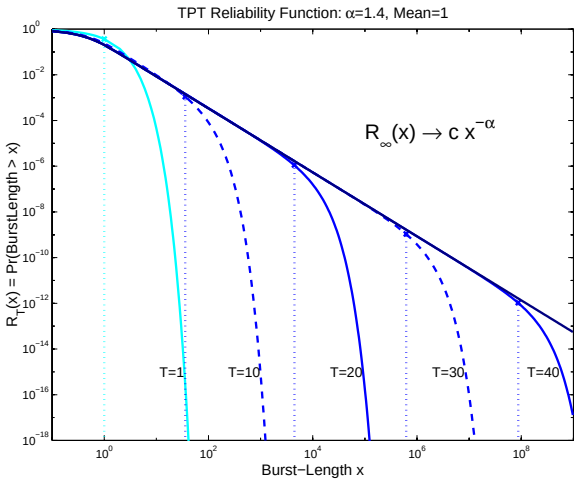

Figure 3.1: **Log-Log Plot of the Reliability Function, $R_T(x) = \mathbb{P}\{X > x\}$, for Truncated Power-Tail Distributions with T Phases:** The TPT reliability function shows a Power-Law behavior – appearing linear with negative slope α – for some range of x, before it drops off exponentially. The PT-range (marked by dotted lines) can be extended arbitrarily by using more phases.

In this thesis, we use a family of Hyper-exponential distributions that was originally introduced by [Greiner et al. 99]:

$$R_{Y_T}(x) = \frac{1-\theta}{1-\theta^T} \sum_{i=0}^{T-1} \theta^i \exp\left(-\frac{\mu_T}{\gamma^i}\, x\right), \tag{3.8}$$

That family of distributions, called *Truncated Power-tails* (TPT), provides a systematic way to obtain PT behavior over a controllable range. With growing number of phases, these distributions asymptotically approach a Power-Tail distribution in the 'weak definition':

$$R_{Y_\infty}(x) \overset{\bullet}{\sim} \frac{1}{x^\alpha}$$

See Sect. 1.3 for the definition of '$\overset{\bullet}{\sim}$'.

In order to show the Power-Law behavior (3.2) with the tail-exponent α, and in order to have mean $\bar{x}$, the constants in (3.8) have to be:

$$\begin{aligned} & 0 < \theta < 1, \\ & \gamma = \left(\frac{1}{\theta}\right)^{1/\alpha}, \\ & \mu_T = \frac{1-\theta}{1-\theta^T} \frac{1-(\theta\gamma)^T}{1-\theta\gamma} \frac{1}{\bar{x}}. \end{aligned}$$

Figure 3.1 demonstrates that the reliability functions of the TPT distributions clearly show a Power-Law behavior (appearing linear on log-log scale) for some range of x before they drop off exponentially. A characterization of the so-called Power-Tail Range is given in the next section. The larger the number of phases, T, the later the exponential drop-off occurs. Together with the a priori set parameter θ (here usually set to $\theta = 0.5$), that is a measure of how dense the individual TPT phases fill in the tail, the number of phases can be arbitrarily extended to provide TPTs with arbitrarily large PT Range. This is the critical advantage of using that systematic approach of the family of TPT distributions as opposed to curve fitting techniques, which are for example used in [Feldman & Whitt 98] to obtain an approximation of PT distributions.

Since the TPT distributions are just a special family of Hyper-exponential distributions, their ME representation follows easily from Eq. 3.8. The matrices are also listed in Appendix B.2. The tail-constants c_{PT} of the TPT distributions are also listed in Appendix I.1.

General characterization of truncated tails: For the models in this thesis, it is sufficient to restrict ourselves to TPT distributions, whose truncation is caused by an *exponential* drop-off. More general concepts are possible, but the specific shape of the drop-off is of minor impact for the performance behavior of the models in this thesis. We can classify the values of a TPT distributed random variable in three different regimes.

- The body of the distribution:
  The reliability function $R(x)$ for values x up to the order of several times the expected value of the random variable is called the body of the distribution. In this thesis, the distribution body is not the major focus, since the goal is to discuss the impact of the distribution tail instead. However, note that (3.8) can be modified to fit the body of any distribution arbitrarily closely.

  In contrast to the distribution body, the tail of the distribution refers to $R(x)$ with $x \gg \mathbb{E}\{X\}$. The tail of a TPT distribution consists of the following two parts:

- The **Power-Tailed** region of the tail:
  The Power-Tailed part of the tail is specified by the *PT exponent* α and the *tail constant* c_{PT}:
  $$R(x) \approx \frac{c_{\mathrm{PT}}}{x^\alpha}.$$
  It comprises values of $x \gg \mathbb{E}\{X\}$ up to the value, at which the tail is truncated, $x < \mathrm{Rng}(X)$, see the next section.

- The **exponential tail**:
  At some point the reliability function starts to drop off exponentially. That exponential tail can be described by its decay rate λ and another tail-constant, d_{TPT}:
  $$R(x) \sim d_{\mathrm{TPT}}\, e^{-\lambda x}.$$

The inverse of the decay rate is called *Power-Tail Range* $\overline{x_T}$ in the next section, since it determines approximately, for what x the regime of the exponential tail is reached.

If the transition from the Power-Tailed region to the exponential tail was sharply at $\overline{x_T}$, the additional exponential tail-constant would be $d_{\mathrm{TPT}} = c_{\mathrm{PT}}/\overline{x_T}^{\alpha}$. However, that is normally not the case in practice since there is a gradual transition between the different tail-regions. In particular, the separate tail constants are necessary for the family of TPT distributions that is used herein.

3.5 Power-Tail Range and Maximum Burst Size

To be able to discuss the impact of truncated Power-Tails, we need to introduce a notion for the location of the truncation. We restrict our investigations to the class of distributions that eventually drop off exponentially (or geometrically in the discrete case). For those distributions, we define the *Power-Tail Range* as the mean of the exponential tail or equivalently the inverse of its decay rate:

$$R(x) \sim \exp(-\lambda\, x) \quad \Longrightarrow \quad \mathrm{Rng}\,(R(x)) := \frac{1}{\lambda}\,. \tag{3.9}$$

The motivation for that definition becomes obvious in particular for the case of Hyper-Exponential distributions: $\mathrm{Rng}\,(R(x))$ is the mean of the largest exponential phase. That largest phase is mainly responsible for the drop-off of the Reliability function. The definition as $\mathrm{Rng}\,(R(x)) = 1/\lambda$ is somewhat arbitrary: any definition as a/λ with a not too large (e.g. $a = 2$) would also be a reasonable choice. However, it is important to pick one definition and use it consistently. Note that the exponential decay rate is the same in the Reliability Function $R(x)$ as it is in the density function $f(x) = -dR(x)/dx$.

It can be easily derived that the PT Range of such distributions can also be obtained by the following limit:

$$\mathrm{Rng}\,(R(x)) = \lim_{x\to\infty} \frac{R(x)}{f(x)} = \lim_{x\to\infty} \frac{R(x)}{-dR(x)/dx}\,, \tag{3.10}$$

where $R(x)$ is the reliability function and $f(x)$ is the density function of the distribution. The main purpose of Eq. (3.10) is to distinguish exponential tails from Power-Tails, and also to locate the exponential tails. For Power-Tail distributions with infinite tails on the other hand, the limit $R(x)/f(x)$ grows unboundedly:

$$\frac{R_{PT}(x)}{f_{PT}(x)} = \frac{x^{-\alpha}}{\alpha\, x^{-(\alpha+1)}} = \frac{x}{\alpha} \;\longrightarrow \infty.$$

Appendix (B.3) shows how to evaluate the range for an arbitrary Matrix-Exponential distribution. However, note that such a computation for distributions of Erlangian type can be misleading since the derived PT Range is smaller than its mean, see also right graph of Fig. 3.2.

Note that the existence of $\mathrm{Rng}\,(R(x))$ does not guarantee that actual PT behavior of $R(x)$ can be observed for $x < x_0$, but it rather locates the start of the exponential tail at x_0.

Although the Power-Tail Range has a clear intuitive meaning, there are several possible meaningful quantitative definitions. One possibility, (3.9), is used in this thesis and it turned out to be successful. However, that definition is restricted to distributions with exponential (or geometric) drop-off. That restriction is sufficient for the models in this thesis. Also, the actual shape of the

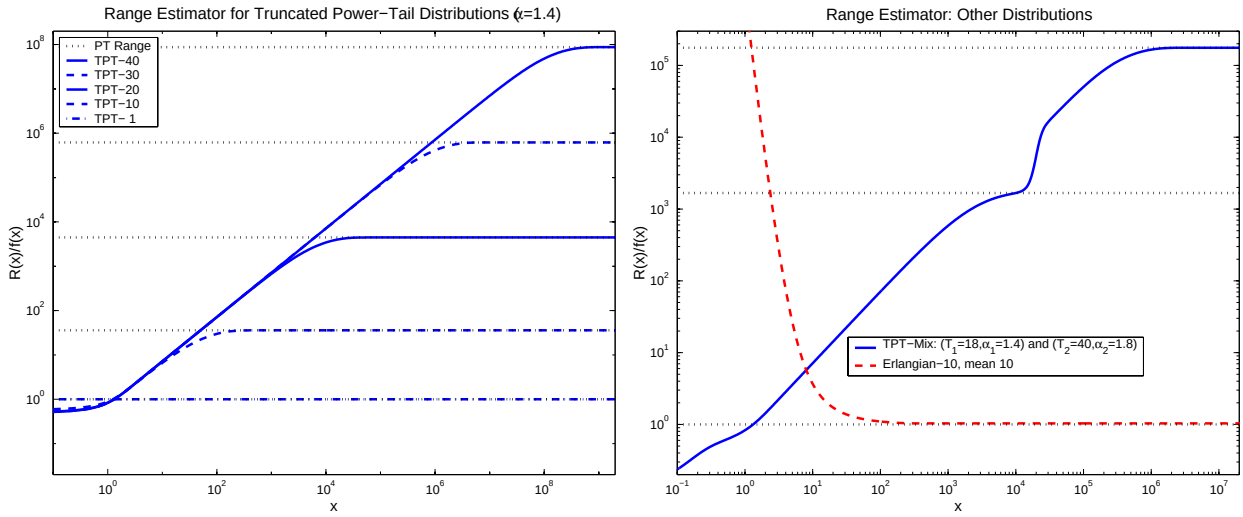

Figure 3.2: **Power-Tail Range of various Distribution Functions:** The ratio $R(x)/f(x)$ for TPT distributions with different number of phases is shown in the left graph. It converges monotonically to the PT Range as computed from the ME representation, see App. B.3. The PT range of an Erlangian distribution is smaller than its mean, see dashed line in the right graph. When computing $R(x)/f(x)$ for a mixture of two TPT distributions, the PT Ranges of both distributions appear as plateaus of $R(x)/f(x)$.

drop-off is of little practical impact, since the Power-Tailed region causes the performance critical behavior. Therefore, any other shape of the drop-off can be approximated by an exponential drop-off, without affecting the practical relevance of the results.

A second possible definition of the PT Range would be the following:

$$\mathrm{Rng}'(R(x)) := \frac{1}{c_{\mathrm{PT}}} \int_0^\infty x^\alpha \, R(x) \, dx \, .$$

Asymptotically for large PT Range, this definition agrees with (3.9) for the TPT distributions of Sect. 3.4. Its disadvantage is that α and c_{PT} must be known to apply this definition. Therefore, we do not investigate this definition any further, but we exclusively use (3.9) in the following.

Due to the simple Hyper-exponential form of the TPT distributions in Sect. 3.4, their PT Range follows easily as

$$\mathrm{Rng}\,(R_{Y_T}) = \frac{\gamma^{T-1}}{\mu_T} =: \overline{x_T} \, . \tag{3.11}$$

Figure 3.2 plots the behavior of $R(x)/f(x)$ for several TPT distributions in the left graph. Also shown in the right graph is an Erlangian-10 distribution and a mixture of two TPT distributions with different exponents and different PT Ranges. Note that the individual PT Ranges of both TPT distributions in the mixture can be observed as plateaus in the R/f curve.

Discrete Distributions: An equivalent definition of a Power-Tail Range for discrete distributions is possible. Instead of an exponential decay, we now talk about a *geometric* decay:

$$p_k \sim a^k = e^{-(-\ln a)k} \quad \text{for} \quad 0 < a < 1.$$

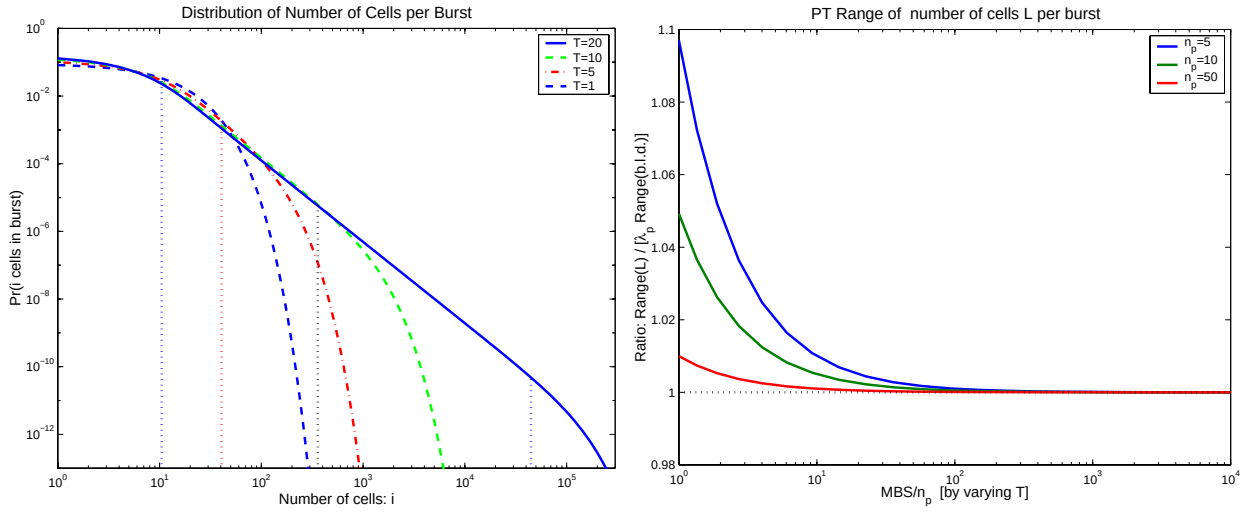

Figure 3.3: **The distribution, ξ_n, of the number of cells per burst for TPT distributed burst duration:** The truncated Power-Tails of the burst-duration in turn cause truncated Power-Tails for the distribution of the number of cells. Asymptotically, the PT Range of ξ_n is given by the expected number of cells in a burst whose duration is the PT Range of the ON time.

The definition of the PT Range for the discrete case is then analogue to the continuous case (3.9):

$$p_k \sim a^k \quad \Longrightarrow \quad \mathrm{Rng}\,(p_k) := \frac{1}{-\ln a}\ . \tag{3.12}$$

Appendix B.3 shows how to compute the PT Range for more general Matrix-Geometric Distributions. Such distributions appear in various places in this thesis, for example for the queue-length distribution of MMPP/M/1 queues or for the distribution of the number of cells per burst, which we consider in the following.

Maximum Burst Size: The ME distribution $\langle \mathbf{p}, \mathbf{B} \rangle$ of the duration of the ON periods in the N-Burst model is of course directly related to the distribution of the number of cells during such a burst. The cells during ON periods are generated according to a Poisson process with rate λ_p Appendix C.4 shows that the distribution of the number of cells is Matrix-Geometric:

$$\xi_n = \mathbf{p}\,\lambda_p^n (\mathbf{B} + \lambda_p \mathbf{I})^{-(n+1)}\,\mathbf{B}\boldsymbol{\epsilon}'\,.$$

Note that there is a positive probability $\xi_0 > 0$

$$\xi_0 = \mathbf{p}\,(\mathbf{B} + \lambda_p \mathbf{I})^{-1}\,\mathbf{B}\boldsymbol{\epsilon}'\,.$$

that an empty burst occurs. At first, such empty bursts seem to be an anomaly. However, if the mean number of cells per burst, $\overline{n_p}$, is not too small, the probability ξ_0 is rather small. Secondly, the calibration methods that are presented in Chapter 4 take the probability ξ_0 into account, so although the empty bursts are possible in the model, they just look like extended OFF periods, but the data is still matched. Empty bursts can be avoided by more complicated SM models that are outside the MMPP class. However, this is beyond the scope of this thesis.

Figure 3.3 shows in its left graph a log-log plot of the distribution ξ_n of the number of cells per burst for TPT distributed ON duration. The plots show that the distribution of the number of cells also shows a truncated Power-Tail. Its PT Range is shown in Appendix B.3 to be

$$\mathrm{Rng}\,(\xi_n) = \frac{1}{\ln\left[1 + 1/(\lambda_p\,\mathrm{Rng}\,(\mathbf{B}))\right]}\,,$$

which converges to $\mathrm{Rng}\,(\xi_n) \sim \lambda_p\,\mathrm{Rng}\,(\mathbf{B})$ for large $\mathrm{Rng}\,(\mathbf{B})$. The right graph of Figure 3.3 shows that convergence occurs rather quickly. We define the *Maximum Burst Size* (MBS) as the approximate PT Range of the distribution of the number of cells

$$\mathrm{MBS} := \lambda_p \cdot\ \mathrm{Rng}\,(\mathbf{B}_T)\,. \tag{3.13}$$

Thus bursts with more than MBS cells only happen with very small (geometrically decaying) probability.

3.6 Autocorrelation Structure of N-Burst Models

It is proven in various papers, e.g. [Greiner et al. 99b], that Power-Tail distributed ON periods have the consequence that the traffic from ON/OFF models (both inter-cell times and counts) is Long-Range Dependent.

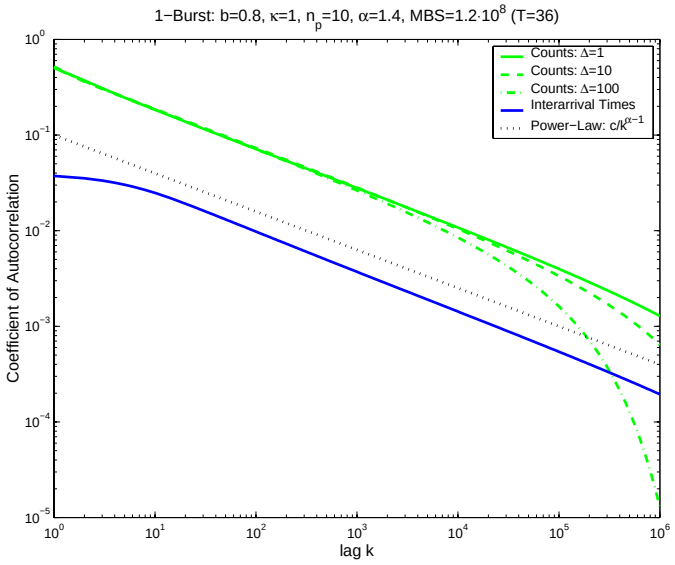

Figure 3.4: **Autocorrelation of Inter-Cell Times and Counts for a 1-Burst model**: truncated Power-Laws with exponent $\alpha - 1$ can be observed for both processes.

For the N-Burst model with TPT distributed ON times, we can compute the auto-correlation numerically in the MMPP representation, see Appendix C.3. Figure 3.4 shows the results for the inter-cell times and the counting process for various interval lengths Δ: All curves show a (truncated) Power-Tail with exponent $\alpha - 1$ in the limit for large T. The truncated ON time in turn causes a truncation of the autocorrelation functions. For the counting process, it can be observed that the larger the interval Δ is, the earlier does the exponential drop-off occur.

3.7 Statistical Methods

We briefly describe a number of statistical tests for Power-Tail distributions and for Long-Range Dependence. Such tests normally provide an estimate of the Power-Tail exponent α which will be necessary in Chapter 4 as part of the parameter estimation procedure for the N-Burst model. However, the list of statistical methods that is given here is in no way complete. See [Beran 94] and [Adler et al. 98] for more information and a more thorough statistical analysis of such estimation procedures.

3.7.1 Test for Power-Tail Distributions

There are several methods to test whether samples (X_i) are generated by a PT distribution, and to estimate the tail-exponent α. We mention two of them here:

Hill Estimator

The Hill-estimator is derived from properties of the order statistics of samples from Power-Tail distributions, We omit the background and just state the estimator here, see [Hill 75] for more details:

$$\hat{\alpha}(k) = \left[\frac{1}{k} \sum_{i=0}^{k-1} (\log X_{[n-i]} - \log X_{[n-k]}) \right]^{-1} , \tag{3.14}$$

where $X_{[n-i]}$ is the ith largest element of the total of n samples. Thus the estimator $\hat{\alpha}(k)$ is based on the k largest samples of the given data-set. For growing values of k, the estimator should eventually stabilize at values close to α.

Mean Excess Function

In Section 3.3, we pointed out that a peculiar property of Power-Tail distributions is that the expected 'overshoot' $X - x_0$ in excess of some threshold x_0 grows with increasing value of the threshold as

$$\mathbb{E}\{X - x_0 \mid X \geq x_0\} \sim \frac{1}{\alpha - 1} x_0 \,. \tag{3.15}$$

This property can be used in an estimate of the expected overshoot of the samples for several threshold values. If the estimate grows roughly linearly, a PT distribution can be expected to be involved and the tail-exponent α can be estimated by linear regression and Eq. (3.15).

This method was used in [Leland & Ott 86] to point out one of the first significant occurrences of PT distributions in computer systems: they found that the execution times of programs show Power-Tailed durations.

3.7.2 Estimators for LRD Properties

LRD manifests itself in a number of ways that can be used for an estimation of the exponent α. All methods described here fit a Power-Law to some data (x, y), which is done by linear regression on $(\log x, \log y)$. The problem in practice is that LRD is an asymptotic property,

so the linear regression has to be done for some sufficiently large set of x. For simplicity, we determine that range of values by inspection – but the resulting estimates can depend quite a bit on the chosen range. No good automatic approaches are known for this problem, and it is not the goal of this thesis to develop those.

In the following, three different estimators for the PT exponent α that is involved in the definition (3.1.4) of LRD are described. The estimators can be applied to either inter-cell times or counts. Here, we exemplarily apply them to inter-cell times that are generated by simulation of N-Burst models.

Autocorrelation Plot

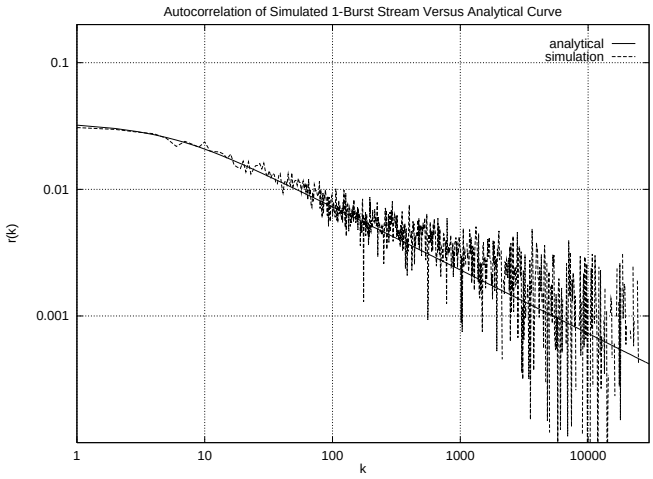

Figure 3.5: **Comparison of an Estimate of the Autocorrelation curve from Simulated 1-Burst Samples with the Exact Numerical Results from the Analytic Model:** the simulation uses $n = 7.8 \cdot 10^5$ samples, which is comparable to the size of several measurements of ATM inter-cell times in Appendix I.2.

The first method is based on the definition of LRD:

$$r(k) \sim k^{1-\alpha}\,.$$

Consequently, the estimated autocorrelation function, $\hat{r}(k_j)$, at lags k_j from the samples X_i of the inter-cell times can provide an estimate for α by linear regression on the points $(\log k_j, \log \hat{r}(k_j))$. Figure 3.5 shows such an estimate for the correlation function for a sample set of $n = 7.8 \cdot 10^5$ inter-cell times. This is approximately the sample-size of the smaller measurements that are used in Chapter 4. For large values of k, the estimated correlation function becomes rather noisy, compared to the straight line that is obtained from the analytic model. Nevertheless, linear regression can be applied to estimate the slope in the log-log plot.

R/S Statistics

The *rescaled adjusted range* or R/S statistics is defined as the ratio $R(n)/S(n)$ for a subset of n samples where

$$R(n) := \max(0, W_1, \ldots, W_{n-1}) - \min(0, W_1, \ldots, W_{n-1})\,, \tag{3.16}$$

$$\text{with} \quad W_k := \sum_{i=1}^{k}(X_i - \bar{X}_{(n)}), \quad \bar{X}_{(n)} := \frac{1}{n}\sum_{i=1}^{n} X_i.$$

W_k with $k < n$ sums up the differences between the first k samples of the block from the sample mean of all the n samples in the block. The denominator

$$S(n) = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(X_i - \bar{X}_{(n)})^2},$$

is just an estimate for the standard deviation of the n samples in the block.

It can be shown (see e.g. [Beran 94]) that for an increasing block-size n, the R/S statistics show the following asymptotic behavior:

$$\frac{R(n)}{S(n)} \sim n^h,$$

where $h = H$ is the *Hurst-parameter* for a self-similar process. For a Long-Range Dependent process whose autocorrelation function behaves as $r(k) \sim 1/k^{\alpha-1}$ with $1 < \alpha < 2$, the ratio $R(n)/S(n)$ grows asymptotically as a Power-Law with exponent $h = (3-\alpha)/2$. For any well-behaved process (short-range dependent), an exponent $h = 0.5$ is observed in the asymptotic growth of the R/S statistics.

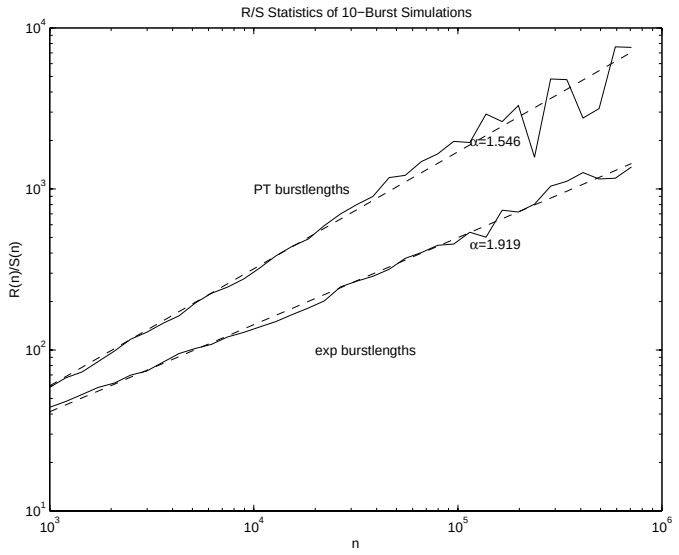

Figure 3.6: **R/S-statistics of 10-Burst Simulations:** The upper curve uses PT-distributed burst-lengths with tail-exponent $\alpha = 1.5$, while the lower curve uses exponentially distributed burst-lengths. For the latter, the resulting $\hat{\alpha}$ is close to 2 which is a sign of lacking Long-Range Dependence.

Figure 3.6 shows the resulting graphs for a simulation of the 10-Burst process. The resulting estimate $\hat{\alpha} = 1.55$ is slightly higher than the true model parameter $\alpha = 1.5$. The exponentially distributed burst-lengths result in $\hat{\alpha}$ of 1.92, thus close to the expected value 2.

The single value of $R(n)/S(n)$ only uses n samples. Therefore the whole sample-space of $7.8 \cdot 10^5$ samples is partitioned in up to 30 equally sized blocks, and the average of the 30 individual values of $R(n)/S(n)$ for the individual blocks is plotted in Figure 3.6. The number of partitions gets less with increasing n since only non-overlapping partitions are used.

Variance-Time Plot

The last method, called *Variance-Time Plot*, uses the aggregated data stream $X_j^{(m)}$:

$$X_j^{(m)} := \frac{1}{m} \sum_{i=(j-1)m+1}^{jm} X_i, \quad j = 1, ..., \left\lfloor \frac{n}{m} \right\rfloor .$$

The variance of the aggregated stream then shows the following behavior (see e.g. [Beran 94]:

$$\frac{\mathrm{Var}(X^{(m)})}{\mathrm{Var}(X)} \sim m^{1-\alpha} \quad \text{for LRD processes.}$$

If the method is applied to a data stream with short-range or no correlation at all, the variance of the aggregated stream decays more quickly as m^{-1} – for the latter case of a renewal process $\mathrm{Var}(X^{(m)}) = 1/m^2 \cdot m \cdot \mathrm{Var}(X) = m^{-1} \cdot \mathrm{Var}(X)$. Therefore, estimates of α close to 2 are a sign of the absence of LRD properties.

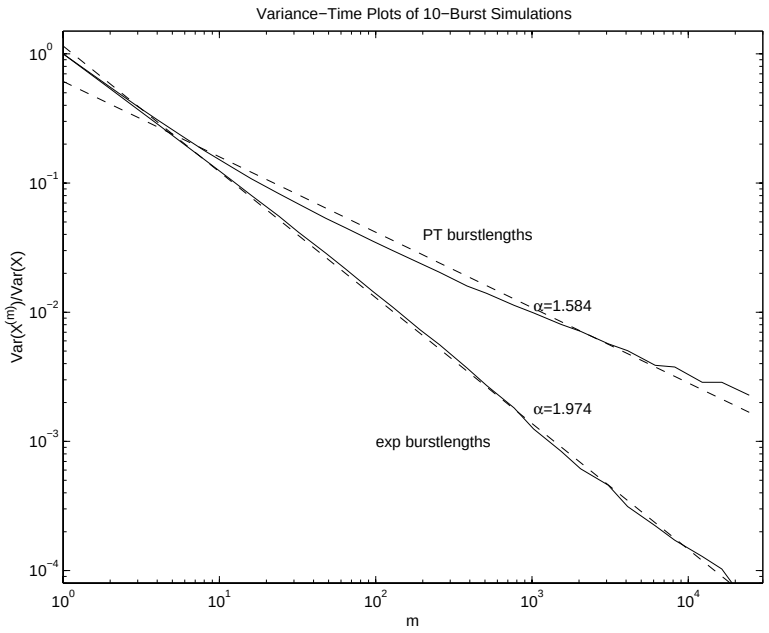

Figure 3.7: **Variance-Time Plots for Simulations of 10-Burst Processes:** The upper curve shows the results for PT-distributed ($\alpha = 1.5$) burst-lengths. It clearly indicates LRD properties, since the variance decays rather slowly. The other simulation uses exponential burst-lengths, therefore the variance-time plot yielded an estimate of α close to 2.

In Figure 3.7, the variance-time plot for simulated samples from a 10-Burst process with Power-Tailed ($\alpha = 1.5$) and exponential ON durations is shown. Aggregation levels of $m = 2^i$, and $m = 3 \cdot 2^i$, $i = 0, 1, ..$ are used in that plot. The aggregation is stopped when the aggregated stream had less than 20 samples (which is already a very thin basis for estimating $\mathrm{Var}(X^{(m)})$). The plot for the LRD model clearly shows a slower decay of the variance than in the exponential model. Note that the estimate $\hat{\alpha}$ is slightly too high. However, since the upper curve in Fig. 3.7 is slightly bent, the actual numerical value of the estimate for α changes slightly with the range of aggregation levels m that is used for the linear regression.

Chapter 4

Measurements and Model Calibration

The best traffic model is useless without a feasible (and reliable) method for parameter estimation. Such a method is developed in this chapter for the 1-Burst model. The method is then applied to a set of actual measured data from [Gogl 00]. The performance analysis of N-Burst/M/1 queues in the subsequent parameters is put in a context of a 'real' scenario by using a parameter-set which is close to the estimates in this chapter.

4.1 Description of Measurements

Figure 4.1: **The topology of the backbone of the German Scientific Network (B-WiN) at the time when the measurements were taken (source: DFN Verein)**: the point of measurement was at the entrance to the backbone in Munich, see next Figure.

Herein we use a set of measurements of inter-cell times that were obtained in [Gogl 00] at the entrance point to the IP-over-ATM backbone of the German Scientific Network (B-WiN) at the end of 1997. Figure 4.1 shows the topology of the backbone at that time. The data was taken at the access point in Munich, where the traffic from three large universities and several smaller institutions entered the backbone. Figure 4.2 shows the detailed structure of the part of the network at the measurement point. More technical details about the network configuration and about the measurement procedure can be found in [Gogl 00]. Also, a lot more analysis of the measured data can be found there. This thesis concentrates purely on the parameter estimation for the N-Burst model.

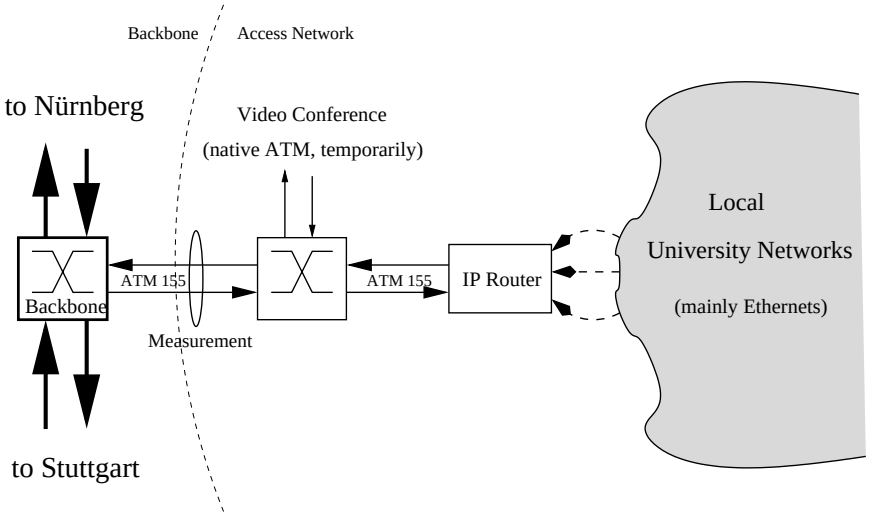

Figure 4.2: **The Point of Measurement:** Inter-cell times of ATM cells were captured at the transition from the access network in Munich to the backbone. See [GOGL 00] for more technical details.

4.2 Daily Profiles

In addition to the high-resolution measurements of individual inter-cell times, long-term measurements of the counting process for time-intervals $\Delta = 2s$ and more were performed in [GOGL 00]. Those long-term measurements clearly show that a regularly changing utilization profile can be observed during the course of a single day. Network traffic is not stationary over the course of the day, but the utilization is noticeably higher during the busy hours, which can be identified as the hours between 10am and 4pm on work days in the B-WiN measurements.

With respect to QoS, those busy hours are of particular interest. The performance analysis, and thus also the parameter estimation procedure, normally concentrates on this time period. The significance of this daily profile is that any stationary description of network traffic can only be valid for that limited time period. Therefore, steady-state performance analysis as presented in Chapter 5 has to be questioned for its practical relevance, if the time it takes for some performance parameter to converge to its steady-state value is of the order or larger than those busy hours. The latter analysis is done in the course of the transient analysis in Chapter 6. Furthermore, it would be meaningless to include features in the model that only become effective if the (stationary) lifetime of the network was much longer than those 6-8 daily busy hours. This argumentation also justifies the use of truncated tails for the burst duration: any burst longer than 8 hours cannot develop its full performance impact in any case, since it will partially fall outside the stationary time-period with high-load.

However, one has to keep in mind that eventually traffic from other parts of the world might travel through the networks, in which case the daily profile can be smoothed out somewhat. Yet at present, this cannot be expected within the next few years due to political and economic reasons.

4.3 1-Burst Calibration

The N-Burst process is fundamentally a more general approach for modeling self-similar traffic. So why bother with the calibration of the 1-Burst model? First, there are more robust methods for the parameter estimation of the 1-Burst process, and the results of the 1-Burst calibration can be a starting point for the general N-Burst process. Secondly, the measurements at the B-WiN entrance were done behind the Customer Service Switch (see Fig. 4.2) and therefore all the outgoing (or incoming) data was already sequentialized in a single stream, see also Sect. 10.4. It is definitely worth testing whether this stream can be adequately modeled by a 1-Burst process.

Unfortunately, it is not possible to recognize ON and OFF periods purely from the inter-cell times X_i with sufficiently high success rate. Using a threshold approach and assuming that an ON period is given as long as the inter-cell time does not exceed a certain threshold, results in a method that is highly dependent on the value for the threshold and even if a 'good' threshold value was chosen, short inter-burst times could be wrongly identified as inter-cell times within a burst or vice versa. Consequently, there is no hope do derive ON time durations from the measured X_i, but the calibration has to rely on more subtle approaches.

Hereafter, we present a calibration method that uses thresholds to filter out inter-cell times within a burst and also inter-burst OFF times, yet only likely candidates for both classes contribute to the estimators, such that it becomes statistically significant for large enough sample-sizes and proper choices of the thresholds (see Sects. 4.3.2 and 4.3.3). Having determined estimates for the cell-rates, the mean ON time is thereafter derived from known formulas for the 1-Burst model.

4.3.1 Marginal Distribution of Inter-Cell Times

The time X_i between two subsequent cells in the 1-Burst model can be assigned two one of the following three categories, each having a different marginal distribution:

1. X_i **Completely contained in a one ON period:** Both cells are generated during the *same* ON period. The cells are generated by a Poisson process with rate $\lambda_p + \lambda_0$, with the condition that the ON period does not end beforehand. When assuming the duration of the ON period as exponentially distributed with rate $1/\overline{x_p}$, the distribution of such an inter-cell time U has the reliability function:

$$R_U(t) = \exp\left(-\Lambda t\right), \quad \text{with} \quad \Lambda := \lambda_p + \lambda_0 + \frac{1}{\overline{x_p}} \cdot$$

   The expected value follows as:

$$\mathbb{E}\left\{U\right\} = \frac{1}{\Lambda} \cdot$$

2. X_i **Started during OFF-Period:** The first cell is generated during an OFF period, which is of course only possible if $\lambda_0 > 0$. The subsequent cell can either be generated in the same OFF period, in the next ON period, or there could be an arbitrary number of ON periods with zero cells in between. The distribution of such an inter-cell time, W, in case of exponentially distributed ON periods is expressed by the following ME distribution:

$$\mathbf{p}_W = [1,\ 0], \quad \mathbf{B}_W = \begin{bmatrix} \frac{1}{Z} + \lambda_0 & -\frac{1}{Z} \\ -\frac{1}{\overline{x_p}} & \frac{1}{\overline{x_p}} + \lambda_p + \lambda_0 \end{bmatrix} = \begin{bmatrix} \frac{1}{Z} + \lambda_0 & -\frac{1}{Z} \\ -\frac{1}{\overline{x_p}} & \Lambda \end{bmatrix} \cdot$$

Since the first cell is generated during the OFF period, at the beginning of the time-span W, the 1-Burst model is still in the OFF state, in $\mathbf{B}_W$ represented by the first state, thus $\mathbf{p}_W = [1, 0]$. From there, the next cell can be generated with rate λ_0 or a transition to the ON state with rate $1/Z$ can occur. Such an ON period can either end with rate $1/\overline{x_p}$ without generating any cell, or the second cell is generated with rate $\lambda_p + \lambda_0$, which ends the time-span W.

The expected value follows as

$$\mathbb{E}\{W\} = \mathbf{p}_W (\mathbf{B}_W)^{-1} \boldsymbol{\epsilon}' = \frac{1 + Z\,\Lambda}{(1 + \lambda_0 Z)\,\Lambda - 1/\overline{x_p}} \cdot \tag{4.1}$$

3. X_i **Started by Last Cell of ON-Period:** The first cell is generated during an ON period, which ends before the second cell is generated. Such an inter-cell time V is the sum of the residual time of the burst-length and the time W as discussed above:

$$\mathbf{p}_V = [1,\ 0,\ 0], \quad \mathbf{B}_V = \begin{bmatrix} \Lambda & -\Lambda & 0 \\ 0 & \frac{1}{Z} + \lambda_0 & -\frac{1}{Z} \\ 0 & -\frac{1}{\overline{x_p}} & \Lambda \end{bmatrix}.$$

The 2×2 sub-matrix on the bottom right is actually the matrix $\mathbf{B}_W$ of case 2. Since V is the sum of the residual ON-period duration and the random variable W, the expected value is simply:

$$\mathbb{E}\{V\} = \mathbb{E}\{U\} + \mathbb{E}\{W\} = \frac{1}{\Lambda} + \frac{1 + Z\,\Lambda}{(1 + \lambda_0 Z)\,\Lambda - 1/\overline{x_p}} \cdot \tag{4.2}$$

If $\lambda_p \gg \lambda_0$ and $\lambda_p \gg 1/Z$ (i.e. the cell-rate during ON dominates), then the inter-cell times U of type 1 generate most of the short inter-cell times, while V and W are mostly influenced by λ_0 or $1/Z$, which means they are more likely large.

Note that the assumption here is that ON periods as well as OFF periods have an exponentially distributed duration, which is of course in general not the case for the ON periods. However, since the calibration methods are based on expected values, the exponential assumption yields a reasonable approximation (see the left graphs in Figs. 4.3 and 4.4). An analysis which is based on the exact distribution of the ON periods would be very complicated, and that exact distribution would be hard to determine in practice.

All three types of inter-cell times collectively have the expected value

$$\mathbb{E}\{X\} = \frac{1}{\lambda} = \frac{Z + \overline{x_p}}{\overline{x_p}\Lambda - 1 + \lambda_0 Z}.$$

The overall fraction of cells during OFF periods is

$$p_W = \frac{Z}{Z + \overline{x_p}} \cdot \frac{\lambda_0}{\lambda} = \frac{\lambda_0\, Z}{\Lambda\, \overline{x_p} - 1 + \lambda_0\, Z}.$$

The fraction of cells that are generated right before the end of an ON time corresponds to the number of non-empty bursts divided by the overall number of cells in some large time-period t_0 (which cancels out)

$$p_V = \frac{1}{\lambda\,(Z + \overline{x_p})} \frac{\Lambda\overline{x_p} - 1}{\Lambda\overline{x_p}} = \frac{1 - \frac{1}{\Lambda\,\overline{x_p}}}{\Lambda\, \overline{x_p} - 1 + \lambda_0\, Z}.$$

All the remaining inter-cell times belong to class 1, thus

$$p_U = 1 - p_V - p_W = \frac{\overline{x_p}\Lambda - 2 + 1/(\Lambda\overline{x_p})}{\lambda\,(Z + \overline{x_p})} = \frac{(\Lambda\overline{x_p} - 1)^2}{\Lambda\,\overline{x_p}\;(\Lambda\,\overline{x_p} - 1 + \lambda_0\,Z)}\,.$$

The fractions p_W, p_V, and p_U also represent the probabilities that an arbitrarily picked inter-cell time X_i is of type W, V, and U, respectively. For validation, the weighted average of the expected values of the three classes of inter-cell times has to be the expected value of all inter-cell times:

$$p_V\,\mathbb{E}\{V\} + p_W\,\mathbb{E}\{W\} + p_U\,\mathbb{E}\{U\} = \mathbb{E}\{X\} = \frac{1}{\lambda}.$$

4.3.2 Threshold Estimate for the Intra-Burst Rate, λ_p

We saw in the last section that the inter-cell times fall into one of three classes with marginal distribution either f_U, or f_V, or f_W. In the usual case of $\lambda_p \gg \lambda_0$ and $\lambda_p \gg 1/Z$, short inter-cell times are most likely generated by the intra-burst class with density $f_U(t) = \exp(-\Lambda t)/\Lambda$. The calibration method makes use of that fact by selecting only the short inter-cell times via a threshold T_0:

$$(U_j^*) = \{X_i | X_i < T_0\}\,.$$

The probability that a sample of type U is generated below threshold T_0 is:

$$F_U(T_0) = 1 - \exp(-\Lambda T_0) = \Lambda T_0 + O(T_0^2),$$

where $g(x) = O(f(x))$ is the notation for $\lim_{x\to 0} g(x)/f(x) = c$ with $|c| > 0$ and $|c| < \infty$, i.e. '$g(x)$ decays as fast as $f(x)$ for $x \to 0$'.

Since the inter-cell times V that are created by the last cell of a burst, are in a broader sense of an Erlangian type (the sum of an exponential r.v. and another r.v.), the probability $F_U(T_0)$ decays more quickly for small T_0:

$$F_U(t_0) = o(T_0) \quad \text{for} \quad T_0 \to 0.$$

The notation $g(x) = o(f(x))$ is equivalent to $\lim_{x\to 0} g(x)/f(x) = 0$, i.e. '$g(x)$ decays faster than $f(x)$ for $x \to 0$'.

Consequently, for a small enough threshold T_0, the fraction of inter-cell types of Class 'V' in (U_j^*) will vanish more quickly than the samples of Class 'U'.

Unfortunately, the same does not hold for samples of Class 'W' if $\lambda_0 > 0$ (if $\lambda_0 = 0$, no such samples exist). Those inter-cell times can be generated 'instantaneously' by the exponential distribution with rate $\lambda_0 \ll \lambda_p$. Consequently,

$$\lim_{T_0\to 0} \frac{F_W(T_0)}{F_U(T_0)} = \frac{\lambda_0}{\Lambda},$$

which is small but still not 0. However, since their proportion in the number of samples in (U_j^*) converges to the (in case $\lambda_p \gg \lambda_0$) even smaller value

$$\frac{\lambda_0\,p_W}{\Lambda\,p_U} = \frac{Z\lambda_0^2\overline{x_p}}{(\Lambda\overline{x_p} - 1)^2},$$

we neglect the samples of Class 'W' in (U_j^*) in the following. With that assumption, for a small enough threshold T_0, $(U_j)^*$ almost surely contains only intra-burst times. The distribution of

$U^*(T_0)$ is then the exponential distribution with rate Λ which is cut off at the value of the threshold T_0:

$$F_{U^*(T_0)}(x) = \mathbb{P}\{U^*(T_0) \le x\} = \frac{1-\exp(-\Lambda x)}{1-\exp(-\Lambda T_0)} \quad \text{for} \quad 0 \le x \le T_0\,.$$

The mean of $U^*(T_0)$ then comes out as

$$\mathbb{E}\{U^*(T_0)\} = T_0 - \frac{T_0}{1-\exp(-\Lambda\, T_0)} + \frac{1}{\Lambda} = \frac{1}{\Lambda} - T_0\,\frac{1}{\exp(\Lambda T_0)-1} =: g_U(\Lambda, T_0)\ . \tag{4.3}$$

Note that $\lim_{\Lambda\to 0} g_U(\Lambda, T_0) = T_0/2$ and $\lim_{\Lambda\to\infty} g_U(\Lambda, T_0) = 0$. Furthermore, $g_U(\Lambda, T_0)$ is monotonically decreasing with Λ.

The measured inter-arrival times $X_i < T_0$ yield an estimate $\widehat{\mu_U}$ from the data. Equation (4.3) can then be solved numerically for an estimate $\widehat{\Lambda}$.

$\Lambda = \lambda_0 + \lambda_p + 1/\overline{x_p}$ still contains two unknown model parameters, λ_p and $\overline{x_p}$, even if we assume the background rate λ_0 to be known in advance. Thus, we have to continue the analysis before we can actually derive the first model parameter. In that sense, the headline of this section is promising too much at this time.

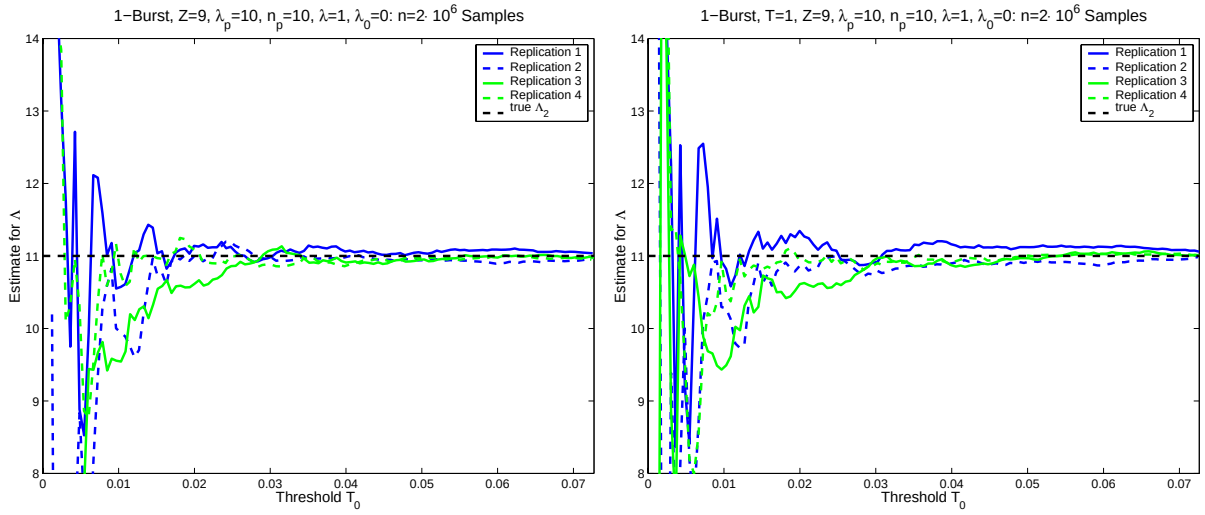

Figure 4.3: **Threshold estimate of Λ (here $\Lambda = \lambda_p + 1/\overline{x_p}$) for different simulations of the 1-Burst process with $n = 10^6$ samples:** If T_0 is too small, then the number of samples below the threshold is too low and the estimator is not reliable. On the other hand, if T_0 is too large, the danger of taking inter-burst times into account increases. For higher T_0 the curve actually drops down (not shown here). The left graph shows the estimators for Power-Tailed burst-length distributions, while the other experiment uses exponential burst-lengths.

Choice of the threshold T_0: The threshold T_0 should be as small as possible in order to exclude all inter-burst times. However, if T_0 is too small, then the number of samples U_j becomes too small and the estimate of $\widehat{\mu_U}$ becomes unreliable. A practical solution would be to compute $\widehat{\Lambda_2}(T_0)$ for a range of thresholds $\min\{X_j\} < T_0 < \text{median}\{X_j\}$, plot the graph, and look for the value of T_0 where the curve shows a horizontal plateau for the first time. Figure 4.3 shows such graphs for simulations of the 1-Burst process. The graph in Fig. 4.3 mainly concentrates on the area for small T_0. For larger values of T_0 the curve indeed drops down because more and more inter-burst times influence the estimator.

If the burst-length distribution is not exponential the distribution of U is much more complicated. However, since we only work with expected values for the estimator, the exponential case can be used as an approximation. In fact, Eq. (4.3) is used successfully in Fig. 4.3 although the burst-length distribution in the simulation is TPT distributed.

In Fig. 4.3, there is no background Poisson rate, so the estimator for Λ is unbiased for $T_0 \to 0$. There can be some distortion through the samples of class 'W', if λ_0/λ_p is not close to zero any more.

4.3.3 Threshold Estimate for Intra-Burst Time Z

A similar threshold approach is used to estimate the intra-burst time Z. This time, a lower bound T_1 is used to filter out the inter-cell times within a burst, leaving the samples

$$(V_j^*) = \{X_i | X_i > T_1\}\,.$$

For large T_1 and $\lambda_p \gg \lambda_0$ and $\lambda_p \gg 1/Z$, the samples in (V_j^*) will be exclusively of type 'W' and 'V', i.e. the fraction of intra-burst times in (V_j^*) goes to zero.

Both V and W are ME distributions, but not purely exponential. The expected value of V^* can of course be computed exactly, but no simple formulas for the calibration method result. Thus, as an approximation, we assume $\mathbb{E}\{V^*(T_1) - T_1\} = \mathbb{E}\{W\}$, as if those there exponential with mean $\mathbb{E}\{W\}$. Hence, we use the relationship

$$\mathbb{E}\{V^*(T_1)\} \approx T_1 + \frac{\overline{x_p} + \overline{x_p} Z \Lambda}{(\overline{x_p} + \overline{x_p}\lambda_0 Z)\,\Lambda - 1}\,. \tag{4.4}$$

Again, the sample mean of all inter-arrival times $X_i > T_1$ yields an estimate $\widehat{\mu_V}(T_1)$ for the expected value of V^*.

A third equation is still necessary to compute the three unknown parameters λ_p, Z, and $\overline{x_p}$. The expected value of all inter-arrival times $\mathbb{E}\{X\}$ with the sample mean $\widehat{\mu_X}$ as an estimator provides the missing equation:

$$\mathbb{E}\{X\} = \frac{Z + \overline{x_p}}{\lambda_p \overline{x_p} + \lambda_0\,(\overline{x_p} + Z)} = \frac{Z + \overline{x_p}}{\overline{x_p}\Lambda - 1 + \lambda_0 Z}\,. \tag{4.5}$$

For the threshold T_1 there is the upside down situation as for T_0. T_1 should be as high as possible since the higher T_1 the smaller the probability of having intra-burst times in the samples (V_j^*). On the other hand, the higher T_1 the smaller gets the number of samples for the estimate of $\widehat{\mu_V}(T_1)$. As a conclusion, the same procedure as for the Λ estimate should be applied: draw the curve $\widehat{\mu_V}(T_1) - T_1$ for $\text{median}\{X_j\} \ll T_1 < \max\{X_j\}$ and look for the last, stable plateau.

The actual estimators for the 1-Burst parameters $\widehat{Z}$ and $\widehat{\overline{x_p}}$ are derived by solving Eqs. (4.4) and (4.5) which leads to the following quadratic equation for $\widehat{Z}$,

$$\begin{aligned} \widehat{Z}^2 \quad & \widehat{\Lambda}\,(1 - \lambda_0 \widehat{\mu_X})\,[\lambda_0\,(\widehat{\mu_V} - T_1) - 1] \;+ \\ + \widehat{Z} \quad & \left[\widehat{\Lambda}\,(\widehat{\mu_V} - T_1) - \widehat{\mu_X}\,(\widehat{\Lambda} - \lambda_0) - 1\right] \;+\; \widehat{\mu_V} - T_1 - \mu_x \;=\; 0, \end{aligned} \tag{4.6}$$

whose positive solution provides the estimate $\widehat{Z}$. Thereafter, $\widehat{\overline{x_p}}$ and $\widehat{\lambda_p}$ are computed as:

$$\widehat{\overline{x_p}} = \frac{\widehat{Z} + \widehat{\mu_X}\,\left(1 - \lambda_0 \widehat{Z}\right)}{\widehat{\mu_X}\widehat{\Lambda} - 1} \quad \text{and} \quad \widehat{\lambda_p} = \widehat{\Lambda} - \lambda_0 - \frac{1}{\widehat{\overline{x_p}}}\,. \tag{4.7}$$

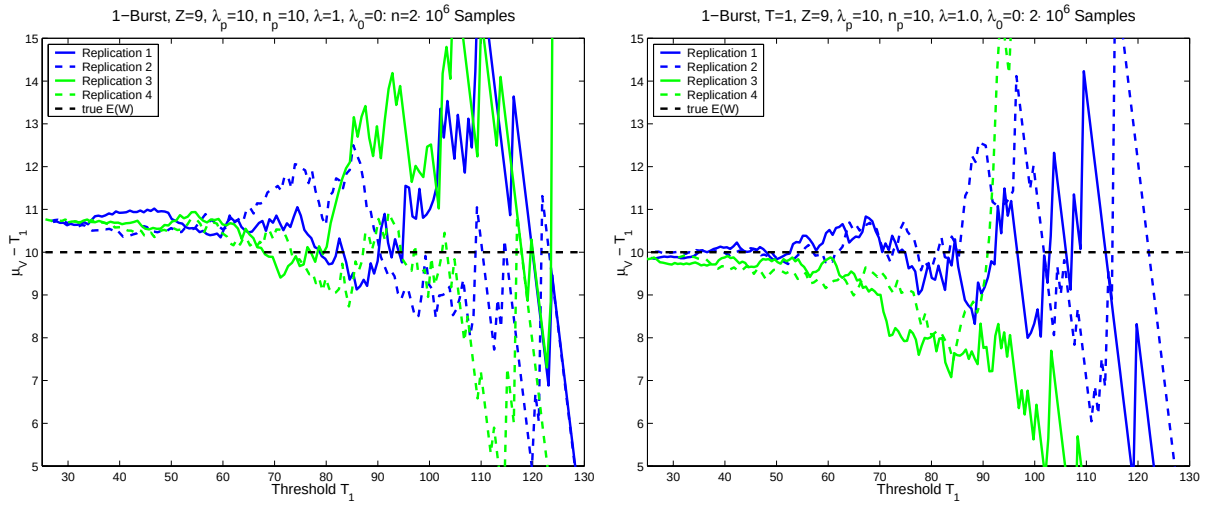

Figure 4.4: **Estimator $\widehat{\mu_V} - T_1$ for the Samples of 1-Burst Simulations:** The dashed line shows the value of the approximation in Eq. (4.4). For the exponentially distributed burst-lengths in the right graph, the estimator is in agreement with the true value $\mathbb{E}\{W\}$. However, for PT distributed burst-lengths (left graph), the estimator is slightly too high.

4.3.4 Determining the Background Poisson Rate λ_0

So far, we assumed that we know the value of the parameter λ_0 beforehand. In practice, this is – of course – rarely the case. There are two fundamentally different approaches, how the parameter λ_0 can be estimated:

- Additional information about the measured traffic can be obtained: For example, the cells can be assigned to certain applications or network management information, which are known to be adequately described by Poisson traffic. In that case, the background Poisson rate λ_0 is determined beforehand by a separate analysis of cells of such classes only.

- Additional properties of the inter-cell times X_i are used to determine an estimator $\widehat{\lambda_0}$. Figure 4.5 shows two potential candidates. As we see in the left graph, the coefficients of correlation of 1-Burst processes that are calibrated to simulation runs reduce when the apriori background Poisson rate λ_0 is increased. The estimate $\widehat{\lambda_0}$ is then obtained from the intersection with the estimated coefficient of variation of the data. However, for the individual replications, the estimate $\widehat{\lambda_0}$ shows a range starting at about 0.05 up to 0.13. In Replication 1, the desired value of $C^2(X) = 9$ is not even reached at all.

  The short-range correlation, $r(1)$, provides a more stable estimator for $\widehat{\lambda_0}$: As the right graph of Fig. 4.5 shows, a larger λ_0 increases the short-range correlation substantially. The intersection with the estimate $\hat{r}(1)$ from the samples X_i occurs very close to the true model parameter $\lambda_0 = 0.1$ of the simulation.

In summary, either additional information beyond the pure inter-cell times X_i need to be obtained for the calibration of the back-ground Poisson rate λ_0 before the remaining parameters are determined. If this is not possible, the estimators $\widehat{Z}(\lambda_0)$, $\widehat{\overline{x_p}}(\lambda_0)$, $\widehat{\lambda_p}(\lambda_0)$ are computed for a range of λ_0 with $0 \leq \lambda_0 < \Lambda$. The short-range autocorrelation, $r(1)$, is then computed for

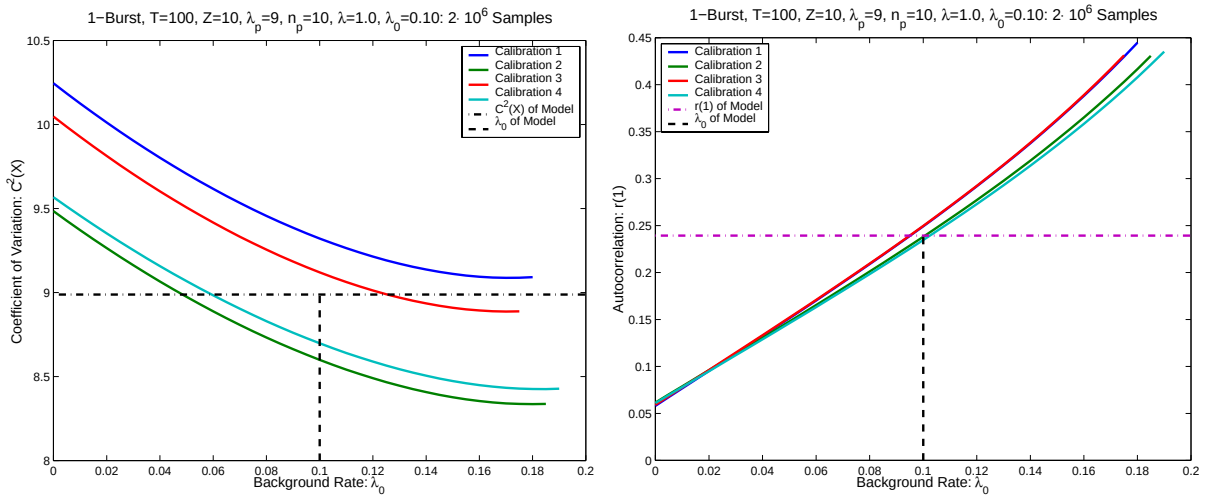

Figure 4.5: **The Coefficient of Autocorrelation and the Short-Range Correlation for different Background Rates:** Note that $\widehat{\Lambda}$, $\widehat{\overline{x_p}}$, and $\widehat{Z}$ are fixed. With varying λ_0 all other N-Burst parameters change according to (4.6) and (4.7).

the analytic 1-Burst model with those parameter-sets. The closest fit of $r(1)$ to the estimated correlation from the samples X_i then yields the estimator $\widehat{\lambda_0}$.

However, note that the short-range correlation depends strongly on the 'body' of the burst-length distribution. With the estimators for the LRD properties in the next section, the tail-exponent of the burst-length distribution is determined – in case it turns out to show LRD properties. However, we always assume the TPT distribution as the underlying Power-Tail distribution, which in practice could be different. Since the duration of the individual ON periods cannot be easily obtained from the inter-cell times X_i, histogram methods cannot be applied to obtain a description of the body of the burst-length distributions.

4.3.5 Application to Measurement Data

Figure 4.6 shows the resulting curves of the estimators $\widehat{\Lambda}$ and $\widehat{\mu_V} - T_1$ for two real measurements while varying the thresholds T_0 and T_1. Two different sets of measurements are discussed in this section: `TX3` captures the inter-cell times of traffic that was sent into the backbone at the access node in Munich. `RX3` on the other hand presents a measurement of ATM cells going the opposite direction, i.e. coming from the backbone into the Munich network. The measured inter-cell times X_i had the following properties:

Measurement	#samples	$\max(X_i)$	$\mathbb{E}\{X\} = \kappa^{-1}$	$C^2(X)$	$r(1)$
`TX3`	784 652	27 ms	61.4 μs	13.7	0.129
`RX3`	784 650	7.2 ms	51.6 μs	22.3	0.0684

The estimators $\hat{\Lambda}(T_0)$ and $\hat{\mu_V}(T_1) - T_1$ clearly show plateaus in Figure 4.6 for both measurements. Therefore the estimates $\hat{\Lambda}$ and $\hat{\mu_V}$ can be determined. From fitting the short-range autocorrelation $r(1)$, the following values for the 1-Burst parameters were obtained:

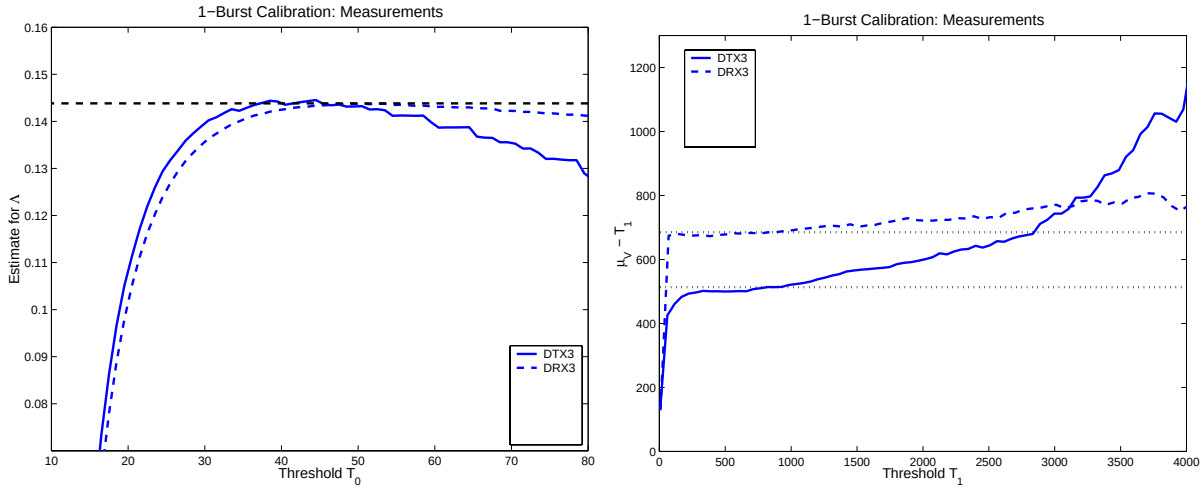

Figure 4.6: **The Threshold Estimates for Measurements `TX3` and `RX3`:** The thresholds that are used for the calibration are $T_0 = 44\,\mu$s and $T_1 = 840\,\mu$s.

Measurement	λ_p [cells/ms]	λ_0 [cells/ms]	$\overline{x_p}$ [μs]	Z [μs]	$\overline{n_p} = \lambda_p \overline{x_p}$	$b = 1 - \kappa/\lambda_p$
`TX3`	132	0.444	82.0	600	10.8	0.88
`RX3`	135	0.142	118	706	15.9	0.86

Note that the intra-burst cell-rate λ_p is almost a factor of 300 (respectively 950) larger than the background Poisson rate λ_0. Also the peak cell-rate is almost the same for the 'Transmit' and 'Receive' direction. Finally, the burstiness parameter b is very high, which is somewhat surprising since the traffic in the measurements is highly aggregated (in the order of 100 to 300 simultaneous IP connections contribute their ATM cells), and one normally expects that higher aggregation reduces the burstiness. Very likely, the burstiness of the aggregated traffic is smaller than for individual source destination pairs, but an analysis on that level is not performed here. Some information about the traffic on application level can be found in [Gogl 00].

See Appendix I.2.2 for a table with results for additional measurements of the same kind.

4.4 Long-Range Dependence and the Tail-Exponent α

So far we have introduced a method to estimate the parameters λ_p, $\overline{x_p}$, Z, and λ_0 for a 1-Burst model from measurements of inter-cell times X_i. The parameters $\overline{n_p} = \lambda_p \overline{x_p}$, $\kappa = \overline{n_p}/(\overline{x_p} + Z)$, and $b = 1 - \kappa/\lambda_p = Z/(Z + \overline{x_p})$ can be derived from the former set of estimated parameters. Still missing is the actual distribution of the ON and OFF times (in addition to their expected values $\overline{x_p}$ and Z). We show in Section 5.7 that the actual distribution of the OFF times is not so critical for performance – by far less critical than the ON time distribution at least. So we just assume the OFF time to be exponential and do not have to worry about it further – the OFF distribution is completely described by its mean Z.

The situation for the distribution of the ON periods is quite different: we will see later that its impact on performance can be large, in particular when PT distributions are involved. The distribution of the ON periods in the 1-Burst model also has a large impact on the correlation

structure of the resulting inter-cell time process: as we showed in Section 3.6, Power-Tail distributed ON times with tail-exponent α cause a slow Power-Law decay of the autocorrelation function, in fact with exponent $\alpha - 1$. If α is small enough, $1 < \alpha < 2$, LRD properties are caused in both inter-cell times and counts. As a consequence, the tests for LRD properties that are briefly described in Section 3.7.2 can be used to find out whether PT (or TPT) distributed ON periods need to be used for the ON periods, and to estimate the tail-exponent α.

In addition to the tail-exponent α, the description of the Power-Tail also requires a tail-constant c_{PT}. For simplicity, we exclusively use the TPT distributions that are described in Section 3.4. Although those TPT distributions could easily be modified to mimic any tail-constant, in the basic version used here, the tail-constant of the TPT distribution is fixed for given α, see Appendix I.1 for a table with numerical values. Thus we do not attempt to match the tail-constant of the ON time distribution in the data.

For truncated tails, additional methods to obtain the parameter MBS (or equivalently the PT-Range $\overline{x_T}$ of the ON time distribution) would have to be developed. We omit the parameter estimation for MBS here completely. Instead, we discuss the scenarios of how a changing value of MBS affects performance. From the truncated behavior of e.g. the correlation function of the measured inter-cell times, the order of magnitude of the MBS in the data can be estimated. However, this estimation requires a huge amount of measured data if that truncation is far out in the tail.

Note that the straightforward approach of applying statistical tests (e.g. for Power-Tails as in Sect. 3.7.1) on the individual ON durations is not possible here: in our scenario, there is no feasible approach to detect the individual ON periods from the measured inter-cell times X_i. As mentioned earlier, any kind of threshold approach to distinguish ON and OFF periods is too unreliable, since the results depend strongly on the actual value of the threshold.

In the following, we apply the methods for LRD estimation of Section 3.7.2 on the measurements `TX3` and `RX3`. Furthermore, the results are compared with a trace of measured Ethernet-data that was already object of several papers (e.g. [Leland et al. 94] and [Willinger et al. 95]) and is therefore well known.

Autocorrelation Plot

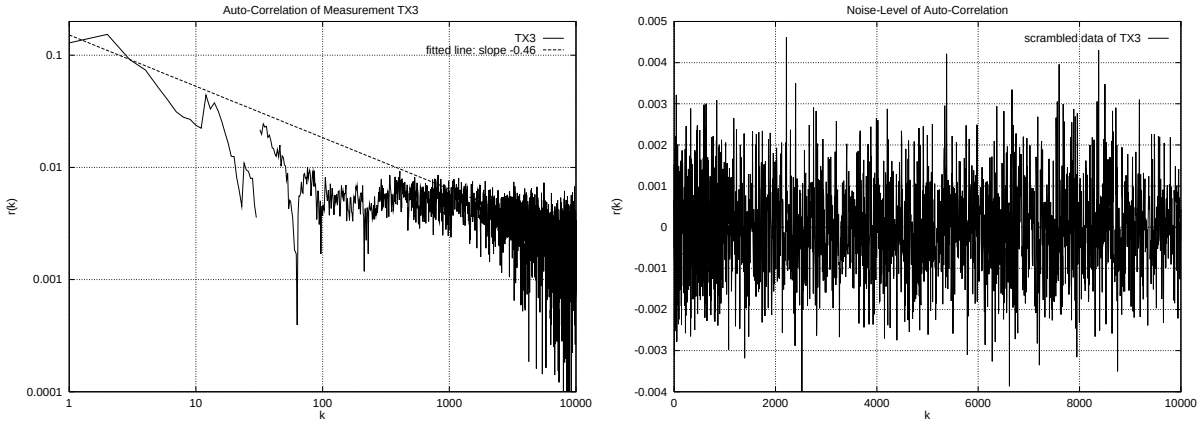

Figure 4.7: **Log-Log Plot of the Autocorrelation in Measurement `TX3`:** LRD properties can be observed in the left graph. The fitted line has a slope of -0.46, i.e. $\hat{\alpha} = 1.46$. The right-hand graph shows the $r(k)$-curve for the scrambled data-set, i.e. the correlation is destroyed; what is left is all noise.

As pointed out in Section 3.7.2, the first estimate for the PT exponent α can be obtained from the autocorrelation of the inter-cell times. Figure 4.7 shows the resulting $r(k)$-curve for measurement `TX3`. The line-fit resulted in an estimated tail-exponent $\hat{\alpha} = 1.46$ for the plotted curve (the value of $\log k$ were not quite regularly spaced, therefore the larger k attract the line slightly stronger). The results for `RX3` and the Ethernet data are given in the table at the end of this section.

In other measurements of the same type, the autocorrelation function was disturbed by the existence of periodic video traffic that was superimposed on the normal (IP over ATM) data traffic on a separate ATM VC, see Fig. 4.2. The resulting autocorrelation function of the measured inter-cell times in this case is shown and discussed in [Schwefel et al. 97]. We just want to mention here that such superimposed traffic of periodic nature introduced an oscillating behavior of the autocorrelation. The test for LRD via the autocorrelation function cannot be used in this case due to the negative parts of the autocorrelation function. Furthermore, the following LRD tests seemingly showed no LRD properties. However, if the video traffic was filtered out, LRD properties comparable to `TX3` were clearly indicated by the same tests for the remaining traffic. Thus, superimposed traffic, in particular if it has periodic nature, can lead to misleading conclusions, since it can obscure existing LRD properties of the non-periodic traffic. For the purpose of demonstrating the impact of such periodic traffic, we use one sample-set `TX1` that includes such video traffic. Otherwise we only consider measurements that were taken during time-periods when no such video traffic was present.

R/S-Statistics

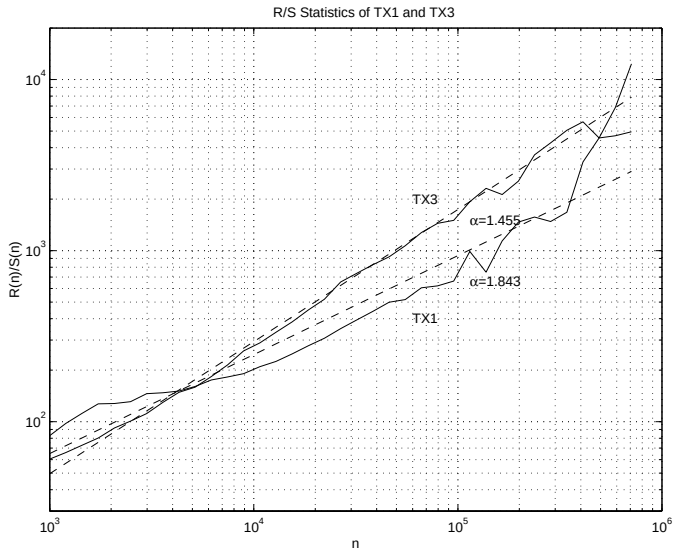

Figure 4.8: **R/S-statistics for measurement `TX3` and `TX1`:** `TX1` contained ATM cells of periodic video traffic in addition to the 'normal' data traffic. The periodic traffic in `TX1` reduces the range $R(n)$ of the data and therefore results in a higher estimate for α.

The results of the application of the R/S statistics (see Sect. 3.7.2) on the measurements `TX3` and `TX1` are shown in Figure 4.8. The R/S statistics for `TX3` clearly appear linearly on log-log scale with a slope that is noticeably larger than 0.5, i.e. LRD properties are indicated in that measurement. The estimate $\hat{\alpha} = 1.46$ is in good agreement with the estimated α from the correlation method.

For the measurement `TX1` that includes the periodic video-traffic, the R/S statistics generate a somewhat bent curve, and the resulting estimate for α is close to the non-LRD case $\alpha = 2$. Without the cells from the video traffic, the result of the R/S statistics are comparable in looks and numerical value of $\hat{\alpha}$ with the `TX3` measurement.

After destroying the autocorrelation of the measured inter-cell time by scrambling, the R/S statistics provide the value $\hat{\alpha} = 1.98$ for `TX3`. So clearly, the strong growth of $R(n)/S(n)$ for the unscrambled measurement is due to the order of the inter-cell times.

Variance-Time Plots

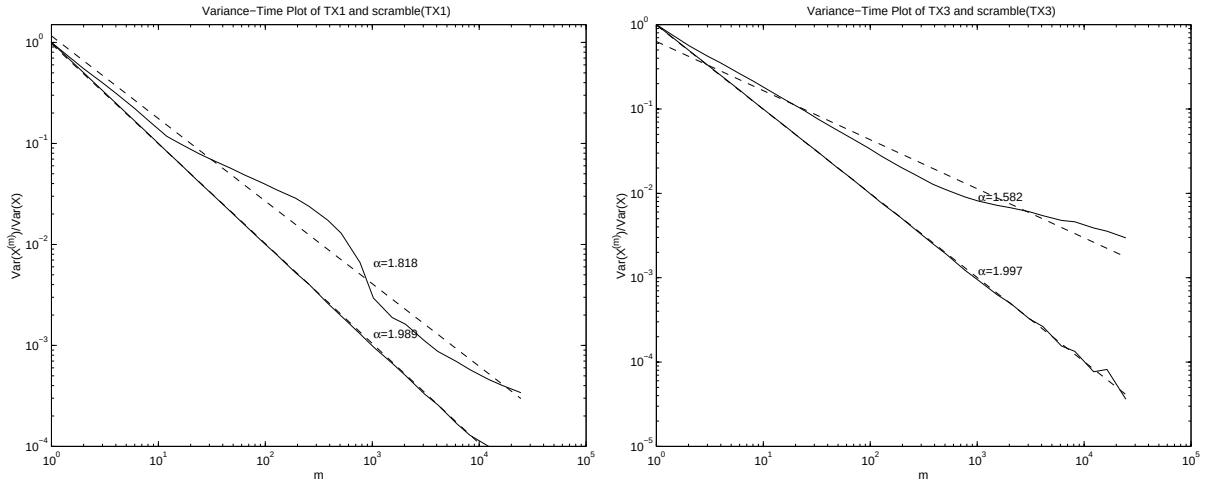

Figure 4.9: **Variance-time Plots for Measurements `TX1` (left) and `TX3` (right):** in both graphs, the curves are shown for the original inter-cell time stream and the scrambled stream. The scrambled stream exhibits a straight line with slope close to -1, i.e. no Long-Range Dependent properties.

Finally, Figure 4.9 shows the results of the application of the last LRD test of Section 3.7.2, the Variance-Time plot. The Variance of the aggregated inter-cell times decreases more slowly with increasing aggregation level than it would be expected in traditional, non LRD data, e.g. as for the scrambled measurements. Again, the periodic video traffic in `TX1` causes larger values for the α estimate.

Summary of Results

The following table summarizes the results of the different estimators for α for three sets of measurements (all without periodic video traffic):

	`TX3`	`RX3`	Ethernet
$\hat{\alpha}$ by $r(k)$	1.37	1.54	1.46
$\hat{\alpha}$ by R/S	1.46	1.51	1.31
$\hat{\alpha}$ by $\mathrm{Var}(X^{(m)})$	1.58	1.76	1.42

The results for `TX3` and `RX3` as well as the Ethernet data clearly indicate an $\alpha < 2$, i.e. infinite variance PT distributions need to be used for the ON periods in the N-Burst model. Other

measurements on application level, e.g. file sizes ([Gribble et al. 98]) or sizes of http transfers [Crovella & Bestavros 96] indicate Power-Tail distributions with comparable exponents.

Note that the linear regression to the logarithmic autocorrelation function can also detect Power-Tails with exponents $\alpha > 2$. Since such larger tail-exponents do not cause LRD properties, the other methods cannot distinguish such PTs with high exponents from plain exponential or other well-behaved tails.

4.5 Remarks on N-Burst Calibration

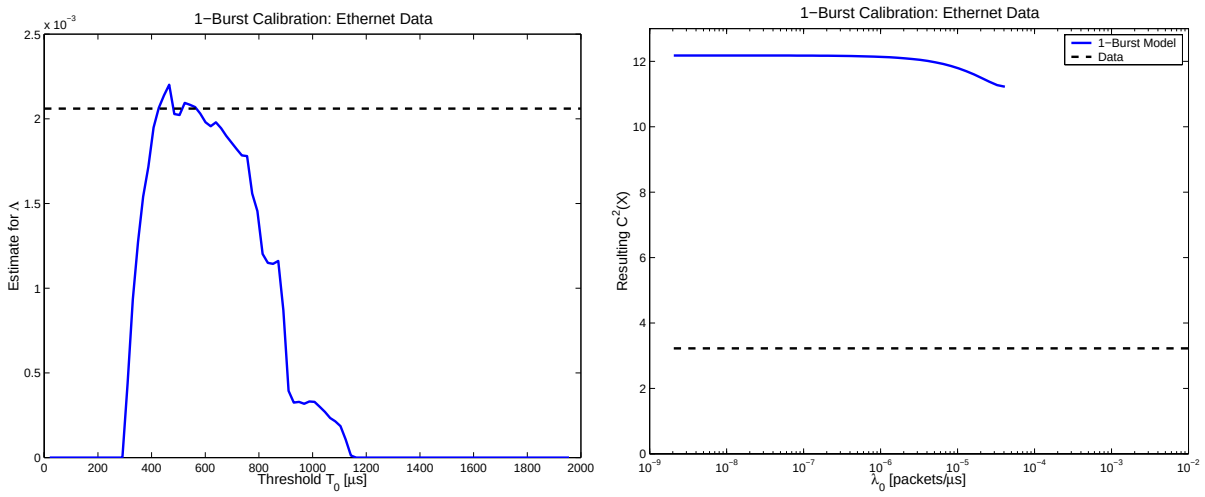

Figure 4.10: **The failure of the 1-Burst calibration method for the Ethernet data:** The estimate $\hat{\Lambda}(T_0)$ does not show a clear plateau in the left graph. Also, the coefficient of variation of the inter-cell times of the calibrated 1-Burst models is far too high compared to the measured data. That coefficient of variation cannot be matched by an increase of the background Poisson rate λ_0.

In principle, any data can be used to calibrate a 1-Burst model with the method that is developed in Sect. 4.3. It is desirable to have some indication, whether the 1-Burst model is really the right type of model for the given data set. In fact such indicators for a mismatch of the model type can be obtained during the calibration procedure. Figure 4.10 shows two of them for the Ethernet data of [Leland et al. 94]:

- The coefficient of variation in the data is by far lower than the one that can be achieved with the 1-Burst calibration, even if a background Poisson rate $\lambda_0 > 0$ is taken into account (right graph of Figure 4.10).
- The estimator $\widehat{\Lambda}(T_0)$ does not show any clear plateaus, but instead it decays rapidly (sometimes stepwise as in the left graph of Figure 4.10).

Consequently, it seems that the 1-Burst model is the wrong model type for the Ethernet data.

In such cases, an N-Burst model with $N > 1$ could be more appropriate. For higher N, the calibration becomes much harder: First, it has to be determined what kind of burst-start process is appropriate, e.g. the Independent Source model or the $[\mathrm{M/G}/N//N]$ model that describes the traffic for a huge ($\gg N$) number of sources, see Chapter 2. Secondly, the threshold approaches that worked for the 1-Burst model cannot be applied any more. A different methodology for parameter estimation has to be developed.

All the measurements of ATM cells at the B-WiN entrance could be described well by calibrated 1-Burst models. Also, the development of robust N-Burst calibration methods would be beyond the scope of this thesis. Thus, no N-Burst calibration methods are discussed here. However, a very heuristic approach for the $[M/G/N//N]$ model with large N is described in [Schwefel et al. 97]: That method uses the variance of cell-rates in sliding sample-windows of some size to determine the average number of simultaneously active bursts. Afterwards, the mean and the variance of the whole measured cell-stream is used to determine the N-Burst parameters $\overline{n_p}$ and λ_p. See [Schwefel et al. 97] for details (however note, that the parameter names are different there!).

Another possibility would be to assign the single cells to the corresponding source-level transfer, e.g. mark all the ATM-cells of one specific `FTP` data transfer. By doing that the cell-stream is separated at the source-level at which a 1-Burst calibration can be tried.

4.6 Summary

A method for estimating the parameters of the 1-Burst model from measured inter-cell times was developed in this chapter. Thresholds are used to separate the measured inter-cell times in three subsets: Small, intermediate, and large inter-cell times. In principle, the small inter-cell times are used to estimate the intra-burst cell-rate λ_p, while the large ones provide an estimate for the mean duration Z of the OFF periods. However, inter-cell times from the 1-Burst model can be partitioned in three classes, each of them with different marginal distributions. A closer look at those marginal distributions showed that the 1-Burst parameters λ_p, Z, and $\overline{x_p}$ cannot be obtained separately, but a more complicated procedure is necessary, which is described in Section 4.3.

That calibration method is successfully applied to several measurements of ATM inter-cell times taken at the entrance to a backbone that connects the German Universities and other research institutes. As part of the calibration procedure, it is necessary to find out the type of distribution that describes the duration of the ON periods in the 1-Burst model. Three different tests for LRD properties indicate that such properties are present in the measured data, and thus the conclusion is that Power-Tail distributions are a more appropriate choice for the duration of the ON periods. Estimates for the tail-exponent α that are obtained from the statistical tests fall in the range $1.4 \leq \alpha \leq 1.6$.

Chapter 5

Steady-State Performance Analysis

In Chapter 2 we set up the N-Burst models and then discussed in Chapter 3 how to include LRD properties in such models. We now start to investigate their performance behavior. This chapter discusses a set of *steady-state* performance parameters for N-Burst/M/1[/B] queues: the expected value mCD and higher moments of the cell-delay distribution, the Buffer-Overflow Probability BOP(B) for a primary buffer of B cells in the infinite-buffer N-Burst/M/1 model, and finally the Cell-Loss Probability CLP(B) in the finite buffer N-Burst/M/1/B loss model. These performance parameters are derived from the steady-state queue-length distribution at cell-arrival times $\mathbb{P}\left\{Q^{(a)} = k\right\}$.

Telecommunications systems have a finite lifetime. Not only is the physical system subject to changes after a few months, e.g. in order to adjust to a higher bandwidth demand, or to include new technology in the network. Even more, a stationary traffic model can only be considered to be realistic during the daily[1] busy hours of the network utilization profile, see Sect. 4.2. Thus, the actual 'stationary lifetime' is limited to a few hours.

Steady-state performance analysis on the other hand assumes that the observation period is long enough, so that the observed distribution of some performance parameter (here queue-length or per-cell delay) has converged to an equilibrium distribution. As we will see later, such convergence within the limited daily busy hours can be doubtful for network traffic with LRD properties. Such criticism of steady-state analysis is the major motivation to perform transient analysis, see Chapter 6. Nevertheless, steady-state analysis provides a first understanding of the performance behavior, and its results can be used as basis for comparison of the transient performance results. Furthermore, in the particular case of N-Burst/M/1 models with LRD properties, it turns out that one of the most important qualitative results, the existence of so-called *blow-up points* (see Sect. 5.1), is a phenomenon common to both steady-state and transient performance analysis.

The steady-state analysis in this chapter starts off by discussing the *blow-up effects* and deriving their location in the parameter-space of the N-Burst/M/1 queue in Sect. 5.1. The investigation of the queue-length distribution at cell-arrivals in Sect. 5.2 provides a full understanding of why those blow-ups occur: except for rather uninteresting cases (called Peak-Rate Allocation), the queue-length distribution turns out to be Power-Tailed, but its tail-exponent changes discontinuously at the blow-up points. The impact of truncated Power-Tail distributions with finite PT Range is discussed in Sect. 5.3: Truncated tails in the burst-length distributions cause a truncation of the PTs of the queue-length distribution. Formulas for the PT-Range of the latter are obtained.

Section 5.4 discusses some practical consequences of the (truncated) Power-Tailed queue-length distribution, one of them is the behavior of BOP(B). Several model variations are discussed in the remaining sections of this chapter: the finite-buffer loss-model in Sect. 5.5; the relation of M/M/1, bulk-arrival, and fluid-flow model to the N-Burst model in Sect. 5.6; N-Burst/M/1

[1] There certainly exist other stationary time-intervals (e.g. during certain night hours), but due to the lower utilization, they are less relevant for network planning than the busy hours.

models with LRD properties caused by PT distributed OFF periods in Sect. 5.7; and finally, N-Burst models with different types of bursts as introduced in Sect. 2.5.2.

The numerical illustrations in this chapter use N-Burst [IS] models whose individual source parameters (b, $\overline{n_p}$, κ) are chosen close to the results of the 1-Burst calibration in Sect. 4.3: $b = 0.9$, $\overline{n_p} = 10$. For better readability, the average cell-rate is (almost) always normalized to $\lambda = 1$. It follows for the 2-Burst model with $b = 0.9$, $\overline{n_p} = 10$, $\lambda = 1$, which is frequently used for illustration purposes in this chapter: $\kappa = 0.5$, $\lambda_p = 5$, and $Z = 18$.

5.1 Blow-up Points

Our first approach to the steady-state analysis of N-Burst/M/1 queues is a one-factor analysis of a very relevant scenario: what happens to the traffic at the bottleneck component if the utilization $\rho = \lambda/\nu$ is varied by experimenting with the service-rate[2] ν?

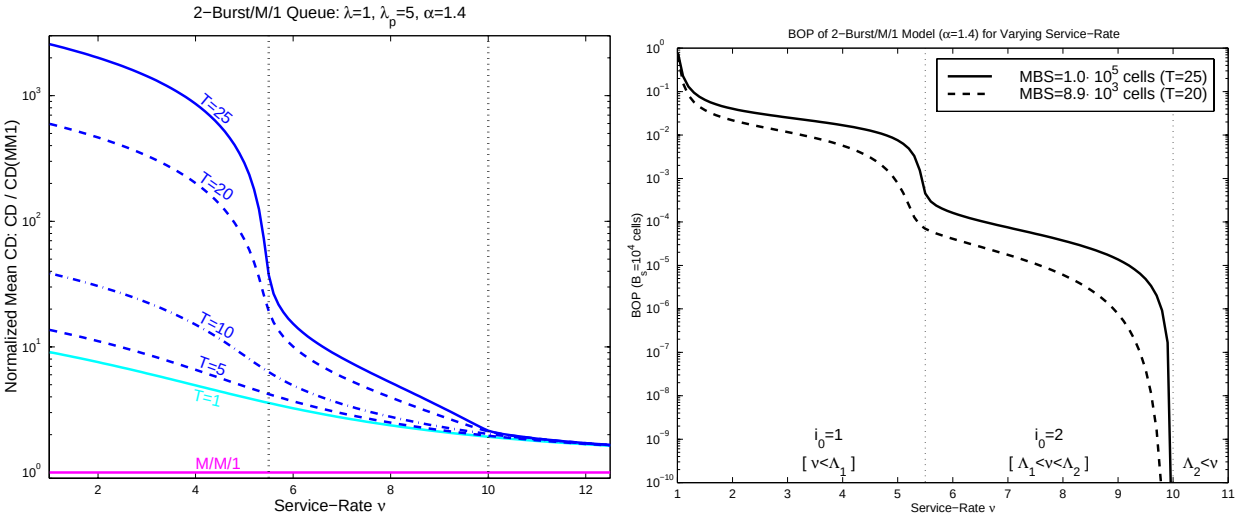

Figure 5.1: **Blow-up Regions in the Parameter Space of the 2-Burst/M/1 Queue:** The QoS dramatically decays when the service-rate, ν, becomes smaller than the blow-up points, Λ_i. In the left graph, the mean cell-delay (or *packet* delay, depending on the protocol) of a 2-Burst/M/1 queue normalized with respect to the mean delay, $1/(\nu - \lambda)$, of an M/M/1 queue is shown. The graph on the right shows the Buffer-Overflow Probability for a buffer of $B = 10^4$ cells. Note that the blow-up points occur at relatively low utilization values, here $\rho_1 = \lambda/\Lambda_1 = 18.2\%$ and $\rho_2 = \lambda/\Lambda_2 = 10\%$, see text.

The performance as measured by mean Cell Delay mCD or Buffer-Overflow Probabilities BOP becomes worse as utilization increases. For the M/M/1 queue, the mCD increases as $\rho/(1-\rho)$, and the BOP increases as ρ^B with increasing utilization ρ. As Figure 5.1 shows in its left graph, although the mCD of a 2-Burst model with exponential burst-lengths ($T = 1$) shows somewhat higher delay values than the M/M/1 queues, the qualitative behavior for increasing ρ (decreasing ν) is not very surprising: the mCD becomes gradually worse.

That observation does not hold any more if LRD properties are introduced into the traffic by the use of TPT distributions. In that LRD scenario for larger T in Fig. 5.1, both the mCD and the

[2] In practice, such a variation is done virtually during the network design phase. In Virtual Private Networks, it can also happen during the operation of the network through the assignment of more physical capacity to the virtual network.

BOP do not increase uniformly with decreasing service-rate any more. Instead, peculiar jumps, called *blow-ups*, can be observed.

The reason for that peculiar behavior is that (even truncated) Power-Tail distributions allow for very long bursts with non-negligible probability. If i of the N sources are in such a 'long-term' active period, then only $(N-i)$ sources are left with ON/OFF behavior during that time-period. Therefore, the model behaves temporarily as an N'-Burst where $N' = N - i$ but with increased background Poisson rate, $\lambda_0' = \lambda_0 + i\lambda_p$. As a consequence, the average cell arrival rate is temporarily raised to

$$\Lambda_i := \lambda_0 + i \cdot \lambda_p + (N-i) \cdot \kappa = \lambda + i \frac{b}{1-b} \kappa, \quad i = 1, ..., N. \tag{5.1}$$

When $\nu < \Lambda_i$, the i long-term active sources lead to an over-saturation period of the switch, i.e. temporarily the mean cell arrival rate exceeds the service rate ($\rho' > 1$ in the transient N'-Burst/M/1 queue). Clearly, substantial growth of the queue-length can result if the over-saturation period is long enough. Hence a potential for overflow situations is also given, even for large buffers.

In the following we assume that the Poisson background cell rate of the N-Burst model is $\lambda_0 = 0$. A positive $\lambda_0 > 0$ can be subtracted from the service rate, $\nu \to \nu - \lambda_0$, for the following discussion.

Since the number of active sources is in the set $i \in \{1, 2, ..., N\}$, it follows that there are N blow-up points that are located by the condition $\nu = \Lambda_i$. The region of the parameter-space for which

$$\Lambda_{i_0-1} < \nu < \Lambda_{i_0} \quad \Longrightarrow \quad i_0 = \left\lceil N \cdot \frac{1-\rho}{\rho} \cdot \frac{1-b}{b} \right\rceil \quad \text{for} \quad \lambda_0 = 0\,, \tag{5.2}$$

is called *blow-up region* i_0, $i_0 \in \{1, ..., N\}$. Within such a blow-up region, i_0 long-term active sources together with the remaining $(N - i_0)$ sources in average behavior are sufficient to produce an over-saturation period, while $i_0 - 1$ long-term active sources do not over-saturate the server. Furthermore, the parameter i_Δ is introduced,

$$i_\Delta := i_0 - N \cdot \frac{1-\rho}{\rho} \cdot \frac{1-b}{b}\,, \tag{5.3}$$

which has a range $0 \le i_\Delta < 1$. The larger i_Δ the closer does the queueing model operate to the blow-up region $i_0 - 1$, where performance is worse.

The blow-up effects at the so-called *blow-up points*, where $\nu = \Lambda_i$, are peculiar to the LRD variant of the N-Burst model[3] Traditional exponential burst-length distributions do not show comparable effects, see Fig. 5.1. Consequently, the number of the blow-up region i_0 is very critical for both delay and overflow probabilities in N-Burst/M/1 queues with Power-tailed burst-length distributions. The next section provides more understanding of the impact of i_0 on steady-state performance.

Beyond the last blow-up point, $\nu > \Lambda_N = N\lambda_p$, or equivalently in the [IS] model, $\rho < 1 - b$, over-saturation of the switch becomes highly improbable. The mCD is rather small and the actual burst-length distribution does not have visible impact, i.e. LRD by Power-Tailed ON-times does not lead to very different QoS than standard exponential burst-lengths. That region of the parameter space is referred to as *Peak-Rate Allocation* (PRA) herein.

[3] The blow-up effects itself can also be observed for other burst-length distributions with high variance, see Sect. 5.8. However, the properties that are discussed in the subsequent sections only hold for PT (or TPT) distributed burst-lengths.

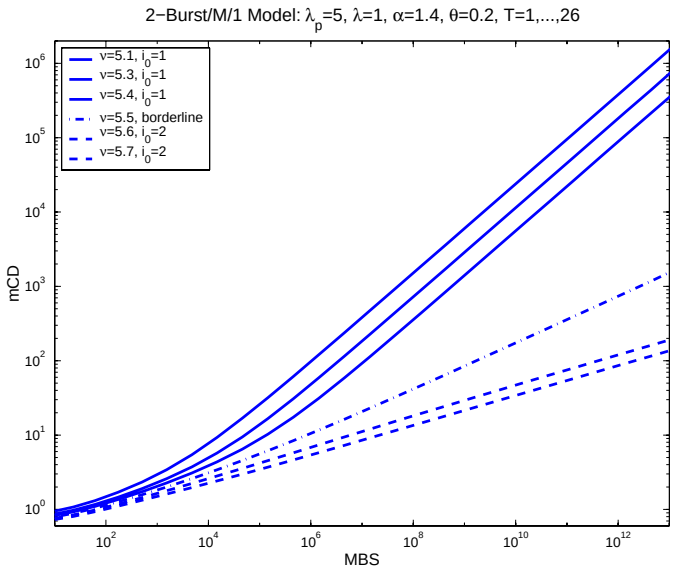

Figure 5.2: **Different Behavior of the Mean Cell Delay in the Blow-Up Regions:** With increasing MBS, the mCD grows as a Power-Law (appearing linearly on log-log scale), but the exponent (slope) changes from one blow-up region to the other. See also Sect. 5.3.2.

Although Eq. (5.1) is quite plausible, our arguments for the location of the blow-up points are not rigorous. Another seemingly plausible location for the blow-up points would be $\hat{\Lambda}(i, N)$, *the mean cell-arrival rate while at least i bursts are simultaneously active*. Λ_i and $\hat{\Lambda}(i, N)$ are *not* identical for most values of i and N.

In Fig. 5.1, the exact location of the first blow-up point (where $\nu = \Lambda_1$) is arguable. $\hat{\Lambda}(i, N) = 5.26$ is obtained for the model of Fig. 5.1, and there is no way to reject such a choice of the blow-up points purely from the results of mCD(ν) and BOP(ν) that are shown in that figure. However, a closer look at the behavior for increasing MBS for service rates ν close to the blow-up points in Fig. 5.2 reveals the exact location of the blow-up. Although the in-depth understanding of the behavior in Fig. 5.2 is provided only in Sect. 5.3, at this time it is enough to observe the behavioral change that occurs at $\nu = \Lambda_1 = 5.5$ in Fig. 5.2. We have made many other calculations of the type shown in Fig. 5.2 for different values of N and i_0 and have found that the blow-up points for large MBS are sharply defined, and invariably occur very near the Λ_i of (5.1). Furthermore, the discussion in the following section provides additional clarification that the behavior of the N-Burst/M/1 queueing model changes at $\nu = \Lambda_i$. Burstiness alone does not account for the observed behavior, but the existence of exceptionally long ON-times is a necessary feature.

Figure 5.1 shows that a dramatic increase of both mean Cell Delay and Buffer Overflow Probability occurs at a transition from blow-up region i_0 to $i_0 - 1$, which happens when the model parameters are changed such that $i_0 - 1$ simultaneous long-term active bursts can oversaturate the switch. One possibility for this blow-up to occur, a decrease in ν (which causes an increased utilization ρ), is shown in Figure 5.1. Another important scenario for blow-ups to occur follows from the rightmost expression in Eq. (5.2): for constant service-rate ν and constant average arrival rate λ, the blow-up can be achieved by an increase of the burstiness b, i.e. the same amount of data is transmitted in shorter ON periods while the OFF periods are stretched accordingly. Such a scenario is discussed in more detail in Sect. 5.6.

For N-Burst models other than the [IS] model, qualitatively the same blow-up effects can be observed. However, Eq. (5.1) has to be replaced by

$$\Lambda_i = \lambda',$$

where λ' is the average over-all cell-rate of an N'–Burst model of the same type with parameters

$$N' = N - i, \quad K' = K - i, \quad \lambda'_0 = i\lambda_p + \lambda_0,$$

$$\text{and} \quad \mu'(n) = \mu(n+i), \quad n = 0, 1, ..., N',$$

while the other parameters $\lambda'_p = \lambda_p$ and $\overline{n_p}' = \overline{n_p}$ remain. See Sect. 2.4 for the description of such more general models. The right-hand side of Eq. (5.2) must also be replaced by a more complicated expression – the parameter b is only meaningful for the [IS] model in any case.

After this brief excursion towards more general models, we restrict ourselves for the rest of the thesis to the [IS] model. The next section reveals that although a change of behavior of the queue-length distribution can be observed at all N blow-up points, the impact on the mean of the delay distribution depends strongly on a particular range for the PT exponent α of the burst-length distribution.

5.2 Power-Tailed Queue-Length Distribution

Obviously, Power-Tail distributed burst-lengths can cause very long over-saturation periods in the N-Burst/M/1 queue. Those over-saturation periods are caused by i_0 (or more) simultaneously long-term active sources and they have a major impact on the performance behavior. Consequently, we need to investigate the distribution of the duration of those over-saturation periods.

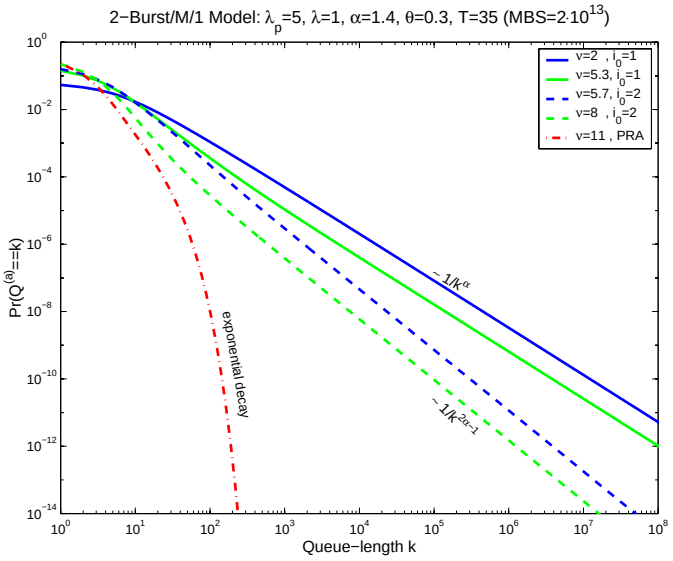

Figure 5.3: **Queue-Length Distribution of 2-Burst/M/1 Queues:** Except for Peak-Rate Allocation (where $N\lambda_p < \nu$), the queue-length distribution at cell-arrival times shows a Power-Tail, whose exponent β depends on the blow-up region i_0 as well as on the tail-exponent α of the burst-length distribution.

Only in the first blow-up region, $i_0 = 1$, does the single long burst correspond to an over-saturation period, whose reliability function $R_1(x) = R(x)$ therefore shows a Power-Tail with

the same exponent α as the burst-length distribution. The scenarios of having an over-saturation period which is caused by several simultaneously active sources lead to different tail-exponents, β, with $\beta > \alpha$.

It was shown in the Sect. 3.3 that the probability, $R_i(x)$, that none of i simultaneously long-term active bursts ends before time x has a Power-Tail with exponent $i\alpha - (i-1)$. Within such an over-saturation period, the queue-length grows with average rate $\Lambda_i - \nu$, which – by the definition (5.2) of i_0 – is positive for $i \geq i_0$. If we assume an empty queue at the beginning of the over-saturation period, the probability that queue-length q is reached within that period is approximately

$$p_i(q) := R_i\left(\frac{q}{\Lambda_i - \nu}\right) \sim \left(\frac{\Lambda_i - \nu}{q}\right)^{i(\alpha-1)+1} .$$

The assumption of an empty queue at the start of the over-saturation period is reasonable when the utilization, ρ, is not too close to 1. After all, the server is idle for the fraction $1 - \rho$ of the time.

It follows that

$$\frac{p_{i_0+k}(q)}{p_{i_0}(q)} \sim \frac{1}{q^{k\,(\alpha-1)}}, \quad \text{for} \quad k = 1, ..., (N - i_0)\,,$$

i.e. the probability that a long queue-length q is caused by an over-saturation period with $i_0 + k$ long-term active bursts becomes negligible compared to the over-saturation periods with only i_0 long-term active bursts. Thus, large queue-lengths $q \gg 0$ are most likely caused by an over-saturation period with only i_0 long-term active bursts. In other words, the duration of the over-saturation periods with i_0 long-term active sources dominates for large q since it has the 'heaviest' tail with exponent

$$\beta := \beta(i_0, \alpha) := i_0(\alpha - 1) + 1, \quad \text{for} \quad \alpha > 1. \tag{5.4}$$

Therefore, it is not the PT exponent α of the individual ON-period, but rather the tail exponent β of the duration of the over-saturation period, that determines the queueing behavior. Consequently, the often used 'Hurst-parameter', $H = (3 - \alpha)/2$, which characterizes the degree of Long-Range Dependence of the traffic, is by itself not sufficient to derive statements about the performance behavior of N-Burst/M/1 queues. Depending on the blow-up region i_0, the performance impact of the LRD properties can range from practically none (in the region of PRA) up to devastatingly high at the other extreme when $i_0 = 1$.

Previous work with fluid ON/OFF models, in particular [Jelenkovic & Lazar 99], showed that if the service-rate is equal to the cell-rate during ON, the reliability function of the queue-length distribution has a Power-Tail with exponent $\alpha - 1$. Their scenario corresponds to $i_0 = 1$, since an individual burst uses up all service capacity.

Compared to the fluid ON/OFF models, the N-Burst/M/1 queue adds further variability by its exponential service times and by the exponential inter-cell times within a burst. However, when considering long queue-lengths caused by long lasting over-saturation periods, fluid-flow and exponentially distributed service times[4] do not show fundamentally different behavior, since the convolution of q exponential random variables (i.e. the Erlangian-q distribution) decreases rapidly as $1/q$ in coefficient of variation. Consequently, the assumption of Poisson arrivals within

[4] The same argument holds for the process that describes the inter-cell times during a burst. They are assumed to be Poisson in the N-Burst model, but deterministic or continuous fluid-flow arrivals during ON periods can be interchangeably used without substantially affecting the behavior during over-saturation periods. See also Sect. 6.5.

a burst and of exponentially distributed service times does not have an impact on the tail of the queue-length distribution, if that tail is caused by Power-Tailed over-saturation periods.

Clearly, the N-Burst/M/1 queue is more general than the fluid-flow model of [Jelenkovic & Lazar 99], since the single burst might not use up all switch capacity, i.e. $i_0 > 1$ is possible. On the other hand, the process by which long queue-lengths are generated in the N-Burst/M/1 queue has the same underlying principle: potentially long over-saturation periods of Power-Tailed duration with i_0 permanently active sources. As a consequence, the tail-behavior as shown by [Jelenkovic & Lazar 99] still applies, but the tail-exponent β of the over-saturation periods replaces α:

$$\mathbb{P}\left\{Q^{(a)} = k\right\} \sim \frac{1}{k^{\beta}}, \quad \text{and} \quad \text{BOP}(B) := \mathbb{P}\left\{Q^{(a)} \geq B\right\} \sim \frac{1}{B^{\beta-1}} \cdot \tag{5.5}$$

$Q^{(a)}$ denotes the queue-length at *cell-arrival times*, but it can be shown that the asymptotic behavior of the distribution of $Q^{(a)}$ in an N-Burst/M/1 queue is the same as the asymptotic behavior of the queue-length as experienced by a random observer:

$$\mathbb{P}\left\{Q^{(a)} = k\right\} = \frac{1}{\rho}\,\mathbb{P}\left\{Q = k+1\right\} .$$

Within the worst blow-up region $i_0 = 1$, Eq. (5.5) is equivalent to the tail-behavior as derived by [Jelenkovic & Lazar 99], since $\beta(i_0 = 1, \alpha) = \alpha$.

Note that a steady-state distribution for $Q^{(a)}$ exists whenever $\rho < 1$. However, since the queue-length distribution has a Power-Tail it can have infinite moments (unless it is truncated, see Sect. 5.3). In particular, its mean is infinite when

$$\beta \leq 2 \quad \overset{(5.4)}{\Longleftrightarrow} \quad 1 < \alpha \leq 1 + \frac{1}{i_0} \cdot \tag{5.6}$$

Since we assume a FCFS (First Come First Serve) service strategy, the cell delay in turn is closely related to the queue-length, $Q^{(a)}$, at cell-arrival: a cell that arrives at a queue with $q_1^{(a)}$ cells in front of it, experiences an Erlangian-$\left(q_1^{(a)}+1\right)$ distributed delay with mean $\left(q_1^{(a)}+1\right)/\nu$, which approaches a deterministic distribution for large $q_1^{(a)}$. Thus, the cell-delay (CD) has the same tail-behavior,

$$\mathbb{P}\left\{\text{CD} > t\right\} \approx \mathbb{P}\left\{Q^{(a)} > t\,\nu\right\} \sim \frac{1}{t^{\beta-1}} \cdot \tag{5.7}$$

The first two moments of the cell-delay distribution can be calculated from the moments of $Q^{(a)}$:

$$\mathbb{E}\left\{\text{CD}\right\} = \frac{\mathbb{E}\left\{Q^{(a)}\right\}+1}{\nu}, \quad \mathbb{E}\left\{\text{CD}^2\right\} = \frac{\mathbb{E}\left\{\left(Q^{(a)}\right)^2\right\} + 3\,\mathbb{E}\left\{Q^{(a)}\right\} + 2}{\nu^2} .$$

In particular, the mean queue-length at cell-arrivals, $\mathbb{E}\left\{Q^{(a)}\right\}$, has a linear relationship to the mean cell-delay, mCD $= \mathbb{E}\left\{\text{CD}\right\}$. Consequently, the mCD becomes infinitely large under the same condition (5.6) as for the mean queue-length. Figure 5.4 illustrates such a divergence of the mCD for growing MBS in case $\alpha < 1 + 1/i_0$.

The arguments given above do not establish a rigorous mathematical proof of the tail behavior for the delay distribution. A rigorous proof for a fluid-flow ON/OFF model was developed by [Dumas & Simonian 00], who provide asymptotic bounds for the tail-behavior of the queue-length distribution, see also below. However, the reasoning in this thesis is confirmed by the

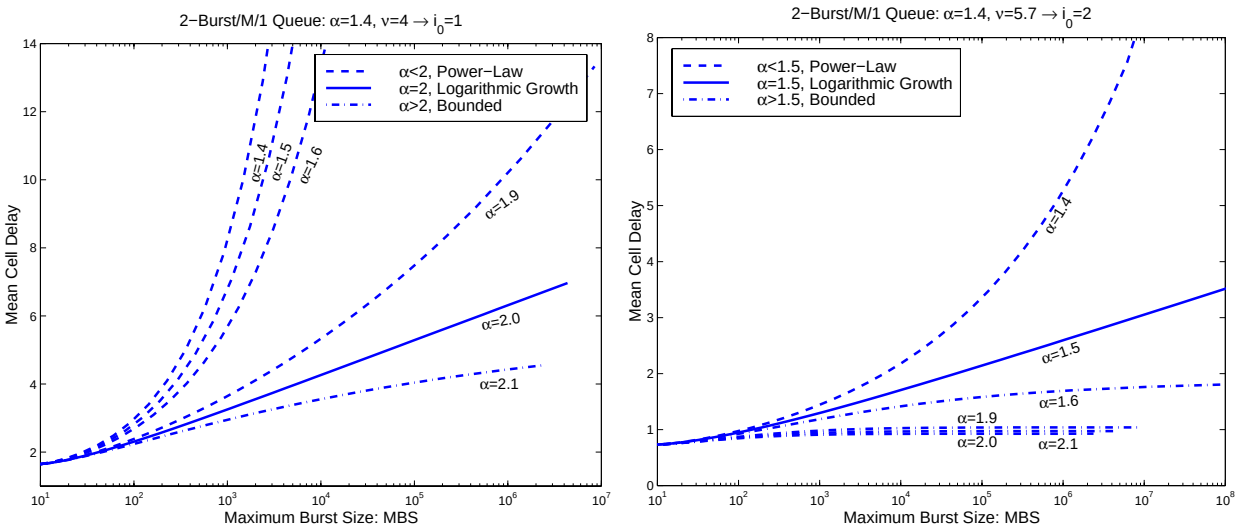

Figure 5.4: **Mean Cell Delay while Increasing the MBS:** Unbounded growth of the mean delay only occurs for $\alpha \leq 1 + 1/i_0$. The set of curves shown in the left graph are calculated for a 2-Burst/M/1 queue operating in blow-up region $i_0 = 1$ ($\rho = 25\%$). The curves on the right show the same model yet $i_0 = 2$ ($\rho \approx 18\%$), hence $\alpha = 1.5$ is the critical value. Note that within each graph, the utilization is constant.

parametric studies of the analytic N-Burst/M/1 models, as shown in Fig. 5.3. The numerical computations of the slopes of the linear parts of the queue-length distribution in the left graph of Fig. 5.5 correspond to the exponent β as derived in (5.4). Shown there is the function

$$\hat{\beta}(k_i) := -\frac{\log Pr\left(Q^{(a)} = k_{i+1}\right) - \log Pr\left(Q^{(a)} = k_i\right)}{\log k_{i+1} - \log k_i},$$

which is the exponent of a fitted Power-Law for the queue-lengths k_i and k_{i+1}.

Even more important, the results displayed in Fig. 5.3 and other parametric studies show that the tail-behavior can already be observed for queue-lengths as short as 10^2 cells. Since buffer-sizes in the order of 10^5 cells are realistic, these effects are expected to have an impact on real systems as well. Analytic arguments about the asymptotic behavior of comparable models are very valuable by itself, but in practice it is necessary to know where the range of applicability begins.

Tail-constants

The knowledge of the tail-exponent β of the queue-length distribution is already sufficient for many qualitative performance results, such as for the explanation of the blow-up effects. For quantitative results however, the tail-constant of the Power-Tail is also necessary. Figure 5.5 shows in its right graph the estimates for the tail-constants of the queue-length distribution that are derived from the numerically computed tail-probabilities.

At present, exact formulas for the direct computation of the tail-constant of the queue-length distribution of the N-Burst/M/1 model are unknown. However, formulas and bounds for related scenarios are known that can be used as approximations for the true values of the tail-constant. See [Jelenkovic & Lazar 99] and [Dumas & Simonian 00] for the mathematical details.

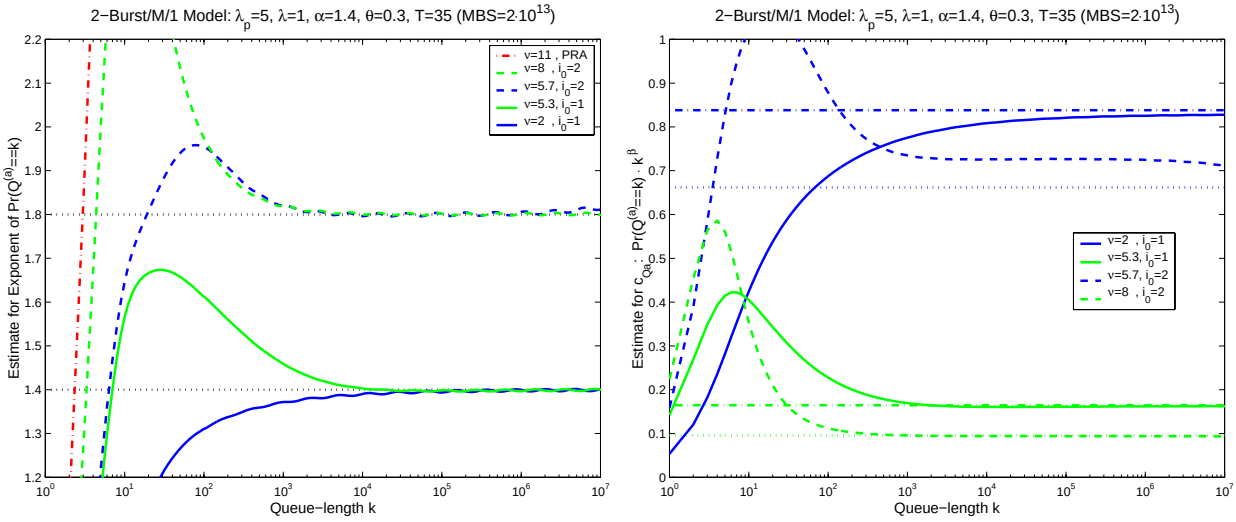

Figure 5.5: **Power-Law behavior of Queue-length Distribution in Blow-up Regions:** The left graph plots an estimate for the PT-exponent β as measured along the curve of the queue-length distribution. The convergence towards the theoretically derived value $\beta = i_0(\alpha - 1) + 1$ becomes obvious. The right graph shows the resulting tail-constants in comparison to the values that are obtained from the approximation formulas (dashed-dotted respectively dotted horizontal lines). Except for the case $\nu = 5.7$, the approximation formula works very well.

A lower bound for the tail-constant of the Buffer-Overflow Probability

$$\mathrm{BOP}(B) = Pr\left(Q^{(a)} \geq B\right) \sim c_{\mathrm{BOP}}\, B^{1-\beta},$$

can be derived from results in [Dumas & Simonian 00]:

$$c_{\mathrm{BOP}} \geq \frac{1}{\rho} \cdot \left[\frac{(1-b)}{\alpha - 1}\, (b\, i_\Delta\, \overline{n_p})^{\alpha-1}\, c_{\mathrm{PT}}^{(1)}(\alpha)\right]^{i_0}, \tag{5.8}$$

where $c_{\mathrm{PT}}^{(1)}(\alpha)$ is the tail-constant of the reliability function of the burst-length distribution when scaled to have mean $\overline{x_p} = 1$. For the TPT distributions used herein, that constant was determined for $\alpha = 1.4$ in [Klinger 97] to be

$$c_{TPT}^{(1)}(\alpha = 1.4) \approx 0.2138\,.$$

A table with tail-constants $c_{TPT}^{(1)}(\alpha)$ for various values of α is given in Appendix I.1.

For an approximation of c_{BOP}, the following extension of the lower bound (5.8) can be used:

$$c_{\mathrm{BOP}} \approx \frac{b^\beta\,(1-b)^{i_0-1}}{1-\rho} \left[\frac{(i_\Delta\, \overline{n_p})^{\alpha-1}}{\alpha - 1}\, c_{\mathrm{PT}}^{(1)}(\alpha)\right]^{i_0}. \tag{5.9}$$

Note that the approximation in (5.9) depends on the parameter N (via i_0 and i_Δ), b, $\overline{n_p}$, α and ρ, but not on λ. In this parameter-set, λ is a pure scaling parameter.

Numerical experiments with N-Burst/M/1 queues showed that (5.9) can be considered exact in blow-up region $i_0 = 1$, where the formula simplifies to:

$$c_{\mathrm{BOP}} \approx \frac{b^\alpha}{1-\rho}\, \frac{(i_\Delta\, \overline{n_p})^{\alpha-1}}{\alpha - 1}\, c_{\mathrm{PT}}^{(1)}(\alpha), \quad \text{for} \quad i_0 = 1. \tag{5.10}$$

In blow-up region $i_0 = 2$, numerical experiments for $\alpha = 1.3$ and $\alpha = 1.4$ suggest that the right-hand side of (5.9) has to be multiplied by another factor $(1 - i_\Delta)^{-0.5}$, otherwise (5.9) underestimates c_{BOP}. However, the exact formula for c_{BOP} for $i_0 > 1$ is unknown at this time, even for fluid-flow models.

The right graph in Fig. 5.5 compares the observed tail-constants from the numerical computation,

$$\widehat{c_{Qa}}(k) := \mathbb{P}\left\{Q^{(a)} = k\right\} \cdot k^\beta,$$

with the results from the approximation formula (5.10) when $i_0 = 1$, respectively with the extended version of (5.9):

$$c_{\mathrm{BOP}} \approx \frac{b^\beta\,(1-b)}{(1-\rho)\,\sqrt{1-i_\Delta}} \left[\frac{(i_\Delta\,\overline{n_p})^{\alpha-1}}{\alpha-1}\, c_{\mathrm{PT}}^{(1)}(\alpha)\right]^2 \quad \text{for} \quad i_0 = 2. \tag{5.11}$$

Note that the tail-constants c_{BOP} and c_{Qa} are related by $c_{Qa} = (\beta - 1)\, c_{\mathrm{BOP}}$. The comparison of the numerical results in Fig. 5.5 shows that the approximation formulas work very well, with some restrictions in the case $\nu = 5.7$, which is close to the actual blow-up point ($i_\Delta = 0.96$ for $\nu = 5.7$). In general, the results of the approximation formulas can deviate from the observed tail-constants when the model operates very close to the blow-up points itself (i_Δ close to 0 or close to 1). At the blow-up points themselves, the behavior is different in any case, so they need to be treated separately. However, since the blow-up points are only singular points, we focus on the behavior in the blow-up regions here.

5.3 Impact of Truncated Tails

In this section, we take a closer look at the impact of the truncations in the tail, regardless of whether they are imposed as part of traffic policing, or they occur due to bounded physical resources, or due to the finite duration of the busy hours.

5.3.1 Truncation of the Queue-Length Distribution

If truncated Power-Tails are used for the burst-length distribution in the N-Burst model, the queue-length distribution still shows Power-Tail properties with exponent β, but at some point the tail is truncated, see Fig. 5.6.

Having introduced the PT Range $\overline{x_T}$ – or equivalently the MBS (see Sect. 3.5) – to characterize the truncated Power Tails of the burst-length distribution, we are now able to discuss the impact on the queue-length distribution in N-Burst/M/1 queues. However, as pointed out in Sect. 5.2, the individual burst does not matter in terms of performance when $i_0 > 1$, but rather the over-saturation period. The PT Range of the duration of an over-saturation period with $i \geq i_0$ long-term active sources follows from Eqs. (B.10) and (B.8):

$$\overline{x_T}[i] := \mathrm{Rng}\left(\mathbf{B}_T^{\oplus i}\right) = \frac{\overline{x_T}}{i}\,. \tag{5.12}$$

The matrix $\mathbf{B}_T^{\oplus i}$ – the Kronecker-Sum with i summands $\mathbf{B}_T$ – is the rate matrix of the ME representation for the duration of the periods with i active bursts, see Appendix B.3.

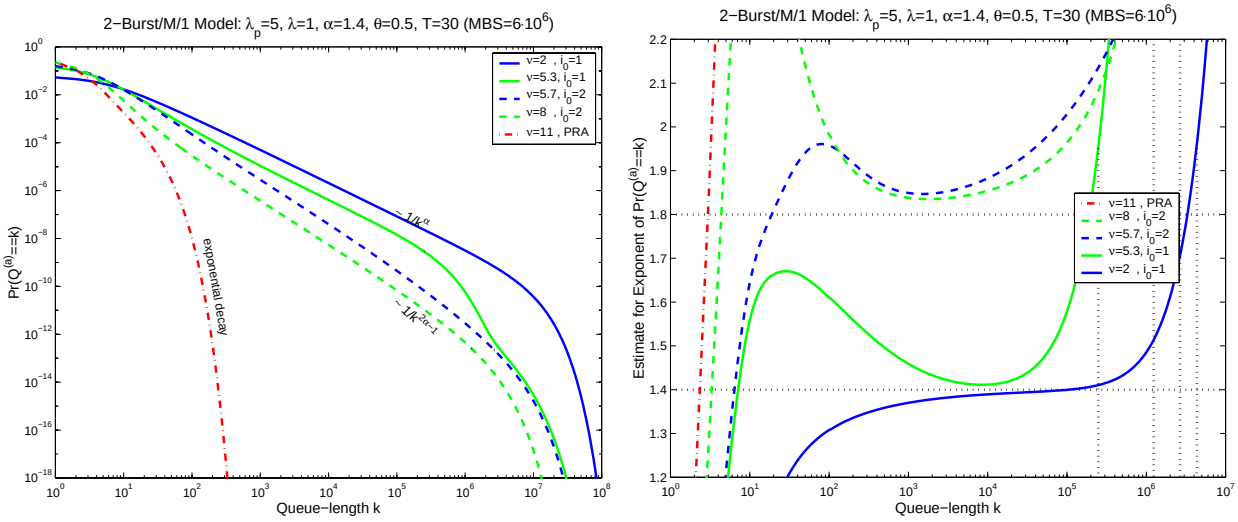

Figure 5.6: **Truncated Power-Law behavior of Queue-length Distribution in Blow-up Regions:** Power-Law behavior of the queue-length distribution with exponent β can still be observed, but such behavior is now truncated at queue-lengths around $k = 10^6$ or $k = 10^7$. Note that the measured tail-exponents in the right graph are slightly higher than the theoretically expected $\beta(i_0 = 1) = 1.4$ and $\beta(i_0 = 2) = 1.8$.

During the long over-saturation periods, cells accumulate in the queue with rate $(\Lambda_i - \nu) > 0$. Consequently, we would expect the queue-length distribution to have a PT Range of

$$\overline{q}[i] := (\Lambda_i - \nu) \cdot \overline{x_T}[i] = \left[i\,(\lambda_p - \kappa) + \lambda \left(1 - \frac{1}{\rho}\right) \right] \cdot \frac{\overline{x_T}}{i} = \text{MBS} \cdot \frac{b}{i} \cdot \left[i - N\,\frac{1-b}{b}\,\frac{1-\rho}{\rho} \right] . \tag{5.13}$$

From Eq. 5.13, it follows that $\overline{q}[i+1] > \overline{q}[i]$. Although according to Sect. 5.2, over-saturation periods with $i > i_0$ long-term active sources contribute only lighter tails with exponent $\beta(i) > \beta(i_0)$ to the queue-length distribution, they show a bigger PT range $\overline{q}[i] > \overline{q}[i_0]$. Consequently, the impact of such i-active over-saturation periods is dominated by the i_0 active over-saturation periods for queue-lengths up to $\overline{q}[i_0]$ only, where

$$\overline{q}[i_0] = b \cdot i_\Delta \cdot \frac{\text{MBS}}{i_0} \cdot \tag{5.14}$$

For larger queue-lengths $\overline{q}[i_0] < k < \overline{q}[i]$, the Power-Tails with larger exponents $\beta(i) > \beta(i_0) = \beta$, $i > i_0$, can show their impact.

Figure 5.7 shows the function R/f (which converges to the PT Range, see Eq. (3.10)) for the queue-length distribution of a 2-Burst/M/1 queue in blow-up regions $i_0 = 1$ and $i_0 = 2$. In the case of $i_0 = N = 2$, there exist only over-saturation periods with $i_0 = 2$ long-term active sources. As a consequence, the fraction R/f converges smoothly to the PT Range $\overline{q}[i_0]$. If $i_0 = 1$ on the other hand, the fraction R/f in Fig. 5.7 shows a plateau at $\overline{q}[1]$, before it converges to $\overline{q}[2]$. It indeed shows a similar behavior to a mixture of different TPTs as shown in the right graph of Fig. 3.2. Note that the PT Range $\overline{q}[N]$ of the queue-length distribution is closely related to the so-called *caudal characteristic* (see [NEUTS 86]), which is the largest eigenvalue of the *Rate Matrix* $\boldsymbol{\mathcal{R}}$ of the Matrix-Geometric solution for the queue-length probabilities, see Appendix D.7.

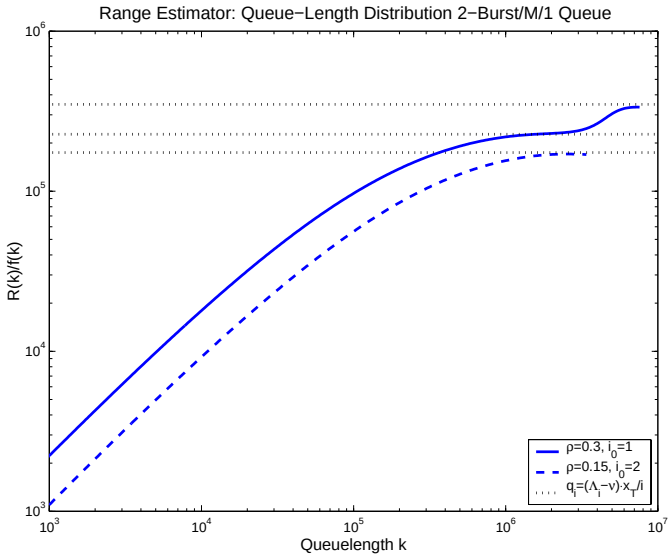

Figure 5.7: **PT Range of the Queue-Length Distribution for 2-Burst/M/1 Queues:** In blow-up region $i_0 = 1 < N$, the quotient $R(k)/f(k)$ shows a plateau at $\overline{q}[1]$ before it finally converges to $\overline{q}[2]$, see text.

For long tails of the burst-length distribution (large $\overline{x_T}$) however, the PT Range $\overline{q}[i_0]$ of the queue-length distribution for the smallest PT exponent β becomes large as well. As a consequence, the probabilities in the Power-Tails with larger exponents $\beta(i) = i(\alpha - 1) + 1$, that are caused by over-saturation periods with $i > i_0$ long-term active bursts, have decayed to small values for queue-lengths $q > \overline{q}[i_0]$. Therefore, although the Power-Tails with larger exponents show a larger PT Range, the probabilities that they contribute can be practically negligible for large $q > \overline{q}[i_0] \gg 0$.

5.3.2 Impact of Truncated Power-Tails on Mean Delay

The behavior of the mean delay in N-Burst/M/1 queues with TPT distributed burst-lengths follows from the queue-length distribution. Obviously, a truncation mainly matters for the scenarios where infinite delay moments result from full Power-tails. According to Eq. (5.6), the mean delay becomes infinite when the model operates in blow-up region i_0 and $\alpha \leq 1 + 1/i_0$. For simplicity, we exclude the boundary value $\alpha = 1 + 1/i_0$ in the following discussion and assume that $\alpha < 1 + 1/i_0$. It follows that the mean queue-length and the mean delay are mainly determined by the (truncated) Power-tails of their distributions. Asymptotically, we can neglect the body of the function and only look at the Power-tail regime from some value $h > 0$ on:

$$\begin{aligned} \mathbb{E}\{\mathrm{CD}\} &= \frac{\mathbb{E}\left\{Q^{(a)}\right\} + 1}{\nu} \sim \frac{1}{\nu} \int_h^{\overline{q}[i_0]} t \, \frac{1}{t^\beta} \, dt \\ &\sim \frac{1}{\nu\,(2-\beta)} \left(\overline{q}[i_0]\right)^{2-\beta} \sim \mathrm{MBS}^{1-i_0(\alpha-1)} \quad \text{if} \quad \alpha < 1 + \frac{1}{i_0} \cdot \end{aligned} \tag{5.15}$$

Thus, we expect the mean delay to grow asymptotically by a Power-Law with exponent $(2-\beta) = 1 - i_0(\alpha - 1)$ when increasing the MBS, but only if $\alpha < 1 + 1/i_0$. The derivation of (5.15) uses the tail-behavior (5.5) of the queue-length distribution up to the PT-Range $\overline{q}[i_0]$ as in Eq. (5.14).

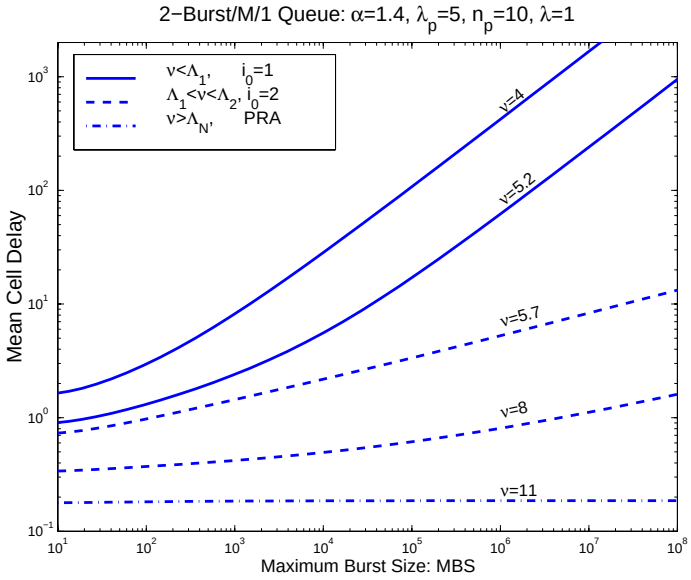

Figure 5.8: **Mean Cell Delay while Increasing the MBS:** Except for Peak-Rate Allocation (PRA, where $N\lambda_p < \nu$), the mean delay asymptotically relates to the MBS by a Power-Law, whose exponent only depends on α and i_0.

Note that if $\alpha = 1 + 1/i_0$, the tail-exponent of the over-saturation period is $\beta = 2$, and thus the integral is taken over $1/t$ and thus grows logarithmically with the upper integration boundary (as observed in Fig. 5.4).

The tail-constant in Eq. (5.15) can be bounded from above asymptotically for large MBS:

$$\text{mCD} \approx \frac{1}{\nu} \int_0^{q_{i_0}} c_{\text{BOP}}\, B^{1-\beta}\, dB < \frac{\rho}{\lambda} \frac{c_{\text{BOP}}}{2-\beta} \left[q_{i_0}\right]^{2-\beta} = \frac{\rho\, c_{\text{BOP}}}{\lambda\,(2-\beta)} \left[\frac{b\, i_\Delta}{i_0}\right]^{2-\beta} \text{MBS}^{2-\beta} \quad \text{for} \quad \beta < 2 \tag{5.16}$$

The Power-Laws as well as the specific exponents are confirmed by the results that are computed for the analytic 2-Burst/M/1 model, shown in Fig. 5.8. Since the plotted graphs are exact results that already show the Power-Law behavior (appearing linear) for a realistic range of values for the MBS, the asymptotic result (5.15) is not just a mathematical limit, but it is applicable in realistic scenarios.

An analogous derivation for the ℓ-th moment of the Cell Delay distribution leads to the more general relationship:

$$\mathbb{E}\left\{\text{CD}^\ell\right\} \sim \text{MBS}^{\ell+1-\beta} = \text{MBS}^{\ell - i_0\,(\alpha-1)} \quad \text{if} \quad \alpha < 1 + \frac{\ell}{i_0} \tag{5.17}$$

5.4 Consequences of (Truncated) Power-Tails

So far in this chapter, the steady-state performance analysis of N-Burst/M/1 queues with infinite buffers has revealed a very peculiar behavior when LRD properties are present in the N-Burst

traffic. In the following, we summarize the findings of the last three sections and point out their practical implications:

1. **Power-Tailed Queue-length Distribution:**
   In blow-up region i_0, the queue-length distribution of an N-Burst/M/1 queue is Power-Tailed with exponent $\beta - 1 = i_0(\alpha - 1)$, see Eq. (5.4). The tail-exponent β changes discontinuously at the N blow-up points, which are characterized by $\nu = \Lambda_i$, $i = 1, ..., N$, see Eq. (5.1). Consequently, β depends on network parameters as well as on traffic parameters. Hence, the careful design of a network and careful traffic management can alleviate the negative performance impacts of LRD, see also Chapter 7.

   Note that the blow-up points for bursty sources (high b) can occur at very low utilization values,

   $$\rho_i = \frac{\lambda}{\Lambda_i} = \frac{1}{1 + \frac{i}{N}\frac{b}{1-b}}, \quad i = 1, ..., N\,.$$

   In the measurements of Chapter 4, we observed burstiness parameters of the order $b \approx 0.9$. Therefore, the blow-up at the transition from PRA ($\nu > N\lambda_p$) to blow-up region $i_0 = N$ always occurs at $\rho_N = 1 - b = 10\%$. When multiplexing for instance $N = 5$ such sources at a switch, the additional four blow-ups occur at utilization $\rho_4 = 12.2\%$, $\rho_3 = 15.6\%$, $\rho_2 = 21.7\%$, and $\rho_1 = 35.7\%$.

2. **Infinite Moments:**
   As a consequence of the Power-Tailed queue-length distribution (5.5), the k-th moment of the queue-length distribution becomes infinite when

   $$\beta - 1 \leq k \quad \overset{(5.4)}{\Longleftrightarrow} \quad 1 < \alpha \leq 1 + \frac{k}{i_0}\,. \tag{5.18}$$

   In particular, the mean queue-length and thus the mean delay become infinite, if $\alpha \leq 1 + 1/i_0$. Therefore, the borderline between infinite and finite mean delay is only at $\alpha = 2$, when $i_0 = 1$, see Fig. 5.4.

   On the other hand, for given $1 < \alpha \leq 2$, the mean delay is infinite only in blow-up regions $i_0 \leq 1/(\alpha - 1)$. When using truncated tails, the mean delay is always finite. However, the blow-up effects for mean delay are especially pronounced in the blow-up regions with $i_0 \leq 1/(\alpha - 1)$, because there, the tail of the delay distribution has a much greater impact on the mean.

3. **Buffer-Inefficiency:**
   The *Buffer-Overflow Probability* (BOP) follows directly from the queue-length distribution,

   $$\mathrm{BOP}(B) := Pr\left(Q^{(a)} \geq B\right)\,.$$

   Consequently, the knowledge of the Power-Tail behavior Eq. (5.5) for the queue-length distribution immediately implies a Power-Law behavior for the BOP:

   $$\mathrm{BOP}(B) \sim c_{\mathrm{BOP}}/B^{\beta-1} = c_{\mathrm{BOP}}\, B^{i_0(1-\alpha)}\,. \tag{5.19}$$

   The practical implication is that additional buffer-space only reduces the BOP very slowly, called *Buffer-Inefficiency*. Such a slow Power-Law decay for LRD ON/OFF traffic does not occur for exponential burst-lengths, where the BOP decreases exponentially fast with increasing buffer-size.

Note that the Buffer Inefficiency occurs for any value of β, including $\beta > 2$ (as opposed to the behavior of mean Cell Delay, which requires $\beta < 2$ to be large). Nevertheless, large values of β alleviate the effects, since the Power-Law drop-off (5.19) is more rapidly.

4. **Truncated Tails:**
   If the maximum size of the ON periods is restricted, i.e. the Power-Tail of the ON time distribution is truncated at some point, the queue-length distribution also shows a truncated Power-Tail. Formulas for the PT Range $\overline{q}[i]$, $i \geq i_0$ of the queue-length distribution could be derived. It follows that $\overline{q}[i] < \text{MBS}$ holds for any $i \geq i_0$. Hence, if the buffer is at least of size MBS, buffer-overflow events become very unlikely, because the switch model operates in the region of the exponential drop-off of the BOP and not in the Power-Law regime.

   As mentioned before, in the scenario of $\beta < 2$ or equivalently $\alpha < 1 + 1/i_0$, the mCD is infinite for full PTs with exponent α for the ON time distribution. If the tails are truncated to MBS cells, the mCD shows a Power-Law relationship

$$\text{mCD}(\text{MBS}) \sim \text{MBS}^{2-\beta}.$$

The four points enumerated above present qualitative results. In order to obtain quantitative results, the exact numerical evaluation of the N-Burst/M/1 model can be used. The much simpler asymptotic behavior for large buffers provide an alternative, but then, the tail-constants of the Power-Law behavior are required in addition to the tail-exponent. Approximations for the tail-constants are provided in Sect. 5.2. Chapter 7 uses both approaches, the exact numerical computation and the approximated asymptotic behavior, for realistic scenarios of the network planning process.

5.5 Finite Buffer Systems

As the left graph of Figure 5.9 shows, the *Cell Loss Probability* (CLP) in a finite-buffer queue also shows a Power-Law behavior with the same exponent $\beta - 1$ as the BOP, but with smaller tail-constant $c_{\text{CLP}} < c_{\text{BOP}}$:

$$\text{CLP}(B) \sim c_{\text{CLP}}\, B^{1-\beta}. \tag{5.20}$$

Thus the conclusions that can be drawn from this Power-Law behavior, in particular the Buffer-Inefficiency, are the same as in the discussion of the BOP.

Results from the transient analysis in Chapter 6 lead to a relationship between the two tail-constants c_{BOP} and c_{CLP}:

$$(1-\rho)\, i_\Delta \, \frac{1 - i_\Delta}{i_0 - i_\Delta}\, c_{\text{BOP}} \leq c_{\text{CLP}} \leq (1-\rho)\, \frac{i_\Delta}{i_0}\, c_{\text{BOP}}. \tag{5.21}$$

The upper and lower bound are identical for $i_0 = 1$, thus

$$c_{\text{CLP}} \approx (1-\rho)\, i_\Delta \, c_{\text{BOP}} \quad \text{for} \quad i_0 = 1. \tag{5.22}$$

The right graph in Figure 5.9 plots the exact values of $\text{CLP}(B)$ in comparison to the Power-Laws $c_{\text{CLP}} B^{1-\beta}$. In this experiment, the tail-behavior of the CLP is close to the upper bound of (5.21), i.e.

$$c_{\text{CLP}} \approx (1-\rho)\, \frac{i_\Delta}{i_0}\, c_{\text{BOP}},$$

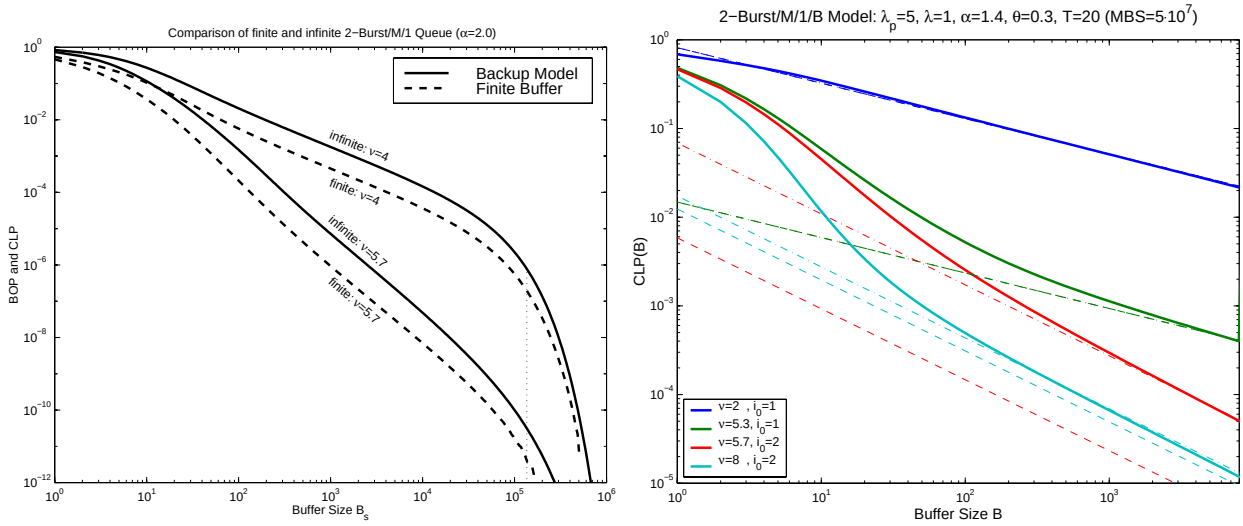

Figure 5.9: **Queue-Length Distribution in Case of Finite Buffers:** A Power-Law decay with same exponent $\beta - 1$ as for the BOP can be observed for the CLP in a finite buffer loss model. Approximations and bounds can be given for the tail-constant c_{CLP} which provide a good approximation when $i_0 = 1$. For $i_0 = 2$ the tail-constant c_{CLP} that is seen in the exact numerical CLP values (solid line) is close to the upper bound for c_{CLP} (which is used to compute the dashed-dotted line). The lower-bound for c_{CLP} (which is used to compute the dashed lines) can be substantially smaller when operating close to the blow-up point (as in the case of $\nu = 5.7$).

where c_{BOP} is derived from Eq. (5.11) for $i_0 = 2$.

Note that in simpler queueing models such as the M/M/1[/B] queue or the M/G/1[/B] queue, the ratio of CLP and BOP approaches

$$\frac{\mathrm{CLP}(B)}{\mathrm{BOP}(B)} \longrightarrow (1-\rho) \quad \text{for large } B.$$

For the N-Burst/M/1[/B] queue with PT distributed burst-lengths that ratio can be arbitrarily small for fixed ρ. In blow-up region $i_0 = 1$,

$$\frac{\mathrm{CLP}(B)}{\mathrm{BOP}(B)} \longrightarrow \frac{c_{\mathrm{CLP}}}{c_{\mathrm{BOP}}} = (1-\rho) i_\Delta \,,$$

with $0 < i_\Delta < 1$. Therefore, in N-Burst/M/1 queues with LRD properties, the assumption of a finite-buffer scenario can have a much greater quantitative impact than in simple M/G/1 models.

Delay in finite buffer systems

In the infinite-buffer model, one conclusion is that although the delay distribution is well defined for $\rho < 1$, it can have an infinite mean for $\alpha < 1 + 1/i_0$. This of course changes in the finite-buffer N-Burst/M/1/B model. The *maximal* delay is Erlangian-B distributed in this scenario and hence the mean delay is limited by the buffer-size. Figure 5.10 illustrates that limiting impact for a buffer of $B = 10^4$ cells. first, the mCD grows as $\mathrm{MBS}^{2-\beta}$ as in the infinite buffer model (Fig. 5.8). But at some point (when the MBS is in the order of the buffer-size B), the mCD(MBS) starts to converge to a horizontal asymptote.

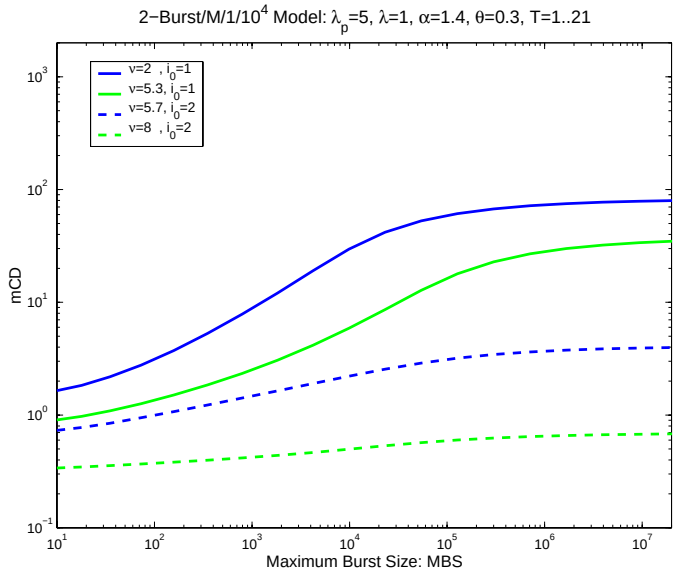

Figure 5.10: **Mean Delay in Finite-Buffer Model:** The Power-Law growth that is observed in the infinite buffer model in Fig. 5.8 is now only observed for an intermediate range of MBS values, before the upper limit from the buffer-size shows its restricting impact. Here a finite buffer of $B = 10^4$ cells is used.

5.6 Approximation Models

We have demonstrated the blow-up effects in Section 5.1 in the scenario of a fixed traffic description (fixed N-Burst parameters N, κ, b, $\overline{n_p}$) but a varying utilization of the queueing model, $\rho = N\kappa/\nu$, due to a change of the service-rate ν. Since the blow-up effects happen at a transition from one blow-up region i_0 to the next region, and

$$i_0 = \left\lceil N\,\frac{1-\rho}{\rho}\,\frac{1-b}{b} \right\rceil = \left\lceil \left(\frac{\nu}{\kappa} - N\right)\,\frac{1-b}{b} \right\rceil,$$

a change in any of the base parameters N, κ, b (but not $\overline{n_p}$) of the N-Burst model can cause such a blow-up effect. Remember that we assume $\lambda_0 = 0$ in our discussion, since a positive λ_0 can be deducted from the service-rate for our purposes.

A variation of the burstiness parameter b while keeping the other N-Burst parameters constant is of particular practical and theoretical interest. Figure 5.11 shows the behavior of the mCD and the BOP in such a scenario. The practical significance of that scenario is that a reduction of b corresponds to a reduced peak-rate λ_p of the individual sources while keeping the overall amount of traffic the same. Such a procedure is commonly called *traffic shaping*, and Figure 5.11 demonstrates clearly that a substantial improvement in QoS can be achieved, if the shaping causes a transition to a better (higher i_0) blow-up region. From the condition $\nu = \Lambda_i$ for the blow-up points and Eq. (5.1), it follows that such transitions occur at

$$b_i = \frac{1}{1 + \frac{i}{N}\,\frac{\rho}{1-\rho}}, \quad i = 1, 2, ..., N.$$

Note that $b_N = 1 - \rho < b_{N-1} < b_1 < 1$.

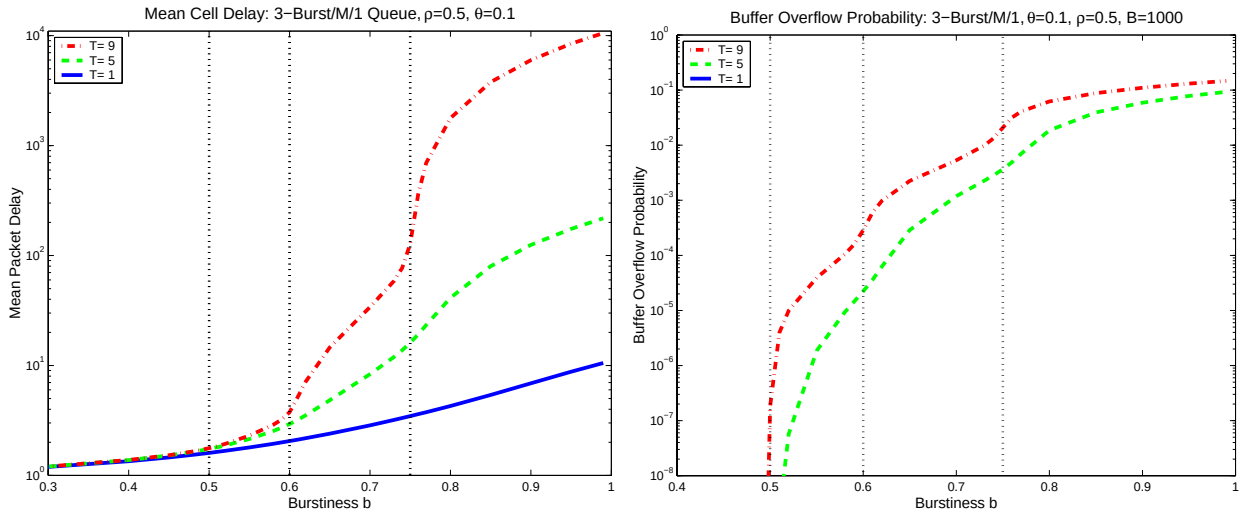

Figure 5.11: **Blow-up Effects from Varying the Burstiness:** A change of traffic parameters, here the burstiness b, can also result in the blow-up effects. For $b = 0$ a limiting M/M/1 queue is obtained, while at the other extreme, $b = 1$, the N-Burst model approaches a bulk-arrival model.

The theoretical (and thus after more elaboration also another practical) significance of Fig. 5.11 appears at the two extremal values of b. As already described in Sect. 2.7, the N-Burst/M/1 model reduces to a simple M/M/1 model for $b \to 0$. At the other end $b = 1$, one obtains a bulk-arrival model. Explicit formulas exist for the mCD of both limiting models:

$$\mathrm{mCD}(b = 0) \quad = \quad \frac{1}{\nu(1-\rho)} \tag{5.23}$$

$$\mathrm{mCD}(b = 1) \quad = \quad \frac{1}{\nu(1-\rho)}\, D \tag{5.24}$$

$$\text{with} \quad D = \left\lceil \frac{\mathbb{E}\left\{L(L+1)\right\}}{2\,\mathbb{E}\left\{L\right\}} \right\rceil . \tag{5.25}$$

Here, L is the random variable that expresses the number of cells per burst, see Sect. 3.5 and Appendix C.4 for the distribution of L in the N-Burst model.

Since performance becomes monotonically worse from $b = 0$ to $b = 1$, the two limit models provide a lower and upper bound for performance. For the intermediate region that contains the blow-up points b_i, the full N-Burst model is necessary. See [Lipsky & Schwefel 00] for a more detailed discussion.

Bulk-arrival models as well as M/G/1 queues are used in [Schwefel99a] to compare the scenario of deterministic service-times of limiting N-Burst/D/1 models with the usual exponential service-time assumption. For TPT distributed burst-lengths with large PT Range, deterministic and exponential service times resulted in practically identical values for the mCD (relative deviation of 10^{-6}).

Fluid-Flow limit

Section 2.7.3 showed how a fluid-flow ON/OFF model can be obtained in the limit as $k \to \infty$ of N-Burst models. k is the 'cell-size' divisor, i.e. each cell is replaced by k smaller units.

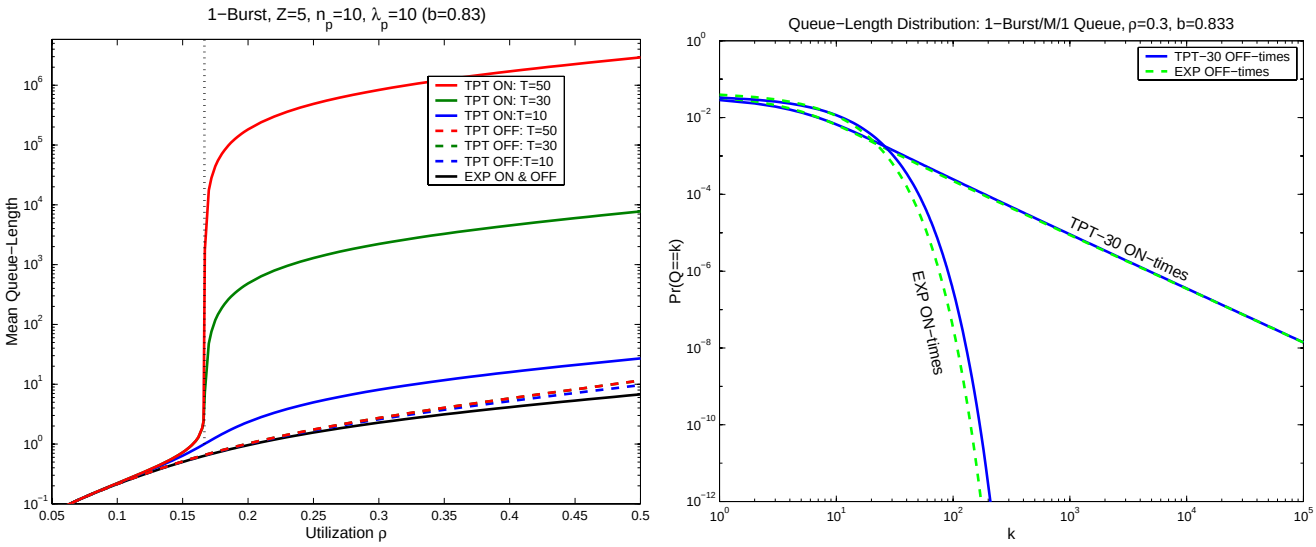

Figure 5.12: **Impact of TPT distributions in ON and OFF times:** Only when the TPT distributions are used for the ON periods (solid lines in left graph), are the characteristic blow-up effects observed that lead to huge mean queue-lengths. TPT distributed OFF periods (dashed lines in left graph) on the other hand show little impact. The impact on the queue-length distribution for fixed $\rho = 0.3$ is shown in the right graph.

See [Lipsky & Schwefel 00] for a more detailed discussion of that limit. We just state the performance behavior for the two boundaries $b = 0$ and $b = 1$:

$$\mathrm{mCD}(b=0,\,k) \;=\; \frac{1/(k\nu)}{1-\rho} = \frac{1}{k}\cdot \mathrm{mCD}(b=0) \tag{5.26}$$

$$\mathrm{mCD}(b=1,\,k) \;=\; \mathrm{mCD}(b=1) - \frac{1/\nu}{1-\rho}\left[\frac{1}{2} - \frac{1}{2k}\right] \tag{5.27}$$

By definition, $\mathrm{mCD}(b,\,k=1) = \mathrm{mCD}(b)$. It also follows that the quantity, $\nu(1-\rho)\cdot\mathrm{mCD}(b=1,\,k)$ is independent of ν and ρ for all k. Since $\mathrm{mCD}(b=1)$ grows unboundedly with the MBS when TPT distributions with $\alpha < 2$ are used (see (5.24)), and since the second term of (5.27) is independent of the MBS, the actual value of k is practically irrelevant for $\mathrm{mCD}(b,\,k=1)$ for TPT distributed burst-lengths with large MBS and $\alpha < 2$. The difference between the fluid-flow bulk arrival limit ($k \to \infty$) and the point-process bulk arrival model ($k = 1$) is given by the limit, $1/[2(\nu - \lambda)]$, of the second term of (5.27).

5.7 Impact of OFF-Time Distribution

So far, we always assumed the OFF times of the ON/OFF sources to be exponentially distributed. It was mentioned in Section 2.3 that the distribution of the OFF time is less critical in terms of performance. We illustrate this claim with numerical results for a 1-Burst model in Fig. 5.12. When TPT distributions are used in the ON periods, the characteristic blow-up effects at $\rho = 1 - b$ can be observed. Within the blow-up regions, huge mean queue-lengths of the 1-Burst/M/1 queue are obtained, see left graph of Fig. 5.12. In contrast to that, the use of TPTs only in the OFF periods increases the mean queue-length only marginally and no blow-up effects happen.

The impact on the queue-length distribution is shown in the right graph of Fig. 5.12: Only when TPT distributions are used in the ON periods, does the Power-Law behavior of the queue-length distribution show up. In that case, the OFF time distribution has very little impact. When the ON-time is exponential, the impact of a TPT OFF-time is at least visible, but the queue-length distribution nevertheless shows a quick exponential drop-off.

Note that a 1-Burst model with exponential ON times and $\lambda_0 = 0$ is a renewal process, even if the OFF periods are PT (or TPT) distributed. Therefore, the inter-cell times show no LRD properties. The counting process on the other hand shows LRD properties for TPT distributed OFF periods. Consequently, this process is an example of an ON/OFF model in which the long-range correlation in the counts causes hardly any performance impact for any value of ρ that is not too close to 1. We have not investigated the scenarios of $\rho \to 1$, since they are of little practical relevance. However, known results from the discussion of TPT/M/1 models *might* be transferable to 1-Burst/M/1 queues with TPT distributed OFF times and exponential ON times: [Greiner et al. 99] showed that the use of TPT distributions in GI/M/1 models shows a strong impact on performance (compared to the M/M/1 model) for $\rho \to 1$. More research about N-Burst/M/1 queues with utilization $\rho \to 1$ is necessary, but such scenarios are more of theoretical interest (networks are not designed to work at ρ close to 1), and therefore they are not investigated further here.

5.8 Blow-up Points for Inhomogeneous Bursts

So far, we investigated the steady-state performance behavior of N-Burst/M/1 queues for the Independent Source [IS] model, in which all N ON/OFF sources are identical and in addition only one value of the cell-rate during bursts (λ_p) occurs. In this section, we present preliminary results of how the performance results extend to traffic mixes with heterogeneous sources. Two different representations of traffic mixes in the N-Burst model were introduced in Section 2.5.2: Each of the N sources could be different, i.e. it generates traffic with different parameters (distribution of ON periods, peak-rates etc.). Alternatively, all N sources can be identical, and each source generates the same traffic mix. We will now briefly look at the steady-state performance behavior of both types of models. We restrict our investigations to the location of the blow-up points. The blow-up effects happen when long-lasting over-saturation periods can occur at the switch and they are a consequence of the fact that a discrete number i_0 of sources has to be in a long ON phase in order to create such an over-saturation period.

Mixture of TPT and exponential sources

The first interesting scenario is the case of a mixture of N_{TPT} sources with (truncated) Power-Tailed ON periods and N_{EXP} sources with exponential (or other well-behaved) ON periods. That model is a heterogeneous $(N_{TPT} + N_{EXP})$-Burst. What behavior can we expect in terms of the blow-up points for this model?

Since only the N_{TPT} Power-Tailed sources can be in long ON periods, only they can contribute to long over-saturation periods. The other N_{EXP} sources always contribute their average rate κ. Consequently, that heterogeneous $(N_{TPT} + N_{EXP})$-Burst model behaves in terms of the location of the blow-up points like an N_{TPT}-Burst with increased background Poisson rate $\lambda_0 + \kappa \cdot N_{EXP}$. Hence only N_{TPT} blow-up points exist and their location is obtained through the condition $\nu = \Lambda_i$, with

$$\Lambda_i = \lambda_0 + \kappa \cdot N_{EXP} + (N_{TPT} - i)\,\kappa + i\lambda_p, \quad i = 1, ..., N_{TPT}.$$

The fact that exponential ON/OFF sources only contribute their mean rate κ to the performance behavior is proven for the fluid-flow ON/OFF model class with $i_0 = 1$ in [Jelenkovic & Lazar 97].

TPT bursts with different peak-rates

Now we consider N-Burst models that allow for different peak-rates during the bursts by either of the two approaches mentioned above and described in more detail in Sect. 2.5.2. According to the discussion in the previous paragraph, we only have to take bursts into account that can show very long durations (e.g. when their duration is TPT distributed).

However, TPT distributed burst-lengths are only one possibility to obtain the blow-up effects. Much simpler distributions can be used instead, as long as they allow for very long bursts to occur. In the following, we use HYP-2 distributions with large variance, in fact we fit their first three moments to a TPT distribution with large PT Range. An N-Burst model with such HYP-2 distributed burst-lengths shows the blow-up effects at exactly the same location as the model with TPT burst-lengths, since the location only depends on Eq. (5.1). However, all the other 'nice' properties, such as LRD, Power-Tailed queue-length distribution with changing exponent, and the Power-Law growth of mCD with MBS do not hold for HYP-2 distributed burst-lengths. But we are not interested in those properties in this section.

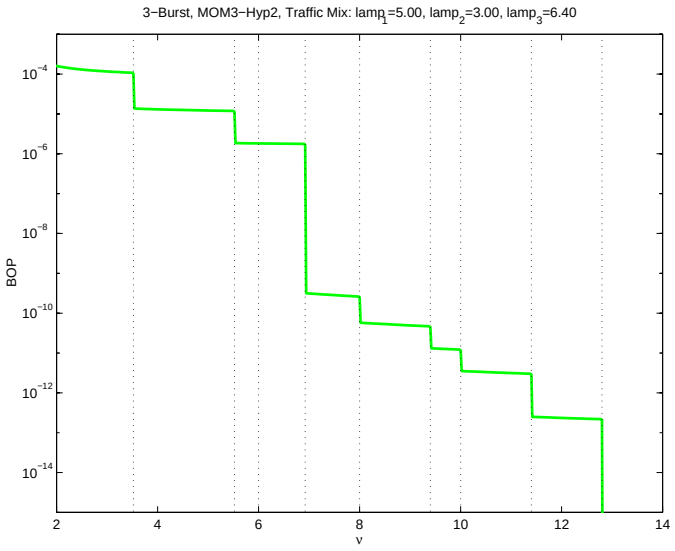

Figure 5.13: **Blow-up Points for Two Identical Sources with Inhomogeneous Bursts:** each source generates a traffic mixes with bursts of three different traffic classes (with different cell-rates, $\lambda_p^{(i)}$). The BOP shows blow-up points at various places, $\Lambda_i^{(j)}$. However, not all possible combinations of cell-rates – such as the third marker from the left – lead to blow-up effects, see text.

Figure 5.13 and Figure 5.14 show calculations of the BOP for both scenarios of heterogeneous bursts respectively sources. The blow-up regions can be clearly observed, yet the number of blow-up points has increased in comparison to the standard case with bursts of only one type. This is not surprising: the radically changing performance behavior is again due to the over-saturation by a number of long-term active bursts. However, due to the increasing possibilities for the intra-burst rates, there is now more than one cell-rate (used to be called Λ_i) which results from a certain number, i, of long-term active bursts. Without going into too much detail in the

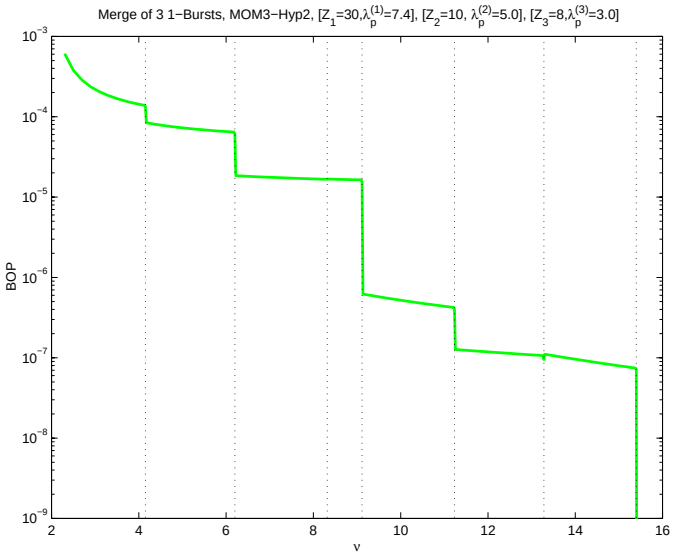

Figure 5.14: **BLow-up Points for Scenarios with Different Sources:** the performance shows the same behavior as in Figure 5.13, yet the number of possible combinations of cell-rates decreases. As before, not all the potential cell-rates lead to blow-up effects, see text.

following, these critical cell-rates will now be called

$$\Lambda_i^{(j)} := \text{Different possibilities for Cell-Rate with } i \text{ long-term active sources,}$$

where the superscript $j = 1, 2, ...$ expresses the different possibilities due to the different burst/source types.

According to Figure 5.13 and Figure 5.14 and a set of further calculations, the following observation seems to hold:

- Potential candidates for blow-up points are described by the condition $\nu = \Lambda_i^{(j)}$.
- Out of those, blow-ups do not happen at $\Lambda_{i+1}^{(j)}$ unless

$$\Lambda_{i+1}^{(j)} > \Lambda_i^{(k)} \quad \text{for all } k,$$

  i.e. the transition to blow-up region $i+1$ for a particular combination of cell-rates only shows up with blow-up effects, when the system is already in blow-up region i for **all** other possible combinations of cell-rates.

These two observations locate the blow-up points. Already that location can be used for network design, see Chapter 7. However, a deeper understanding of the impact of traffic mixes requires that all investigations of the previous chapters be redone using heterogeneous bursts. Also, additional questions, such as the impact of the proportions of the individual burst-types and the implications of Power-Tailed bursts with different tail-exponent, α, need to be investigated. The intention of this section is to provide some preliminary results for traffic mixes. The in-depth discussion is left open for future research.

5.9 Summary

In contrast to standard bursty models (with exponentially decaying probability of large bursts), N-Burst models with LRD properties through TPT distributed burst-lengths show a number of so-called blow-up regions with dramatic changes in performance (as measured by mCD, BOP, and CLP) at their boundaries, called blow-up points.

Those blow-up points are located by Eq. (5.1), and already the knowledge about their location can be useful for network design purposes.

The reason for the blow-up effects is that N-Burst models with TPT distributed burst-lengths can generate long over-saturation periods during which the mean cell-rate is temporarily higher than the service rate. In blow-up region i_0 (see Eq. (5.2)), i_0 long-term active sources are required to create such a long over-saturation period. It is shown in Sect. 5.2 that such over-saturation periods have a Power-Tailed duration with exponent $\beta = i_0(\alpha - 1) + 1$, where α is the PT exponent of the burst-length distribution.

The most important result of this chapter is that it is the exponent β that matters for performance, and not solely α, which determines the degree of LRD of the traffic. Consequently, the performance impact of such LRD traffic can be alleviated by network parameters.

The queue-length distribution of N-Burst/M/1 queues turns out to be Power-Tailed with exponent $\beta - 1$. The fact that the exponent β changes discontinuously at the blow-up points is the cause of the blow-up effects.

In addition to the knowledge about the tail-exponent $\beta - 1$, approximations for the tail-constant of the queue-length distributions are discussed in Sect. 5.2. In comparison with the exact numerical performance results for N-Burst/M/1 queues, those approximations turned out to be useful to obtain quantitative results for the performance parameters in large-buffer systems from simple, closed-form formulas (without matrix computations as necessary for the exact numerical results).

The use of truncated tails, or equivalently a Maximum Burst Size (MBS), results in a truncation of the queue-length distribution. A formula for the PT-Range of the queue-length distribution is obtained in Sect. 5.3. In the scenario $\beta < 2$, when the expected value of the Cell Delay is infinite for full PTs, a Power-Law relationship between mCD and MBS is obtained:

$$\mathrm{mCD}(\mathrm{MBS}) \sim \mathrm{MBS}^{2-\beta}, \quad \text{when} \quad \beta < 2.$$

The numerically computed values for mCD show that such Power-Laws can be observed for a realistic range of MBS values.

As a consequence of the Power-Tailed queue-length distribution, the BOP drops off slowly as $\mathrm{BOP}(B) \sim B^{\beta-1}$ with increasing buffer-size B. Thus, additional buffer-space is an ineffective method to reduce high overflow probabilities – even for $\beta > 2$. The CLP in a finite-buffer model shows a Power-Law behavior with the same exponent, thus the same conclusions hold. A relationship between the tail-constants c_{BOP} and c_{CLP} is given in Sect. 5.5, but the derivation of this relationship requires results of the transient analysis, therefore it is delayed until Chapter 6.

It is shown in Sect. 5.6 that the M/M/1 queue and a bulk-arrival queue can be seen as limit models of N-Burst queues with $b \to 0$ respectively $b \to 1$. The simple closed-form expressions for various performance parameters of these limit models can thus be used as boundaries for N-Burst/M/1 performance. A short discussion of the fluid-flow limit is presented as well, more details can be found in [Lipsky & Schwefel 00].

Finally, the claim that the distribution of the OFF period is not critical for performance is confirmed in Sect. 5.7, and preliminary results about the location of the blow-up points for inhomogeneous bursts are presented in Sect. 5.8.

To conclude the steady-state analysis, we summarize the practical consequences for the network design process in the scenario of LRD ON/OFF traffic in the following:

Only in case of Peak-Rate Allocation, when the service rate of the switch is higher than the peak-rates of all sources together, can the bad effects of such traffic be avoided. However, Peak-Rate Allocation usually leads to an unnecessary waste of network capacity, since for badly behaved sources the network would have to be run at utterly low utilization, e.g. less than 10%, which is usually hardly justifiable to the financial departments, since capacity costs money.

However, adding more capacity is only one of four possibilities to alleviate the negative effects of the Power-Tailed distributed bursts. Summarizing the results of this chapter, the remedies are:

1. **Limit the Maximum Burst Sizes**: The LRD properties are thereby truncated, and the PT Range of the buffer-occupancy distribution can be controlled via 5.14. See Section 5.3 for details.

2. **Add Switch Capacity,** ν: If enough capacity is added, a better blow-up region (higher i_0) can be achieved, and the PT exponent β that determines the performance behavior is increased, see Eq. 5.4.

3. **Higher Degree of Multiplexing**: If the traffic of a larger number of sources is multiplexed while keeping the utilization ρ constant, it follows from Eq. (5.2) that a better blow-up region can be the result. Again, the PT exponent β of the queue-length distribution is increased thereby. However, this requires that the service-rate ν is increased proportionally to N.

4. **Traffic Shaping**: By decreasing the peak-rate λ_p during bursts as done in Sect. 5.6, the location of the blow-up points Λ_i is shifted, and eventually a better blow-up region is the result. However, note that there can arise performance problems (Loss, Delay) at the traffic shaper itself[5]

These remedies improve QoS already for standard, non LRD traffic. However, since they can achieve a better blow-up region for LRD N-Burst traffic, the improvements can be very drastic for such traffic. On the other hand, neglecting the results of this chapter can be fatal, since very small changes in the network (e.g. adding another source) can totally mess up the network's performance when such changes lead to a transition to a worse blow-up region than before.

Also, an important practical result of this section is that the utilization by itself is not sufficient for QoS oriented network planning, since the blow-up effects can occur for very low utilization values: a properly designed network can run at $\rho = 0.8$ and show a better performance as a poorly planned network that runs at utilization $\rho = 0.2$, if the latter is operating in the worst blow-up region, $i_0 = 1$.

[5]If the traffic is described by a 1-Burst model, the shaper can be modeled by a 1-Burst/M/1 queue that necessarily operates in $i_0 = 1$, since its purpose is to reduce the peak cell-rate of the traffic stream, i.e. $\nu(\text{shaper}) < \lambda_p$.

Chapter 6

Transient Performance Analysis

In the previous chapter, the impact of ON/OFF traffic with LRD properties on the steady state-values for the QoS parameters *mean Cell Delay* (mCD) , *Cell Loss Probability* (CLP), and *Buffer Overflow Probability* (BOP) was investigated. The analysis showed that QoS does not reduce uniformly with, for instance, decreasing service rate. Instead, well-defined so-called *blow-up points* have been observed at which the QoS parameters deteriorate dramatically.

The analysis in Chapter 5 was based on steady-state values for the QoS parameters. Steady state is the limit of the *distribution* of some parameter when the influence of the initial state has disappeared.[1] An actual measured system reflects exactly one possible state at a particular instant in time, and as such the real system itself can never *be* in steady-state. However, when the number of samples is increased, a collection of observations is expected to converge stochastically to the steady-state distribution, provided that it exists. Therefore, the system never reaches steady-state, but rather does the growing collection of samples eventually represent the steady-state distribution closely enough. Note that we only deal with *ergodic* systems in this thesis, and the investigations in this chapter implicitly assume ergodicity, even when put in a general context.

The question arises, how large does the observation period have to be to yield a reasonable representation of steady-state behavior? In most practical applications of stochastic models, such a question is not asked and steady-state results are applied without further reasoning. As we will see later, this is problematic for our LRD network traffic models: the large variability in the traffic together with the long-range correlations significantly reduce the speed of convergence towards steady-state. Since the critical high-utilization periods for telecommunication networks are typically limited to 5-8 hours per day, some QoS parameters might never come close to their steady-state values within that time-interval. In such cases, steady-state analysis is not sufficient, but it has to be complemented or even replaced by *transient* analysis.

For individual data-connections or high-resolution measurements, smaller time-scales in the range of minutes or seconds are of interest. Clearly, the steady-state values of the QoS parameters will not be observed in such short time intervals. In fact, the results of the individual measurements will fluctuate greatly. A network engineer has to be aware of those fluctuations to be able to interpret his measurement results correctly. Such insight can only be obtained from transient analysis.

One of the striking results in the steady-state analysis is that N-Burst/M/1 queues with infinite PT Range of their burst-length distribution show an infinite mean cell delay in blow-up region i_0 when $\alpha < 1 + 1/i_0$. An infinite mean delay can never be observed during the finite (stationary) life-time of a telecommunication systems. Therefore, the transient analysis in this chapter starts off in Sect. 6.1 by an investigation of the mean cell delay in finite observation intervals. The

[1] In some models, the steady-state limit can depend on the initial state. However, this is not the case for the models that we investigate.

analysis in Sect. 6.1 is based on simulation models as well as on approximating M/G/1 queues, as introduced in Sect. 2.7.4.

One particular property of both delay and overflow-events in our queueing models is that the individual events are highly correlated. Once the buffer has filled up, a substantial number of subsequent cells experiences high delay or causes losses/overflows. This is true for any queueing model but it is accentuated greatly by the LRD properties in the cell-arrival process. Therefore, it is very relevant to investigate the time until the buffer fills up for the first time, called *First Passage Time* (FPT), which is discussed in detail in Section 6.2. Due to the correlated nature of overflow events, scenarios are possible in which many observation intervals show no overflow events at all, while a large fraction of cells cause overflows in the few remaining observation intervals that have at least one overflow. The probability $\gamma(t_0, B)$ that there is at least one overflow in an observation period of duration t_0 can be derived from the FPT distribution, see Sect. 6.3 for a discussion. The fraction of cells, called *conditional Buffer-Overflow Ratio* (*conditional Cell-Loss Ratio*), that causes overflows (losses) in the intervals with at least one such event is investigated in Sect. 6.4. The LRD properties in the N-Burst traffic cause a very peculiar behavior of those conditional, transient performance parameters, namely they grow asymptotically with increasing buffer-size. The transient overflow probability, γ, together with the conditional transient overflow/loss ratios provide a much better description of the performance behavior than a steady-state BOP or CLP, since the transient parameter pair is able to express the fluctuations that occur in different observation intervals.

Frequently, an important part of the transient analysis is the analysis of the so-called *Busy Period*, i.e. the time interval that starts when a customer enters an empty queue, and ends as soon as the switch becomes idle again for the first time. Since the practical application of results from the Busy-Period analysis requires a careful interpretation (the Busy Periods are correlated due to the memory – the number of active bursts – in the N-Burst arrival process), we do not present a discussion of the Busy Period analysis in this thesis. See [Schwefel 99e] for some preliminary results, and Chapter 10 for a discussion of the relevance of future research in this area.

Investigating the transient behavior by simulations can be very costly. In addition, the fluctuations in the simulation results can be large when LRD arrival processes are involved. On the other hand, the transient analysis of the analytic models (see Appendix F.3 and Appendix H for algorithms) might be costly in terms of computing time as well, but the results are exact and no fluctuations have to be dealt with. Therefore, much more insight can be expected when using both methods together.

Although the transient analysis of even simple systems such as the M/M/1 queues is still a current research topic (see e.g. [Leguesdron et al. 93] and [Coevering 95]), the growing interest in the application of analytic models to telecommunications systems has produced the need for more complicated models: [Ren & Kobayashi 95] looked at the transient queue-length probabilities in a fluid flow model with N exponential ON/OFF sources. [Tanaka et al. 95] and [Sericola 98] also use fluid queues, but with more general, Markov modulated arrival rates. Finally, [Kulkarni & Li 98] investigate the transient behavior of queues with a particular MMPP arrival process, but still restricted to exponential state-times. So far, little work was done with respect to the transient behavior for arrival processes with LRD properties.

Most examples in this chapter use N-Burst [IS] models with the individual source parameters chosen in correspondence to an early parameter estimation for measurement `TX3` in [Schwefel et al. 97]:

$$\kappa = 16.3\,\text{cells/ms}, \quad \overline{n_p} = 9.1, \quad b = 0.88\,.$$

The improved calibration method in Section 4.3 leads to slightly different results, but the differences are of negligible magnitude. From these parameters, a peak-rate $\lambda_p = 140\,\text{cells/ms}$ and a mean OFF time of $Z = 0.49\,\text{ms}$ follow. The background Poisson rate is always assumed to be $\lambda_0 = 0$. This parameter set is referred to as `TX3` source in this and the two following chapters, although it slightly deviates from the `TX3` parameter set of Section 4.3.

During a busy hours period of $t_0 = 5\,\text{h}$, such a `TX3` source generates on average about $3 \cdot 10^8$ cells or about 16 GigaBytes (GB) of data. Medium time-scales such as $t_0 = 5\,\text{min}$, as for example for longer individual connections, correspond to about 5 million cells ($\approx 270\,\text{MB}$) per `TX3` source. Finally, on time-scales of short connections, such as $t_0 = 5\,\text{s}$, about 80 000 cells ($\approx 4.5\,\text{MB}$) are generated on average by one such source.

6.1 Transient Mean Delay

As a first approach to the transient analysis, we investigate the mean cell delay that is observed in a finite observation period of duration t_0. The scenario in which the steady-state mCD is infinite is of special interest here since, obviously, that steady-state value can never be observed. First, we use a simulation model to illustrate how problematic the interpretation of individual measured traces can be. Thereafter, we discuss the transient mean queue-length – as a related performance measure to mean delay – in approximating M/TPT/1 models.

6.1.1 Simulation Experiments

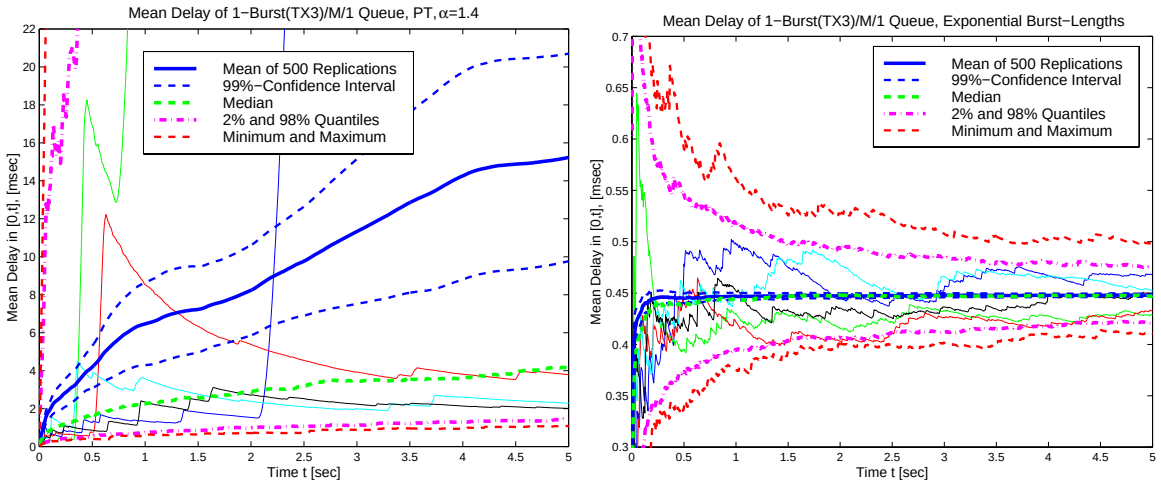

Figure 6.1: **Simulated Mean Cell Delay in** $[0, t_0]$**:** Large fluctuations can be observed in the individual simulation runs, five of them plotted as thin lines in the left graph, which uses the LRD 1-Burst model in blow-up region $i_0 = 1$. Although both graphs are generated from simulation experiments of 1-Burst/M/1 Queues with $\rho = 0.5$, the exponential burst-lengths that are used for the right graph cause much smoother behavior. Note that the region with delay values below 0.3 msec is suppressed in the right graph. See text for a detailed discussion of the graphs.

Figure 6.1 presents the results of two simulation experiments illustrating clearly that although the use of steady-state results can be very reasonable for traditional traffic models, this is not necessarily true for traffic from the LRD version of the N-Burst models. Both graphs show the mean cell delay $\text{mCD}(t)$ averaged over all cells that left the switch up to time t. The arrival

process is in both graphs a 1-Burst model with parameters according to measurement `TX3`, but exponential ON periods are used in the right graph as opposed to PT distributed ON periods in the left graph. The steady-state utilization of the 1-Burst/M/1 queues is in both cases $\rho = 0.5$.

The steady-state value of the mCD for the exponential 1-Burst model is in fact meaningful in the finite observation intervals, since the range of observed mCD(t) values in 500 replications converges rather quickly towards the steady-state value. For the average and the median of the 500 replications, convergence is even quicker.

The observations are totally different for the 1-Burst model with LRD properties in the left graph of Fig. 6.1: Although about 80000 cells contribute to mCD(t) after 5 seconds, the observed value fluctuates greatly from one replication to the other. Within an individual replication, huge jumps of the mean delay occur, which implies that a large number of subsequent cells experience huge delay values. Even when averaging over all 500 replications, no stability is observed and the resulting confidence intervals for the true mean are huge.

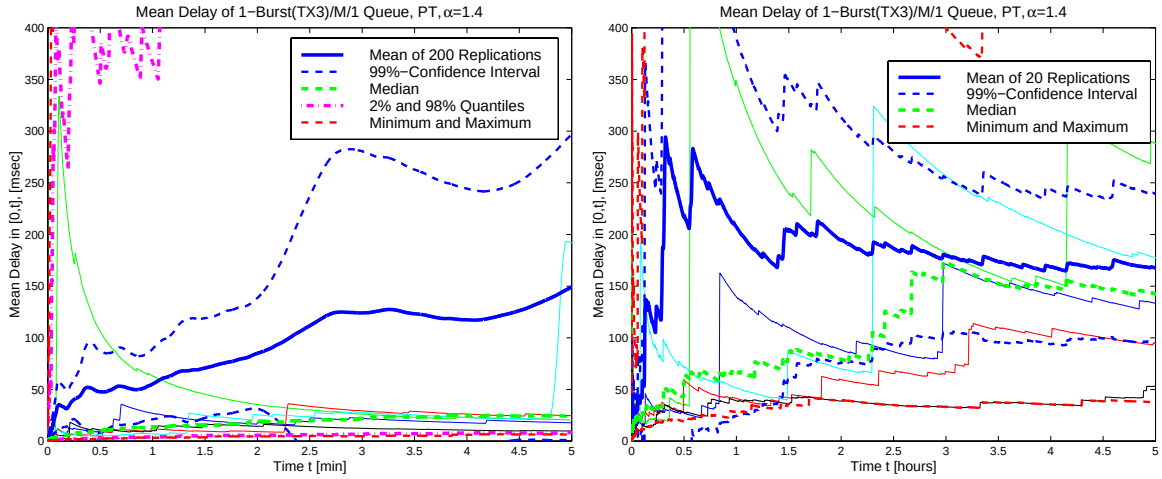

Figure 6.2: **Simulated mean Delay for 1-Burst/M/1 Queue, Medium and Large Time-Scales:** Even after 5 hours of simulated time in the right graph, large fluctuations can be observed. The mean of 20 replications seems to converge in the right graph. However, this is only due to the limited number of replications: as a consequence of the infinite steady-state mCD, the true transient mean delay mCD$((t)$ grows unboundedly with t.

Even when we move to larger time-scales of $t_0 = 5\,\text{min}$ and $t_0 = 5\,\text{h}$ no stability is observed in the delay values for the LRD 1-Burst model, see Fig. 6.2. Sharp increases of mCD(t) are still observed in the individual replications after a few hours of simulation time, although several hundred million cells contribute their individual delay value to that plotted average. Of course, the mCD in blow-up region i_0 for $\alpha < 1 + 1/i_0$ is a particular bad parameter to observe, since its steady-state value is infinite. Other parameters of the delay-distribution, such as the median, behave better.

In summary, measurements of transient performance parameters are very well behaved only for N-Burst/M/1 queues whose traffic does *not* have LRD properties, or for the scenario of Peak-rate Allocation ($\nu > N\lambda_p$). In such systems, fluctuations diminish after a few seconds for the transient mean delay and other transient performance parameters. Consequently, steady-state behavior can be observed very quickly.

In contrast to the exponential case, the observations are totally different for LRD traffic as well as for actual measured traces (see [Schwefel 99e]). Fluctuations in individual replications are huge and they do not disappear in the mCD(t) estimators, even when taking into account the

whole busy period of 5 hours. Even worse, the fluctuations do not disappear when averaging over a large number of replications. Confidence intervals for the true expected value over all replications are rather large and thus the estimated values have to be interpreted with great care.

6.1.2 Analytic Approximating Models

Due to the large fluctuations in the observed performance parameters, simulation experiments are not really a very suitable method to obtain reliable values. Their main purpose in the previous section is to illustrate those fluctuations rather than to obtain actual numbers. Numerical results from analytic models are not subject to such fluctuations. Instead of analyzing the transient mean cell delay, we look at the transient queue-length distribution in this section.

It is possible to compute the distribution of the queue-length Q_t at time t for an N-Burst/M/1 queue (see (24) in [Lucantoni 93]), yet it involves inverting a Laplace transform and a z-transform of a function containing matrix expressions. Therefore, it is rather complicated and computationally demanding.

Instead, we analyze the approximating M/TPT/1 queue as described in Sect. 2.7.4. Note, that the M/G/1 model approximates the N-Burst/M/1 queue at burst-level. Therefore, the cell-arrivals are not spread out over time, which is an approximation of the N-Burst/M/1 queue with $b \to 1$, i.e. automatically $i_0 = 1$. Thus, the steady-state mean queue-length of the $M/G/1$ queue is infinite for full Power-Tails with $\alpha \leq 2$.

Herein, we focus on the expected value, $\mathbb{E}\{Q_t\}$, of the queue-length at time t, given that the queue was empty at time 0. Note that this does not involve averaging over the observation period as for the cell delay in the simulation model of the previous section. We look at time-averages later on in this section.

The Laplace Transform of the mean queue-length of the M/G/1 queue with arrival rate λ and mean service time, $\overline{x}$, is given in [Cohen, 82], Equation (4.52):

$$\hat{f}(s) = \int_0^\infty e^{-st}\, \mathbb{E}\{Q_t | Q_0 = 0\}\, dt = \frac{\lambda}{s^2} - \frac{1-\mu(s,1)}{s\,[1+s/\lambda-\mu(s,1)]} \cdot \frac{\beta(s)}{1-\beta(s)}, \quad \mathrm{Re}(s) > 0,$$

where $\mu(s,1)$ is the root z of

$$z - \beta\,[s + \lambda(1-z)] = 0, \tag{6.1}$$

with smallest absolute value. $\beta(s)$ is the Laplace Transform of the service-time distribution. If that distribution has an ME representation $\langle \mathbf{p}, \mathbf{B} \rangle$ (see Appendix B.1), its Laplace Transform can be computed by $\beta(s) = \mathbf{pB}(s\mathbf{I} + \mathbf{B})^{-1}\boldsymbol{\varepsilon}'$, see [Lipsky 92]. The roots of Eq. (6.1) can be determined numerically in a number of ways. We used a variation of Newton's algorithm in the following. The numerical inversion of the Laplace transform is described in Appendix H.5.

Figure 6.3 shows the curves of the expected queue-lengths of an M/TPT/1 queue that has features similar to the 1-Burst/M/1 model, whose arrival process is calibrated on measurement `TX3`, and whose utilization is $\rho = 0.5$. Curves are plotted for different values of the truncation, T. The steady-state mean queue-length and the time to reach 90% of that value are marked by dotted lines for each T. The truncations correspond to Maximum Burst-Sizes and Power-Tail Ranges of the service-times of the M/TPT/1 queue as follows:

Truncation T	1	5	10	15	20	30	37
MBS [cells]	9.1	36.5	326	$3.5 \cdot 10^3$	$4.1 \cdot 10^4$	$5.6 \cdot 10^6$	$1.8 \cdot 10^8$
PT-Range (ST)	0.28 ms	1.1 ms	10 ms	0.11 s	1.2 s	173 s	1.5 h

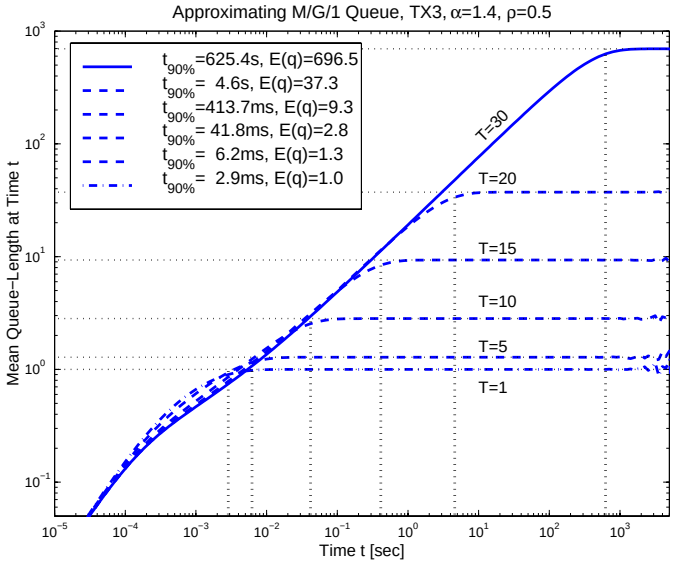

Figure 6.3: **Mean Queue-Length at Time t of an Approximating M/TPT/1 Queue:** The dotted horizontal lines show the steady-state value of the mean queue-length for the corresponding truncation T. The dotted vertical lines mark the time, at which the mean queue-length has reached 90% of its steady-state value. For large T, $\mathbb{E}\{Q_t\}$ appears linear on log-log scale for some range. The slope of the linear part turns out to be $0.6 = 2 - \alpha$.

The higher the value of T, the larger the steady-state mean queue-length becomes, but it also takes a longer time to converge to the steady-state value, see the values in Figure 6.3. For $T = 37$ it takes approximately a period of $t_0 = 5$ h (about the whole stationary, daily high-load period) to come close to the steady-state value.

For large T, the curves show a linear appearing segment before they come close to their steady state-value. Measurements of the slope on log-log scale yield a value of $2 - \alpha$. If this is true in genera;, then it follows for the behavior of the mean queue-length:

$$\mathbb{E}\{Q_t\} \sim t^{2-\alpha} \quad \text{for} \quad \alpha < 2.$$

Another very important conclusion can be drawn from the observation of $\mathbb{E}\{Q_t\}$ in Fig. 6.3: In a limited time-horizon, t_0, there is some upper limit on the truncation value (equivalently the MBS) of the burst-length distribution, beyond which practically no further impact on the expected value, $\mathbb{E}\{Q_t\}$, of the queue-length distribution can be observed until time t_0. In other words, $\mathbb{E}\{Q_{t_0}\}$ (MBS) converges, and the smaller t_0 is, the earlier does the convergence occur.

Note that the individual customers in the approximating M/TPT/1 queue represent whole bursts rather than single cells. Nevertheless, similar observations can be expected for the queue-length of the full N-Burst/M/1 model.

Time average of queue-length

In practice, performance evaluation by measurements or simulations very frequently does not look at the queue-length at a specific time t, but more often at the time-average,

$$\bar{Q}_t = \frac{1}{t} \int_0^t Q_{t'}\, dt'.$$

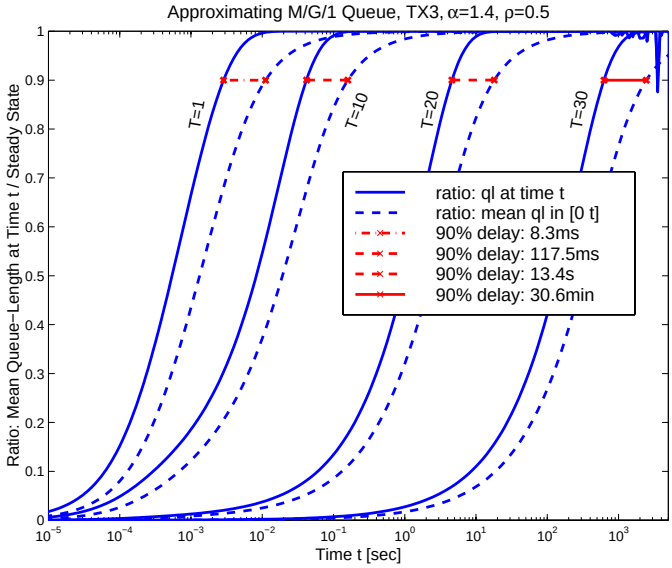

Figure 6.4: **Ratio to Steady-State Value of Mean Queue-Length:** The expected time-average of the queue-length (dashed lines) approaches its steady-state value slower than the solid lines of $\mathbb{E}\{Q_t\}$. The wiggly lines in the upper right corner are due to numerical instabilities.

The expected value of the time-average can be calculated analytically by switching the order of the integrals – which is allowed under very loose mathematical conditions:

$$\mathbb{E}\{\bar{Q}_t\} = \mathbb{E}\left\{\frac{1}{t}\int_0^t Q_{t'}\,dt'\right\} = \frac{1}{t}\int_0^t \mathbb{E}\{Q_{t'}\}\,dt'.$$

Therefore, the time average of the queue-length is subject to the same Power-Law behavior

$$\mathbb{E}\{\bar{Q}_t\} \sim t^{2-\alpha} \quad \text{for} \quad \alpha < 2,$$

but as Figure 6.4 shows, it converges slower to the steady-state mean queue-length than $\mathbb{E}\{Q_t\}$.

The convergence towards a finite steady-state value is due to the truncation of the Power-Tails. On those time-scales, the classical theory for Markovian Processes can be applied. As it is stated e.g. in [Lipsky 92], transient state probabilities approach their steady-state value exponentially fast as $e^{-t/R}$ (R is called *Relaxation Time*), while time averages approach the steady-state value only as $1/t$, which is much slower. In the light of those results, the shift of the time average curves in Fig. 6.4 is not surprising.

The important conclusion for measurements is that the commonly done averaging process over sequential samples of a measurement requires a much longer measurement period to become stable. For the M/G/1 approximation with $T = 30$ in Fig. 6.4, the measurement time would have to cover 41 minutes instead of 10 minutes, to reach 90% of the expected steady-state value.

6.2 First Passage Process

So far in this chapter, we investigated the transient mean queue-length, which is closely related to the transient mean delay. Long delay times are caused by long queue-lengths, and since

long queues only drain gradually, it is obvious that the cell delay of subsequent cells is highly correlated. This is even more obvious for Buffer-Overflow events: once the buffer has filled up beyond the level B, a number of subsequent cells overflow. In order to avoid overflows completely, the very first overflow event is of critical importance. The random variable that expresses the time until that first overflow occurs is called *First Passage Time* and it is discussed in detail in this section.

6.2.1 Motivation & Definition

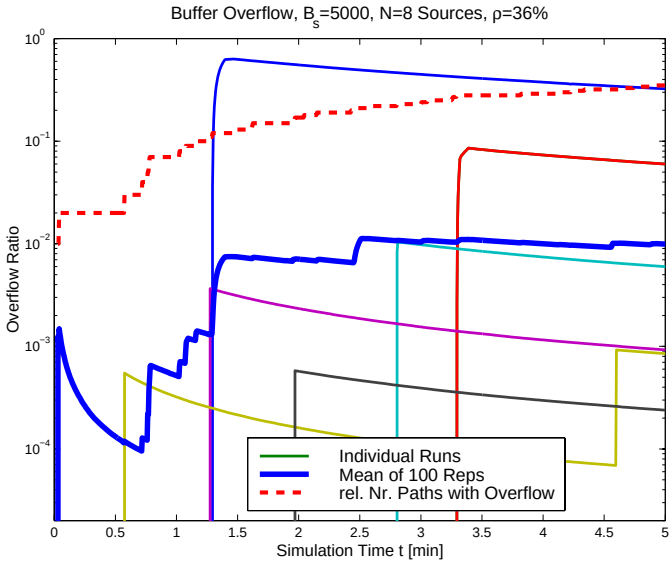

Figure 6.5: **Simulated Fraction of Overflowed Cells, BOR(t, B), for an 8-Burst/M/1 Queue with $i_0 = 2$ and $B = 5000$:** Also shown is the average of all the 100 independent replications (thick line), and the fraction of replications that showed at least one overflow by time t (dashed line).

A simulation similar to that performed for the transient mean cell delay in Sect. 6.1.1 is shown for the Buffer-Overflow Probability in Fig. 6.5: an 8-Burst/M/1 model is simulated for a finite time t_0, starting with an initially empty queue and the beginning of the first ON-period at $t = 0$. During the simulation run, the fraction of overflowed cells is observed, called *Buffer Overflow Ratio*:

$$\mathrm{BOR}(t, B) := \frac{\text{number of overflowed cells in [0,t]}}{\text{number of arriving cells in [0,t]}}.$$

Each run terminates after $t_0 = 5$ minutes during which – on average – about 40 million cells arrive. The curves for the individual simulation runs show that the overflow events are highly correlated: About 64% of the 100 simulation runs did not show any overflows at all, but if overflows occurred, then a large number of cells overflowed. For $\rho < 1$, the BOR converges stochastically to the steady-state BOP for long observation periods, t_0:

$$\mathrm{BOR}(t_0, B) \overset{t_0 \to \infty}{\longrightarrow} \mathrm{BOP}(B).$$

However, even though about 40 million arrivals were simulated in each replication in Fig. 6.5, the random variable $\mathrm{BOR}(t_0, B)$ still shows large variance, with the values ranging from 0 to about 35%.

Later in Sects. 6.3 and 6.4 we introduce transient performance parameters that are able to describe such fluctuating behavior. Those parameters are based on an investigation of the very

first overflow event. The time at which the buffer occupancy reaches level n for the first time is expressed by the random variable

$$\tau_n := \min_{t>0} \{t | Q_t = n\} \ . \tag{6.2}$$

τ_n is called *First Passage Time* (FPT). Here we always assume an initially empty queue. Note that the FPT is independent of the buffer-model, i.e. there is no difference between the finite-buffer loss model and the infinite backup buffer model, as long as $n \leq B$.

6.2.2 Mean First Passage Times

The FPT is a random variable with some underlying distribution. Before we look more closely at the actual distribution in the next section, we first investigate its expected value, called *mean First Passage Time* (mFPT):

$$\text{mFPT}(B) := \mathbb{E}\{\tau_B\} \ .$$

A Matrix-Analytic algorithm to compute mFPT(B) for MMPP/M/1 queues is presented in Appendix F.3. An alternative method for the computation of mFPT(B) can be obtained via the use of martingales. A summary of that method is presented in Appendix H, see [Asmussen et al. 00a] for the detailed derivation.

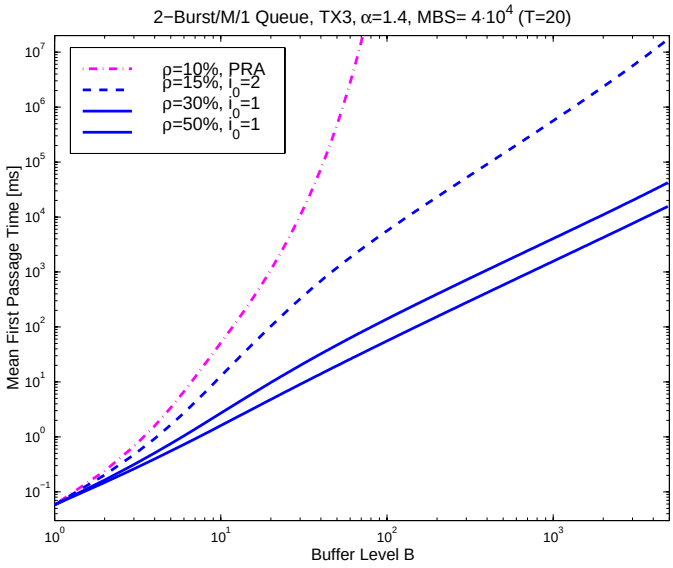

Figure 6.6: **Mean First-Passage Times of 2-Burst/M/1 Queue in Different Blow-Up Regions:** Power-Laws can be observed in the blow-up regions $i_0 = 1$ (solid lines) and $i_0 = 2$ (dashed lines), but not for Peak-Rate Allocation (dash-dot).

The most important result of the steady-state analysis of N-Burst/M/1 queues in the previous chapter is the existence of N blow-up regions, $i_0 = 1, ..., N$ (plus one PRA region), in which the performance behavior is radically different: the steady-state queue-length distribution shows a Power-Tail with exponent $\beta = i_0(\alpha - 1) + 1$, and β changes at the transition from one blow-up region to the other, see Sect. 5.2.

The computation of mFPT(B) for N-Burst models in different blow-up regions and in the PRA region in Fig. 6.6 reveals that the blow-up regions are also of critical impact for the behavior of

$\mathrm{mFPT}(B)$. Within the blow-up regions $i_0 = 1, ..., N$, a Power-Law growth

$$\mathrm{mFPT}(B) \sim c_{\mathrm{mFPT}}\, B^{\beta}\,, \tag{6.3}$$

of $\mathrm{mFPT}(B)$ is observed with the same exponent β that we obtained in the steady-state analysis, see Eq. (5.4). As opposed to the blow-up regions, geometric growth of $\mathrm{mFPT}(B)$ occurs in the PRA region, i.e. when $N\lambda_p < \nu$.

Note that the slower Power-Law growth of the $\mathrm{mFPT}(B)$ in the blow-up regions is another manifestation of the Buffer-Inefficiency that is discussed in Sect. 5.4, but now with respect to a transient performance parameter: additional buffer-space is a very ineffective method to increase the mFPT in order to delay the first overflow-event beyond the time-scales of interest. In the PRA region, as with traditional traffic models without LRD properties, the faster geometric growth of $\mathrm{mFPT}(B)$ leads to much better improvement from the use of bigger buffers.

Later on in Section 6.4, the Power-Law behavior (6.3) is 'proven' (although not rigorously), and a relation for the tail-constant c_{mFPT} is obtained. That derivation is based on the Power-Law behavior (5.19) and (5.20) of the steady-state $\mathrm{BOP}(B)$ and $\mathrm{CLP}(B)$, and on the asymptotic behavior of another transient performance parameter, which is introduced and discussed in Sect. 6.4.

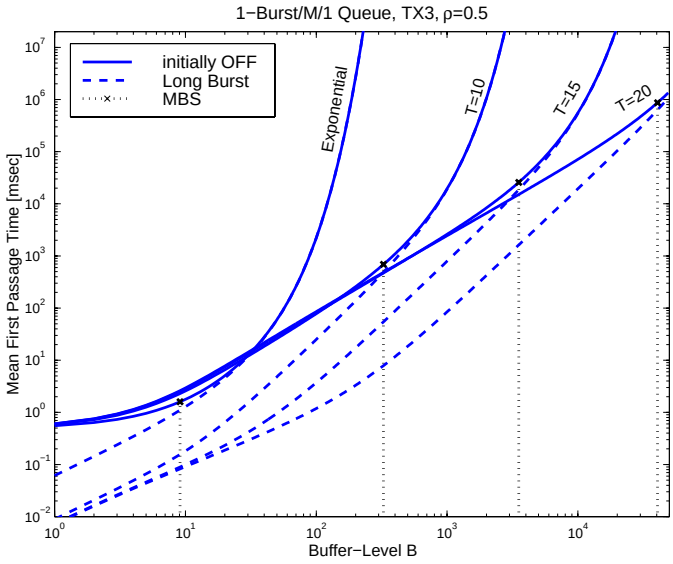

Figure 6.7: **Mean First-Passage Times of 1-Burst/M/1 Queue for different Truncations:** For buffer-levels somewhat bigger than the MBS (marked by dotted lines), the mFPT grows geometrically.

As already observed and discussed in the steady-state analysis of the queue-length distribution in Sect. 5.3, the use of truncated PTs for the ON-time distribution also limits the Power-Law behavior (6.3) to a certain range of buffer-sizes, after which the $\mathrm{mFPT}(B)$ starts to grow geometrically, shown in Figure 6.7. Since we introduced the Power-Tail Range in Sect. 3.5 for complementary distribution functions (which are monotonically decaying), that PT-Range definition cannot be applied without modifications to the function $\mathrm{mFPT}(B)$. In Sect. 6.3, we introduce a transient overflow-probability that is amenable to an analysis of its PT-Range. A brief discussion of the PT-Range is provided in that section.

Figure 6.7 also illustrates the impact of the initial condition on the behavior of $\mathrm{mFPT}(B)$. It is always assumed that the buffer-occupancy of the N-Burst/M/1 queue is zero at the beginning

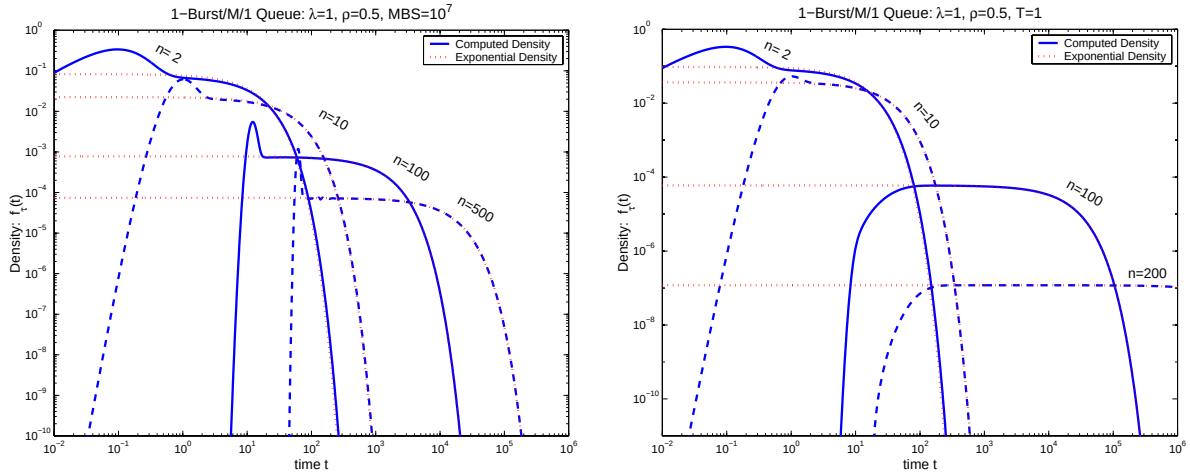

Figure 6.8: **The Density Function of the Distribution of τ_n for a 1-Burst/M/1 Queue with $\rho = 0.5$ ($b = 0.9$, $i_0 = 1$):** Truncated Power-Tails with MBS $= 10^7$ are used for the burst-length distribution in the left graph. The right graph shows the same model with exponential burst-length distribution. For a comparison with the M/M/1 queue, see the left graph of Figure A.2 in Appendix A.2.

of the observation interval. However, the initial state of the N-Burst arrival process[2] also affects the observed behavior. See [Schwefel 99e] for a more detailed discussion.

6.2.3 Distribution of First Passage Times

So far, we have investigated the behavior of the expected value of the FPT. In order to obtain transient overflow probabilities, as they are introduced and discussed in the next section, we need to know the whole distribution of the FPT, τ_n. The method that is developed in [Asmussen et al. 00a] and summarized in Appendix H allows the Laplace transform of the distribution of the FPT to be computed. Via numerical inversion of the Laplace Transform (see Appendix H.5), exact values for the density $f_{\tau_n}(t)$ can be obtained.

Figure 6.8 shows the log-log plots of computed densities $f_{\tau_n}(t)$ for 1-Burst/M/1 queues, both with exponential burst-lengths (right graph) and with TPT distributed burst-lengths (left graph). Except for small t, the density functions are very close to an exponential distribution with same mean mFPT(n). Therefore, an exponential distribution with mean mFPT(n) can be used as an approximation for the true $f_{\tau_n}(t)$. The asymptotic theory that is discussed in [Asmussen et al. 00a] confirms the suitability of the exponential approximation and provides rigorous conditions for the convergence of the exact First Passage Probabilities towards the probability that is derived from the exponential approximation. Also, other asymptotic approximations are obtained in [Asmussen et al. 00a]. Appendix H.3 contains an overview.

A noticeable difference in the density functions for the 1-Burst model with TPT distributed ON periods is the pronounced peak that occurs close to time $n/(\lambda_p - \nu)$ $(= n/8$ in the scenario of Figure 6.8), which is the mean time it takes the queue to reach buffer-level n given that the arrival process started with a very long ON period (with duration longer than $n/(\lambda_p - \nu)$). For

[2] The initial state in Fig. 6.7 is the OFF state for the solid curves, and the largest TPT phase of the ON period for the dashed curves.

large $n \geq 100$, no such peak at all occurs for the exponential 1-Burst model. More research is necessary to gain a full understanding of the shape (width and height) of that peak. However, since it can be expected that the shape of the peak depends critically on specific model details (the assumption of Poisson arrivals during ON times), this thesis does not investigate this further.

6.3 Transient Buffer Overflow Probabilities

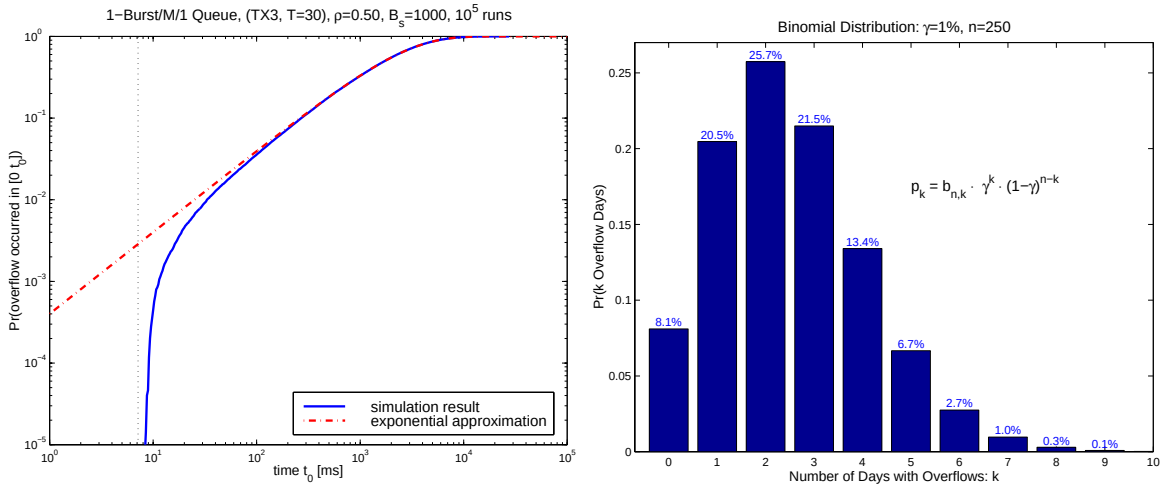

Figure 6.9: **The probability $\gamma(t_0, B)$ and its relevance:** The left graph shows the simulation estimate for $\gamma(t_0, B)$ for a fixed buffer-size $B = 10^3$ cells. When assuming a value of $\gamma = 1\%$ for the whole daily high load period (e.g. $t_0 = 5\,\text{h}$), the number of days in a year during which overflow events occur are distributed binomially as plotted in the right graph.

From the knowledge of the distribution of τ_n, the transient overflow probability

$$\gamma(t_0, B) := \mathbb{P}\left\{\text{At Least One Overflow in } [0, t_0]\right\} = Pr\left(\tau_{B+1} \leq t_0\right) . \tag{6.4}$$

can be derived. As pointed out in the previous section, an exponential approximation can be used for that transient overflow probability in case of large buffers B:

$$\gamma(t_0, B) \approx \mathbb{P}\left\{\text{FPT}(B) \leq t_0\right\} \approx 1 - \exp\left(-\frac{t_0}{\text{mFPT}(B)}\right) . \tag{6.5}$$

Figure 6.9 shows in its left graph simulation estimates for $\gamma(t_0, B = 1000)$ in comparison to the computed exponential approximation from (6.5) for a 1-Burst/M/1 model. For short observation intervals $t_0 < B/\left(\Lambda_{i_0} - \nu\right)$, the exponential approximation over-estimates the true probability $\gamma(t_0, B)$.

The practical relevance of a certain value of e.g. $\gamma = 1\%$ for t_0 as large as the daily busy hours is shown in the right graph of Fig. 6.9: the number of days within a year (assuming the year contains 250 work days), on which at least one overflow-event can be observed, is distributed binomially with parameter γ. In the right graph of Fig. 6.9, the binomial distribution is plotted when assuming $\gamma = 1\%$. In that case, with probability 8.1%, no overflow-days at all can be observed during the cause of a year, while the probability of exactly one overflow-day is 20.5%.

More efficient methods for simulation estimates of small probabilities $\gamma(t_0, B)$ are discussed in [Asmussen et al. 00b]. In that paper, *Importance Sampling* is used to obtain estimators

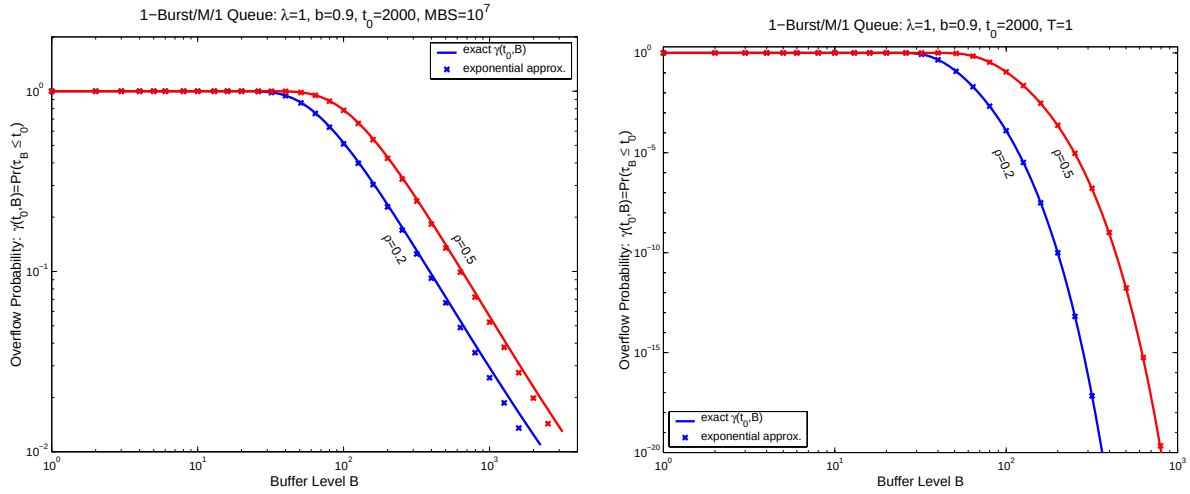

Figure 6.10: **Transient Overflow Probabilities for 1-Burst/M/1 Queue,** $i_0 = 1$**:** For the TPT distributed burst-lengths in the left graph, the overflow probabilities drop off slowly by a Power-Law. In contrast to that, the overflow probability drops off exponentially for exponential burst-lengths in the right graph. Mind the different scale of the y-axis in the two graphs!

with small variance. The numerical computation of exact probabilities $\gamma(t_o, B)$ can also be obtained by the algorithms in Appendix H, since the Laplace transformation of the cumulative distribution function of τ_n follows easily from the Laplace Transformation of the density function, see Appendix H.5.

If the ratio $t_0/\mathrm{mFPT}(B)$ is very small, the exponential function in Eq. (6.5) can be approximated by the linear term of its Taylor series:

$$\gamma(t_0, B) \approx \frac{t_0}{\mathrm{mFPT}(B)} \quad \text{for} \quad \frac{B}{\Lambda_{i_0} - \nu} \ll t_0 \ll \mathrm{mFPT}(B)\,. \tag{6.6}$$

Using the asymptotic behavior (6.3) of the mFPT, it follows that

$$\gamma(t_0, B) \approx \frac{t_0}{c_{\mathrm{mFPT}}} B^{-\beta}, \quad \text{for} \quad 0 \ll \frac{B}{\Lambda_{i_0} - \nu} \ll t_0 \ll \mathrm{mFPT}(B)\,. \tag{6.7}$$

If $B/(\Lambda_{i_0} - \nu)$ approaches or exceeds t_0, the exponential approximation (6.5) becomes invalid. On the other hand B must be large enough for two reasons: first, B needs to be in the range of buffer-size for which Power-Law behavior (6.3) of $\mathrm{mFPT}(B)$ can be observed. Secondly, B needs to be large enough such that $t_0/\mathrm{mFPT}(B)$ is small in order to be able to use the linear Taylor approximation $1 - \exp(-x) \approx x$ for the exponential function. As a consequence of the restrictions on B, in a mathematically rigorous derivation, two simultaneous limits $B \to \infty$ and $t_0 \to \infty$ have to be used in a way that

$$\lim_{B\to\infty} \frac{t_0(B)}{B} = \infty \quad \text{and} \quad \lim_{B\to\infty} \frac{t_0(B)}{B^\beta} = 0\,.$$

We do not investigate such an asymptotic theory here, but instead we compare the approximation (6.7) with the exact numerical results that we obtain from the algorithms in Appendix H.

Figure 6.10 shows the exact values for the overflow-probability $\gamma(t_0, B)$ in comparison to the exponential approximation (6.5) for 1-Burst models with different burst-length distributions. The

Power-Law behavior (6.7) of the transient overflow probability can be observed in the left graph, where MBS $= 10^7$ is used. The exponential approximation provides almost indistinguishable results. Only for larger buffer-sizes $B > 500$, a small deviation from the exact values can be observed for the Power-Tailed model.

The Power-Law behavior (6.7) does not hold for arbitrary large buffers. Eventually, a quick drop-off of $\gamma(t_0, B)$ will occur, which can be due to two reasons, whichever occurs first:

- Truncated Tails: If TPT distributions with a certain MBS are used for the burst-lengths in the N-Burst model, the PT Range of the queue-lengths that are reached during an over-saturation period are limited to $\overline{q}[i_0] = bi_\Delta \text{MBS}/i_0$, see (5.14). Thus, the transient probability $\gamma(t_0, B)$ is also expected to drop off for buffer sizes $B > \overline{q}[i_0]$.
- Limited by time-interval t_0: If the buffer becomes so large that a long over-saturation period cannot fill it up in time t_0 even if it starts right away, the overflow-probability $\gamma(t_0, B)$ decays rapidly. This happens approximately for buffer-sizes beyond

$$B_{t_0} := t_0\ (\Lambda_{i_0} - \nu) = \frac{b}{1-b}\, i_\Delta\, t_0 \kappa\,.$$

In the left graph of Fig. 6.10, the drop-off due to the truncation of the burst-length distribution would occur for buffers beyond $\overline{q}[i_0] = 4.7 \cdot 10^6$ for the case $\rho = 0.2$, and $\overline{q}[i_0] = 7.4 \cdot 10^6$ for $\rho = 0.5$. A quick decay of $\gamma(t_0, B)$ due to the finite observation interval would be observed already much earlier, since $B_{t_0} = 10^4$ when $\rho = 0.2$ respectively $B_{t_0} = 1.6 \cdot 10^4$ when $\rho = 0.5$.

6.4 Conditional Buffer Overflow Ratio

The transient overflow probabilities $\gamma(t_0, B)$ separates the 'good' observation intervals with no overflows or losses form the 'bad' intervals, during which some overflows occur. Although for low values of γ, the number of 'bad' overflow intervals is very low, Fig. 6.5 points out that the behavior in those few bad intervals might be extremely bad. To investigate such behavior, we introduce another transient performance parameter, the *conditional Buffer-Overflow Ratio*:

$$\text{BOR}_c(t_0, B) := \mathbb{E}\left\{\text{BOR}(t_0, B) \,|\, \text{BOR}(t_0, B) > 0\right\}.$$

Note that $\text{BOR}(t_0, B)$ is a random variable (see Sect. 6.2.1), while $\text{BOR}_c(t_0, B)$ is defined here as an expected value.

The transient overflow-probability $\gamma(t_0, B)$ is independent of the buffer-model: it does not matter whether cells are discarded or stored in the infinite secondary buffer since only the first overflow or loss event matters. On the other hand, the fraction of lost cells is smaller in the finite-buffer loss model than the fraction of overflowed cells in the infinite-buffer model. Hence, an analogous definition for the *conditional Cell-Loss Ratio* in the finite N-Burst/M/1/B loss model is necessary:

$$\text{CLR}_c(t_0, B) := \mathbb{E}\left\{\text{CLR}(t_0, B) \,|\, \text{CLR}(t_0, B) > 0\right\}.$$

The two transient performance parameters express the expected fraction of overflowed or lost cells, conditioned on at least one observed overflow/loss in a time-interval with length t_0. Standard stochastic definitions provide a relation between BOR_c and γ:

$$\text{BOR}_c(t_0, B) = \mathbb{E}\left\{\text{BOR}(t_0, B) \mid \text{BOR}(t_0, B) > 0\right\} =$$

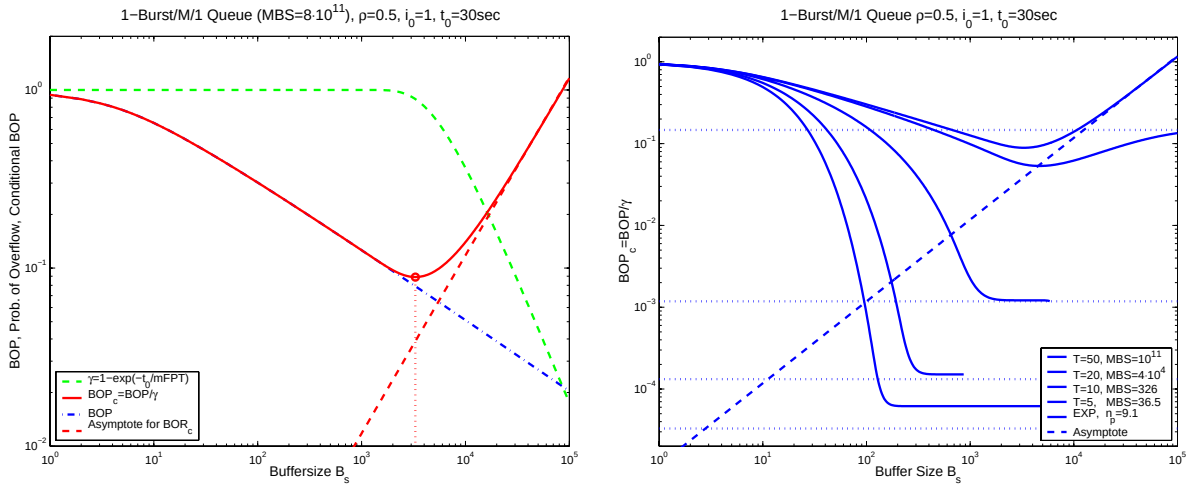

Figure 6.11: **Computation of the conditional Buffer Overflow Probability in 1-Burst/M/1 Models:** At first surprisingly, $\mathrm{BOP}_c(t_0, B)$ grows with increasing buffer-sizes, see left graph. The asymptote for its growth (dashed straight line) is determined later in the text, see Eq. (6.8). For truncated tails, the right graph shows that the BOP_c converges to a constant instead, approximated by the dotted horizontal lines computed from the lower bound of Eq. (6.11). $\gamma(t_0, B)$ and thus $\mathrm{BOP}_c(t_0, B)$ are computed for a time-interval of $t_0 = 30\,\mathrm{sec}$, which corresponds to about $5 \cdot 10^5$ arrivals at the queue.

$$= \int_{x=0^+}^{1} x \cdot \frac{\mathbb{P}\left\{\mathrm{BOR}(t_0, B) = x\right\}}{\gamma(t_0, B)}\, dx = \frac{\mathbb{E}\left\{\mathrm{BOR}(t_0, B)\right\}}{\gamma(t_0, B)}.$$

Thus, $\mathrm{BOR}_c(t_0, B)$ can be substantially larger than $\mathbb{E}\left\{\mathrm{BOR}(t_0, B)\right\}$ when an overflow event occurs only with small probability $\gamma(t_0, B)$ within $[0; t_0]$. Therefore, the conditional overflow ratio is mainly interesting for reliable networks with low γ. However, a low γ can also be the case when looking at short time-scales of e.g. short connection lengths. The analog arguments hold for the CLR_c.

6.4.1 Asymptotic Behavior of BOR_c for Large Buffers

Since it is hard to compute $\mathbb{E}\left\{\mathrm{BOR}(t_0, B)\right\}$ in the analytic model, we compute

$$\mathrm{BOP}_c(t_0, B) := \frac{\mathrm{BOP}(B)}{\gamma(t_0, B)}$$

instead of BOR_c. However, it turns out that $\mathbb{E}\left\{\mathrm{BOR}(t_0, B)\right\} \longrightarrow \mathrm{BOP}(B)$ rather quickly (when $t_0 > B/(\Lambda_{i_0} - \nu)$), so that $\mathrm{BOP}_c(t_0, B) \approx \mathrm{BOR}_c(t_0, B)$, if B is not too large with respect to t_0, see Sect. 6.4.4. Figure 6.11 shows in its left graph the computation of an analytic 1-Burst/M/1 model in $i_0 = 1$: shown are the steady-state $\mathrm{BOP}(B)$, the transient probability $\gamma(t_0, B)$ (by Eq. (6.5)), and finally $\mathrm{BOP}_c(t_0, B)$ as the ratio of the two others. We observe that first $\mathrm{BOP}_c \approx \mathrm{BOP}$, before at about $B = 3000$, $\gamma(t_0, B)$ starts to drop off, which causes BOP_c to *grow*. For large B, the BOP_c curve looks like a straight line with slope 1 on log-log scale, i.e. it grows linearly. The right graph of Fig. 6.11 makes clear that such a linear growth of the BOP_c is peculiar to the N-Burst/M/1 queue with Power-tailed ON periods, since early truncations of the ON-time distribution lead to a converging BOP_c.

It is shown in Appendix G.2 that the conditional overflow ratio, $\mathrm{BOR}_c(B)$ in an N-Burst/M/1 model grows asymptotically linearly with increasing buffer-size B,

$$\frac{1}{\alpha-1}\cdot\frac{1}{i_\Delta}\cdot\frac{1}{1-\rho}\cdot\frac{B}{\lambda\, t_0} \;\le\; \mathrm{BOR}_c(t_0,B) \;\le\; \frac{1}{i_0\,(\alpha-1)}\cdot\frac{N^{\frac{1-b}{b}\frac{1}{\rho}}}{i_\Delta\,(1-i_\Delta)}\cdot\frac{B}{\lambda\, t_0}\,, \tag{6.8}$$

when the model operates in blow-up region i_0 and the burst-length distribution shows sufficiently long Power-Tails. Formula (6.8) is derived by estimating the number of overflow-events that are caused by a single, long over-saturation period of duration X', where X' is larger than the time, $x_0 = B/(\Lambda_{i_0} - \nu)$, which is necessary to fill up the buffer, see Appendix G.2. After the over-saturation period ends, the queue needs another time period of duration $T_{(dr)}$ to drain until the queue-length is back at level B, after which no further overflow events are caused.

The two asymptotic bounds in (6.8) are identical for $i_0 = 1$:

$$\mathrm{BOR}_c(t_0,B) \;\sim\; \frac{1}{\alpha-1}\cdot\frac{1}{i_\Delta}\cdot\frac{1}{1-\rho}\cdot\frac{B}{\lambda\, t_0} \quad \text{for} \quad i_0 = 1 \tag{6.9}$$

The fact that there are best case and worst case assumptions for the rate of draining in the infinite buffer model leads to the two bounds for the asymptotic behavior, see Appendix G.2. The derivation of the asymptotic behavior of the CLR_c in the finite-buffer loss model is actually simpler, since there is no drain period involved. Appendix G.1 shows that we obtain for the asymptotic behavior

$$\mathrm{CLR}_c(B,t_0) \;\sim\; \frac{1}{i_0\,(\alpha-1)}\cdot\frac{B}{\lambda\, t_0}\;. \tag{6.10}$$

Eq. (6.10) is not quite mathematically rigorous. A simultaneous limit $B \to \infty$ *and* $t_0 \to \infty$ would have to be considered, see the remark in Appendix G.1 for further explanation.

Truncated Tails

For exponentially truncated Power-tails in the burst-length distribution, the asymptotic behavior of $\mathrm{BOR}_c(t_0,B)$ and $\mathrm{CLR}_c(t_0,B)$ is derived in Appendix G.3:

$$\frac{b}{1-\rho}\cdot\frac{\mathrm{MBS}}{\lambda\, t_0} \;\le\; \mathrm{BOR}_c^{(TPT)}(t_0,B) \;\le\; \frac{N}{i_0}\cdot\frac{1-b}{\rho\,(1-i_\Delta)}\cdot\frac{\mathrm{MBS}}{\lambda\, t_0}\;. \tag{6.11}$$

The two bounds are again identical in $i_0 = 1$. The larger the MBS is, the tighter the bounds turn out to be, see the horizontal dotted lines in the right graph of Fig. 6.11.

As before, since there is no drain-period involved in the finite-buffer model, a single expression results for the CLR_c:

$$\mathrm{CLR}_c^{(TPT)}(t_0,B) \;\sim\; \frac{b\, i_\Delta}{i_0}\cdot\frac{\mathrm{MBS}}{\lambda\, t_0}\,, \tag{6.12}$$

Note that in contrast to the scenario with infinite PTs, the formulas in Eqs. (6.11) and (6.12) do not depend on the buffer-size.

Conclusion: Additional buffer space decreases the number of 'bad' days, while it *increases* the conditional Overflow Ratio for LRD ON/OFF traffic from the N-Burst model. For traditional, exponential burst-lengths the conditional Overflow Ratio converges. However, note that larger buffers do of course prevent some overflows, but they on the other hand even increase the observed fluctuations in daily overflow-ratios: Either no overflows at all, or a huge number of overflows.

Example

Let us consider a 2-Burst model where each of the sources is described by the `TX3` calibration, as described in the beginning of this chapter, at utilization $\rho = 15\% \Longrightarrow i_0 = 2,\ i_\Delta = 0.5088$. Then, the asymptotic behavior of the BOR_c for full PTs is bounded by

$$5.78 \cdot \frac{B}{\lambda\, t_0} \leq \mathrm{BOR}_c(t_0, B) \leq 8.77 \cdot \frac{B}{\lambda\, t_0},$$

and in case of truncated tails

$$1.04 \cdot \frac{\mathrm{MBS}}{\lambda\, t_0} \leq \mathrm{BOR}_c^{(TPT)}(t_0, B) \leq 1.57 \cdot \frac{\mathrm{MBS}}{\lambda\, t_0}.$$

The following table computes the asymptotic values together with the BOP, mFPT, and BOP_c for $t_0 = 5\,\mathrm{min}$:

B	1000	5000	10^4	$5 \cdot 10^4$
upper bound: $\mathrm{BOR}_c^{(PT)}$	$9.0 \cdot 10^{-4}$	$4.5 \cdot 10^{-3}$	$9.0 \cdot 10^{-3}$	4.5%
lower bound: $\mathrm{BOR}_c^{(PT)}$	$5.9 \cdot 10^{-4}$	$3.0 \cdot 10^{-3}$	$5.9 \cdot 10^{-3}$	3.0%
$\mathrm{BOP}(B)$	$5.6 \cdot 10^{-4}$	$1.4 \cdot 10^{-4}$	$7.6 \cdot 10^{-5}$	$1.6 \cdot 10^{-5}$
$\mathrm{mFPT}(B)$	7.7 min	2.4 h	8.3 h	7.0 d
$\gamma(t_0, B)$	48%	3.5%	1.0%	$5.0 \cdot 10^{-4}$
BOP/γ	$1.2 \cdot 10^{-3}$	$4.1 \cdot 10^{-3}$	$7.6 \cdot 10^{-3}$	3.3%

Thereby, the values for the BOP, mFPT are computed exactly for $T = 30$ which corresponds to $\mathrm{MBS} = 5.6 \cdot 10^6$. $\gamma(t_0, B)$ is computed by the exponential approximation (6.5).

6.4.2 Asymptotic Behavior of the Mean First Passage Time

The knowledge of the asymptotic behavior of the BOR_c can be used to derive several corollaries: one of them is a confirmation of the observed asymptotic behavior of the mFPT in Sect. 6.2.2.

Since the behavior of the steady-state BOP is known from Sect. 5.2 to be $\mathrm{BOP}(B) \sim c_{\mathrm{BOP}} \cdot B^{1-\beta}$, and we now found the asymptotic behavior of $\mathrm{BOR}_c(B)$ to be linearly increasing with B, we can derive the asymptotic behavior of the mFPT. By using the exponential approximation (6.5) for $\gamma(t_0, B)$ and the approximation $1 - \exp(-x) \approx x$ for small x, we obtain:

$$\mathrm{BOR}_c(t_0, B) \approx \frac{\mathrm{BOP}(B)}{\gamma(t_0, B)} \approx \frac{\mathrm{BOP}(B)}{1 - \exp(-t_0/\mathrm{mFPT}(B))} \approx \frac{\mathrm{BOP}(B) \cdot \mathrm{mFPT}(B)}{t_0}.$$

Since $\mathrm{BOR}_c^{(PT)}(B) \sim B$, we get for the mean First Passage Time:

$$\mathrm{mFPT}(B) \approx \frac{t_0 \cdot \mathrm{BOR}_c^{(PT)}(t_0, B)}{\mathrm{BOP}(B)} \sim \frac{B}{B^{1-\beta}} = B^\beta.$$

Therefore, for ON/OFF traffic with Power-tailed ON periods with exponent α, the mFPT grows asymptotically by a Power-Law with exponent $\beta = i_0\,(\alpha - 1) + 1$. Computations of mFPTs for N-Burst models in Sect. 6.2.2 confirm such an asymptotic behavior.

The constant in the asymptotic behavior of $\mathrm{mFPT}(B)$ can be derived as follows:

$$\mathrm{BOR}_c^{(PT)}(t_0, B) \approx \frac{c_{\mathrm{BOP}}\, B^{1-\beta} \cdot c_{\mathrm{mFPT}}\, B^\beta}{t_0} \quad \Longrightarrow$$

$$\frac{1}{c_{\mathrm{BOP}}} \cdot \frac{1}{\lambda} \cdot \frac{1}{(1-\rho)\,(\alpha-1)\,i_\Delta} \;\leq\; c_{\mathrm{mFPT}} \;\leq\; \frac{1}{c_{\mathrm{BOP}}} \cdot \frac{1}{\lambda} \cdot \frac{N\,(1-b)}{i_0\,\rho\,(\alpha-1)\,b\,i_\Delta\,(1-i_\Delta)}\,. \tag{6.13}$$

Again, the two bounds are identical when $i_0 = 1$:

$$c_{\mathrm{mFPT}} = \frac{1}{c_{\mathrm{BOP}}} \cdot \frac{1}{\lambda} \cdot \frac{1}{(1-\rho)\,(\alpha-1)\,i_\Delta}, \quad \text{for} \quad i_0 = 1. \tag{6.14}$$

Thereby, a relationship between the constants in the asymptotic behavior of the steady-state parameter BOP and the transient mFPT is established. Since it is generally harder to compute $\mathrm{mFPT}(B)$ for large B due to the iterative algorithm (see Appendix F.3), the established relationship is very useful.

Furthermore, a similar relationship between the constant in the tail-behavior of the finite-buffer CLP (see Sect. 5.5) and the mFPT can be derived from Eq. (6.10):

$$c_{\mathrm{CLP}} = \frac{1}{c_{\mathrm{mFPT}}\,\lambda\,i_0\,(\alpha-1)}\,. \tag{6.15}$$

Consequently, the combinations of the formulas for c_{mFPT} with Eq. (6.15) yield an upper and lower bound for c_{CLP}:

$$\frac{\rho\,b\,i_\Delta\,(1-i_\Delta)}{N\,(1-b)}\,c_{\mathrm{BOP}} \;\leq c_{\mathrm{CLP}} \;\leq\; (1-\rho)\,\frac{i_\Delta}{i_0}\,c_{\mathrm{BOP}}\,. \tag{6.16}$$

As usual the two bounds are identical when $i_0 = 1$:

$$c_{\mathrm{CLP}} = (1-\rho)\,i_\Delta\,c_{\mathrm{BOP}} \quad \text{for} \quad i_0 = 1. \tag{6.17}$$

Equation (6.17) was discovered empirically in [Schwefel & Lipsky 99b]. The relationship between c_{BOP} and c_{CLP} is already mentioned in Section 5.5, but the derivation required the results of the transient analysis in this chapter.

Finally, the approximation Eq. (5.10) for the tail-constant c_{BOP} that was discussed in the previous chapter can be used to obtain an explicit formula for c_{mFPT}:

$$c_{\mathrm{mFPT}} \approx \frac{\overline{n_p}}{(b\,i_\Delta\,\overline{n_p})^\alpha} \cdot \frac{1}{c^{(1)}_{\mathrm{PT}}(\alpha)} \cdot \frac{1}{\lambda}\,, \quad \text{when} \quad i_0 = 1. \tag{6.18}$$

6.4.3 Minimum of $\mathrm{BOR}_c(B)$

We can observe in Fig 6.11 that $\mathrm{BOR}_c(B)$ first follows $\mathrm{BOP}(B)$, since $\gamma(t_0, B) \approx 1$ for small buffers B. Therefore, $\mathrm{BOR}_c(B)$ decreases first, before it grows asymptotically linearly as in (6.8). Consequently, there is some buffer-size B_{min} for which a minimum of the BOR_c is achieved. To derive the location of the minimum, let us start with

$$\mathrm{BOR}_c(t_0, B) \approx \frac{\mathrm{BOP}(B)}{1 - \exp\left[-t_0/\mathrm{mFPT}(B)\right]}\,.$$

At the minimum, the first derivative is zero:

$$\frac{d\mathrm{BOR}_c(t_0, B)}{dB} = 0 \quad \Longleftrightarrow$$

$$\frac{d\mathrm{BOP}(B)}{dB} \cdot \left[1 - \exp\left(-\frac{t_0}{\mathrm{mFPT}(B)}\right)\right] +$$

$$+\,\mathrm{BOP}(B)\cdot\exp\left(-\frac{t_0}{\mathrm{mFPT}(B)}\right)\cdot\frac{t_0}{\mathrm{mFPT}(B)^2}\cdot\frac{d\mathrm{mFPT}(B)}{dB}=0\,.$$

Using the asymptotic behavior (5.19) of $\mathrm{BOP}(B)$ and (6.3) of $\mathrm{mFPT}(B)$, we get the equation for B_{min}:

$$\exp\left(-\frac{t_0}{c_{\mathrm{mFPT}}}\,B_{min}^{-\beta}\right)\cdot\left[1+\frac{1}{c_{\mathrm{mFPT}}}\,\frac{\beta}{\beta-1}\,t_0 B_{min}^{-\beta}\right]=1.$$

Substituting

$$z_0:=\frac{t_0}{c_{\mathrm{mFPT}}}\,B_{min}^{-\beta}\,,$$

the equation

$$\exp(-z_0)=\frac{1}{1+\frac{\beta}{\beta-1}\,z_0} \tag{6.19}$$

can be solved numerically for its root $z_0(\beta)$. A table for numerically computed values of $z_0(\beta)$ for different β is provided in Appendix I.3 .

Having obtained $z_0(\beta)$ numerically,

$$B_{min}=\sqrt[\beta]{\frac{t_0}{c_{\mathrm{mFPT}}\,z_0(\beta)}}\,, \tag{6.20}$$

is the asymptotic location of the minimal BOR_c – asymptotic in the sense that it is a valid approximation if the value of B_{min} is not too small. Note that c_{mFPT} depends on N, λ, ρ, $\overline{n_p}$, b, and α, thus it depends also on $\beta=i_0(\alpha-1)+1$. However, c_{mFPT} is independent of t_0. Consequently, B_{min} grows proportional to $\sqrt[\beta]{t_0}$.

In the first blow-up region, $i_0=1$, Eq. (6.18) can be used:

$$B_{min}(i_0=1)\approx b\,i_\Delta\,\overline{n_p}\,\sqrt[\alpha]{\frac{c_{\mathrm{PT}}^{(1)}(\alpha)}{\overline{n_p}\,z_0(\alpha)}\,\lambda\,t_0}\,.$$

Example: For the 2-Burst `TX3` model with $\rho=0.15$ ($i_0=2$ and $\beta=1.8$), where $c_{\mathrm{BOP}}=0.15$ and $c_{\mathrm{mFPT}}=1.5\,\mathrm{ms}$, the $\mathrm{BOR}_c(t_0,B)$ shows minima at the following buffer-sizes:

t_0	5s	5min	1h	5h	8h
B_{min}	74	717	2851	6971	9051
$\gamma(t_0,B_{min})$	100%	70%	9.3%	1.9%	1.2%
$\mathrm{BOP}_c(t_0,B_{min})$	$5.4\cdot10^{-3}$	$1.1\cdot10^{-3}$	$2.5\cdot10^{-3}$	$5.5\cdot10^{-3}$	$7.0\cdot10^{-3}$

However, the value $B_{min}(5\,\mathrm{s})=74$ might not locate the actual minimum well. Its value is so small that the underlying assumptions of the asymptotic Power-Law behaviors in the derivation of (6.20) are questionable.

6.4.4 Validation via Simulation

Finally, we validate the derived asymptotes for the BOR_c in several simulation experiments.

Both graphs in Figure 6.12 show that the simulation estimates for $\mathbb{E}\{\mathrm{BOR}\}$, γ and BOR_c correspond well with the computation of the BOP, γ and BOP_c in the analytic N-Burst model.

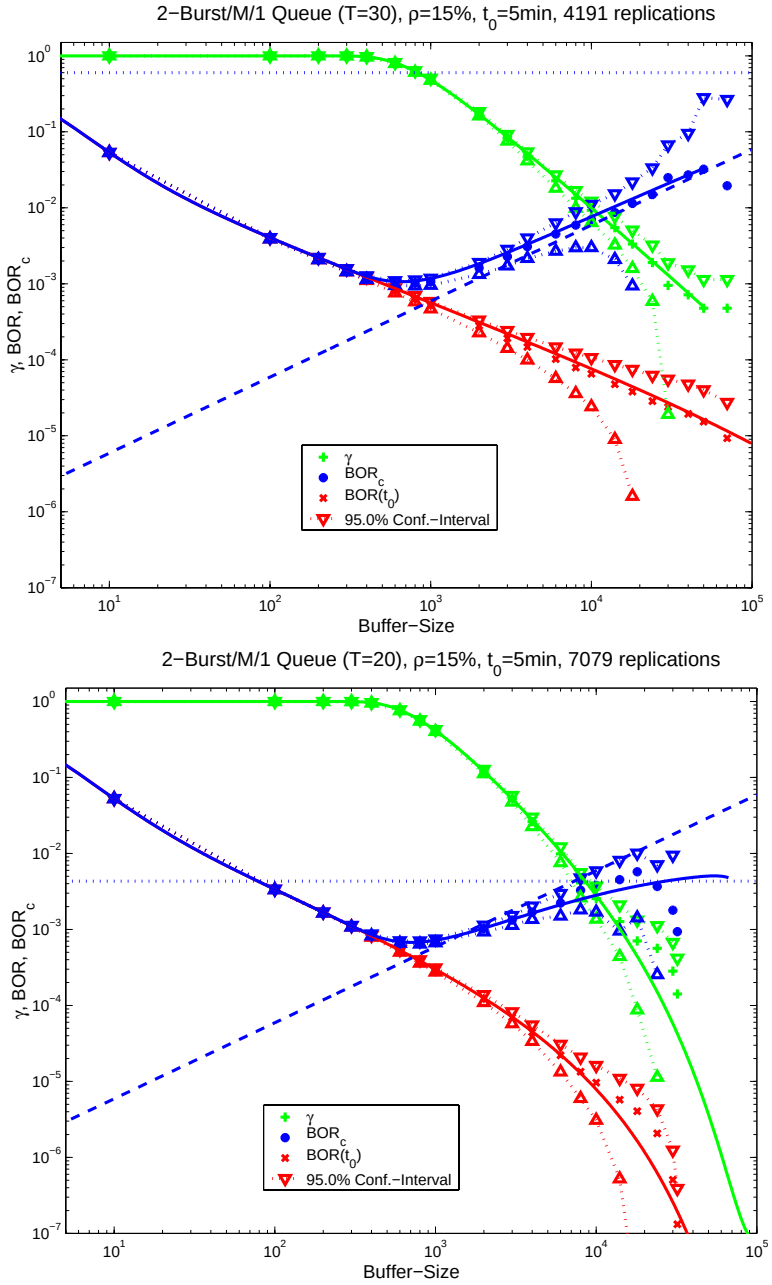

Figure 6.12: **Validation of the Behavior of the Conditional Overflow Ratio for 2-Burst/M/1 Models via Simulation:** Several thousand simulation runs of $t_0 = 5\,\text{min}$ simulated time with on average almost 10 million arriving cells per replication were simulated and the expected overflow ratio $\mathbb{E}\{\text{BOR}(t_0, B)\}$ (marked by '×'), the probability $\gamma(t_0, B)$ ('+'), and the conditional overflow ratio $\text{BOR}_c(t_0, B)$ (bullets) were estimated from the simulation results. The analytic computation (solid lines) of the BOP, the approximation for γ via the mFPT, and the resulting BOP_c turned out to be well within the 95% confidence intervals of the estimators (marked by triangles). The asymptotic approximations Eqs. (6.8) and (6.11) are also plotted by the dashed and dotted lines.

Note that it is the steady-state BOP that is compared to the simulation estimate for $\mathbb{E}\{\mathrm{BOR}\}$. However, since the observation interval t_0 is sufficiently long, those two values are close. For a counter-example, see below.

The difference between the upper and lower graph in Fig. 6.12 is that the impact of the truncation can be observed in the lower graph for the plotted range of buffer-sizes. For $T = 20$, the truncation has a visible impact, while $T = 30$ in the upper graph does not: In the latter case, $\mathbb{E}\{\mathrm{BOR}\}$ and γ decay as Power-Laws and BOR_c grows approximately linearly with the asymptote from (6.8). The following values of $\overline{q}[i_0]$ indicate that only for buffer-size bigger than 10^6 does the truncation $T = 30$ in the upper graph show its impact:

$$\begin{array}{lllll} \text{Upper Graph:} & T = 30, & \mathrm{MBS} = 5.6 \cdot 10^6, & \overline{q}[i_0] = 1.3 \cdot 10^6 \\ \text{Lower Graph:} & T = 20, & \mathrm{MBS} = 4.06 \cdot 10^4, & \overline{q}[i_0] = 9.2 \cdot 10^3 \end{array}$$

The Power-Laws, both for $\mathbb{E}\{\mathrm{BOR}\}$ and γ turn into a quick exponential drop-off for $B \approx 10^4$ in the lower graph, and the BOR_c converges approximately to the horizontal asymptote which is predicted by (6.11). Interestingly, the BOR_c curve for $T = 20$ is influenced by both asymptotes: at first, the BOR_c grows when $B > 800$, but later the impact of the truncation causes convergence. Thus, the knowledge of both asymptotes provides a good picture of how $\mathrm{BOR}_c(t_0, B)$ evolves with the buffer-size. Since the computation of the asymptotes is straightforward without any complicated matrix manipulations, they can well be used in practical design problems.

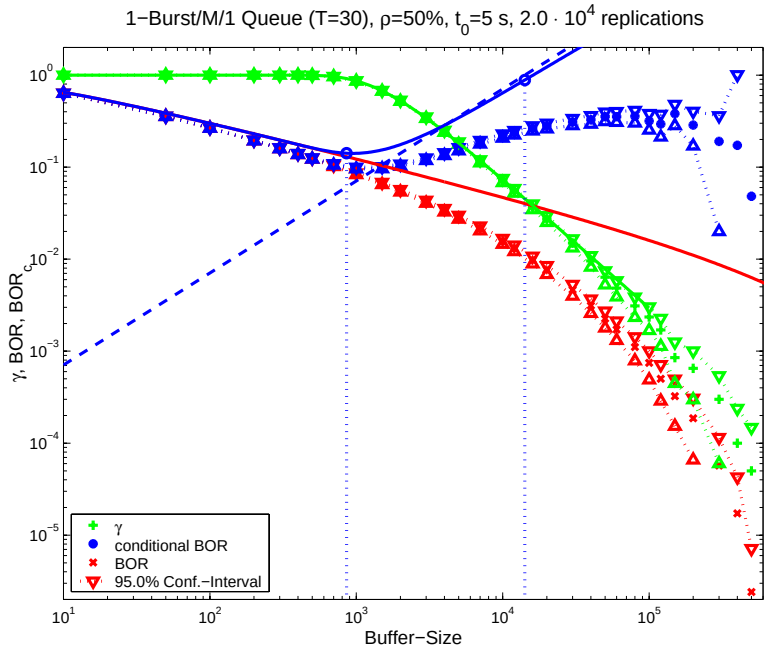

Figure 6.13: **Simulation Results for the BOR_c for large buffers and short t_0:** Although the exponential approximation for $\gamma(t_0, B)$ works reasonably well for the whole range of B, the time t_0 is so short that $\mathbb{E}\{\mathrm{BOR}(t_0, B)\}$ is noticeable smaller then the steady-state $\mathrm{BOP}(B)$.

In Figure 6.12, the observation interval t_0 was long enough such that $\mathbb{E}\{\mathrm{BOR}(t_0, B)\} \approx \mathrm{BOP}(B)$. Fig. 6.13 shows the result of another simulation experiment, in which t_0 is so small, that it causes $\mathbb{E}\{\mathrm{BOR}(t_0, B)\} \ll \mathrm{BOP}(B)$ for the large buffer-sizes in the plotted range. Note that the computed value of γ still matches the simulation results, but the BOR_c does not grow linearly. Instead it converges to some value around 30%. The predicted asymptote reaches 100% at $B = 1.4 \cdot 10^4$ (marked by second dotted vertical line), therefore it is clear that it is invalid for buffer-sizes beyond that point, due to the observation interval t_0 which is too short.

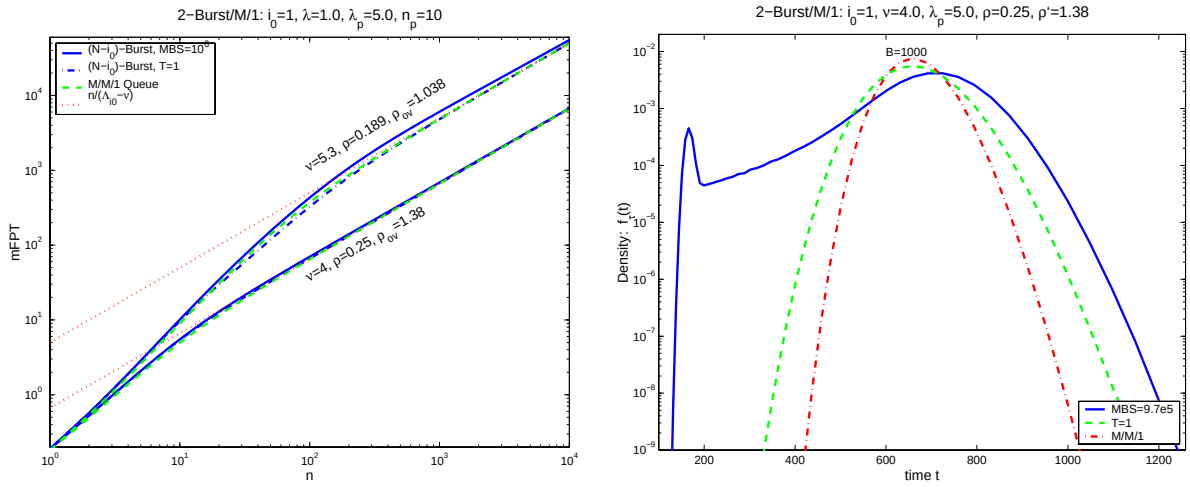

Figure 6.14: **Transient Behavior During Over-saturation Periods:** A 2-Burst/M/1 queue in blow-up region $i_0 = 1$ behaves like an unstable 1-Burst/M/1 queue with increased background rate during an over-saturation period.

6.5 Transient Analysis of Over-Saturation Periods

In various locations during the steady-state and the transient analysis, we have pointed out the relevance of very long over-saturation periods that occur in the N-Burst model with PT distributed burst-lengths. The analysis that is presented in Sect. 5.3 relies on the argument that during an over-saturation period with i_0 long-term active bursts, the queue-length grows with average rate $\Lambda_{i_0} - \nu$. With the help of the transient analysis that has been developed in this chapter, we can now look more closely at the behavior during the over-saturation periods.

First, let us consider the somewhat simpler case $i_0 = N$. The behavior of an N-Burst model during an over-saturation period with N permanently active sources is equivalent to an unstable M/M/1 queue with arrival rate $\Lambda_N = \lambda_0 + N\lambda_p > \nu$. The First Passage Time to some buffer-level n is approximately Gamma distributed with parameters

$$\mathrm{mFPT}(n) \sim \frac{n}{\Lambda_N - \nu}, \quad C^2\left[\tau_n\right] \sim \frac{\Lambda_N + \nu}{\Lambda_N - \nu}\frac{1}{n}$$

for large n, see [Asmussen et al. 00a]. Graphs of the FPT distribution for the M/M/1 queue are shown in Appendix A.2. The assumption of an average growth with rate $\Lambda_N - \nu$ is therefore reasonable.

When we look at blow-up regions, $1 \leq i_0 < N$, the situation becomes more complicated. During an over-saturation period with i_0 permanently active sources, the N-Burst/M/1 model behaves as an unstable N'-Burst/M/1 model with $N' = N - i_0$ and $\lambda_0' = \lambda_0 + i_0\lambda_p$ while $\lambda_p' = \lambda_p$, $b' = b$ and $\overline{n_p}' = \overline{n_p}$ remain unchanged.

The behavior of $\mathrm{mFPT}(n)$ for such an unstable N'-Burst/M/1 queue is shown in the left graph of Figure 6.14. Three models are thereby compared: the M/M/1 queue, the N'-Burst/M/1 queue with exponential burst-length, and finally the N'-Burst/M/1 queue with TPT distributed burst-lengths with large PT Range. In the lower set of curves ($\nu = 4$), all three models result in practically identical curves for $\mathrm{mFPT}(n)$. However, if we look at an N-Burst/M/1 model, which

operates close to the blow-up point, the utilization of the resulting N'-Burst/M/1 queue that reflects the behavior during the over-saturation periods is close to 1 (but necessarily still above 1). In that case, the LRD version of the N'-Burst/M/1 queue shows slightly larger values for mFPT(n) in the left graph of Figure 6.14.

When we look at the whole distribution of FPT(n) for a certain n in the right graph of Fig. 6.14, the difference in the models becomes obvious: although the less critical case with $\rho' = 1.38$ well above 1 is shown there, the distributions first show different variance, and secondly, the peculiar peak just below $t = 200$ appears in the LRD model. For large enough n, the asymptotic theory in [Asmussen et al. 00a] states that the FPT distribution converges towards a normal distribution for all three models. However, the convergence for the LRD model appears to be much slower.

As a conclusion, although it is true that the queue-length on average grows with rate $\Lambda_{i_0} - \nu$, large fluctuations in that average growth-rate are possible for the LRD model even for very long over-saturation periods.

6.6 Summary

The steady-state analysis in Chapter 5 provides quite a few insights into the behavior of N-Burst/M/1 queues with LRD properties. In traditional models of telecommunication systems, as for example for the N-Burst/M/1 queue with exponential burst-lengths, steady-state behavior is observed rather quickly. However, in a finite observation period, the behavior of the performance model with LRD traffic can be quite different from the steady-state results. Therefore, transient performance parameters can provide a much better description of the system's behavior.

Large cell delays as well as overflow or loss events are not independent events: Since they are both due to full buffers, subsequent cells experience large delays or cause buffer overflows. This is true in any queueing model, but it is particularly accentuated in models with LRD arrival processes. In order to investigate whether such scenarios with large buffer-occupancy occur in an observation interval, the analysis of the First Passage Time (FPT) is very useful.

It turns out that although for given buffer-size B the actual distribution of FPT(B) is rather well-behaved (close to exponential, see Sect. 6.2.3), the Power-Tailed burst-length distributions have a strong impact on the behavior of its expected value, mFPT(B): Within the blow-up region i_0 the asymptotic Power-Law relationship (6.3)

$$\text{mFPT}(B) \sim c_{\text{mFPT}} B^{\beta} \quad \text{with} \quad \beta = i_0(\alpha - 1) + 1,$$

is observed, whose exponent β is already familiar from the steady-state analysis in Sect. 5.2. One of the straightforward consequences is that the blow-up effects also occur for transient parameters (in particular mFPT and $\gamma(t_0, B)$).

With the help of the FPT distribution, the probability $1 - \gamma(t_0, B)$ that no overflow/loss event occurs during an observation interval of duration t_0 can be obtained. This transient overflow probability is highly relevant in practice, e.g. for connection admission purposes, see 7.4. Using the approximately exponential distribution of the FPT, an approximation

$$\gamma(t_0, B) \approx \frac{t_0}{\text{mFPT}(B)}$$

is obtained for large buffers, see Sect. 6.3. Together with the asymptotic Power-Law for mFPT(B), another Power-Law relationship follows for the transient overflow probability γ.

The transient overflow probability γ already describes part of the fluctuations that are observed in the fraction of lost (or overflowed) cells in finite observation intervals: Approximately a fraction $1 - \gamma$ of the observation intervals show no overflow events at all. However, in the intervals with at least one overflow event, a large number of cells might cause overflow events. The expected fraction of overflowed (lost) cells in the intervals with at least one such event is described by the conditional Buffer Overflow Ratio, BOR_c, (conditional Cell Loss Ratio, CLR_c), see Sect. 6.4. By an analysis of the number of overflowed cells during and immediately after a long oversaturation period, the asymptotic behavior of those conditional transient performance measures is derived, see Appendix G. At first surprisingly, those conditional performance measures increase with larger buffers, although the absolute number of overflows is of course reduced. This result describes nicely that the observed fluctuations – either no overflow, but if any then a large number – even increase with larger buffers for N-Burst traffic with LRD properties. The parameter pair $\langle \gamma(t_0, B),\ \mathrm{BOR}_c(t_0, B) \rangle$ is therefore recommended as transient replacement of the steady-state BOP.

The knowledge of the asymptotic behavior of the conditional performance measures allows to derive a relationship between the tail-constants in the Power-Law behavior of the mFPT and in the steady-state BOP respectively the CLP, see 6.4.2. Also, it is possible to derive a formula to compute an optimal buffer-size $B_{min}(t_0)$, at which the BOR_c shows a minimum. However, the practical relevance of this optimal buffer-size is not fully understood at this time.

Chapter 7

Applications

The results of the steady-state and transient analysis of N-Burst/M/1 queues in the previous chapters are highly relevant for the network planning task. In particular the asymptotic results are, due to their simplicity, well suited for estimates of required resources or limitations on the incoming traffic for intermediate systems with large buffers.

The integration of such solutions for local design problems (i.e. only with respect to a single bottleneck node) in the global network planning process is beyond the scope of this thesis. See Chapter 10 for a discussion of such a step. However, note that the commonly used M/M/1 approach in such global planning tools leads to performance results that are by far too optimistic. In particular in large buffer-systems, the M/M/1 analysis yields overflow or loss probabilities that are practically zero (in the order of ρ^B) unless the utilization, ρ, is extremely close to 1. Since we are mainly concerned with large-buffer ($B > 100$) systems at average utilization ρ substantially smaller than 1, a comparison with the M/M/1 queue is only included for the mean delay in Sect. 7.6.

Since all the examples in the previous two sections use realistic values for the parameters of the N-Burst model, the individual illustrations can be seen as applications of the N-Burst model to practical questions of network design, e.g. capacity planning in Fig. 5.1. Most of those illustrations use the exact analytic N-Burst model to obtain the performance results. Of course, the exact analytic model could be used in general for practical network design tasks. However, such computations are rather complex and time-consuming. Instead, the derived asymptotic relationships, together with the approximations for the tail-constants, provide a more feasible way of providing numerical values for several QoS parameters quickly. The inaccuracy of those approaches is in most cases negligible in comparison to the inaccuracy in measuring and forecasting the parameters of the traffic that is expected in the planned network.

7.1 Summary of Formulas

At network components with large buffers that experience ON/OFF traffic with LRD properties, the QoS behavior in terms of delay and cell losses or overflows is mainly determined by the tail of the queue-length distribution. Consequently, asymptotic results can well be used for practical design problems. This section summarizes all the important formulas for both, steady-state and transient performance parameters. Although all those formulas are derived in Chapters 5 and 6, new equation numbers are assigned here in order to make this chapter self-contained.

N-Burst Parameters	$N,\ \lambda,\ \overline{n_p},\ b,\ \alpha,\ c_{\mathrm{PT}}^{(1)}(\alpha),\ \mathrm{MBS}$	
	$\kappa = \dfrac{\lambda}{N},\ \lambda_p = \dfrac{\kappa}{1-b}$	
Switch Parameters	$B,\ \rho,\ \nu = \dfrac{\lambda}{\rho}$	
Blow-up Region	$i_0 = \left\lceil N \cdot \dfrac{1-\rho}{\rho} \cdot \dfrac{1-b}{b} \right\rceil \quad \text{for} \quad \rho > 1-b$	(7.1)
	$i_\Delta = i_0 - N \cdot \dfrac{1-\rho}{\rho} \cdot \dfrac{1-b}{b}$	(7.2)
	$\beta = i_0\,(\alpha - 1) + 1 \qquad [\text{effective PT exponent}]$	(7.3)
Steady-State Behavior:		
Buffer Overflow Prob.:	$\mathrm{BOP}(B) \sim c_{\mathrm{BOP}}\, B^{1-\beta}$	(7.4)
Cell Loss Prob.:	$\mathrm{CLP}(B) \sim c_{\mathrm{CLP}}\, B^{1-\beta}$	(7.5)
mean Delay (full tail):	$\mathrm{mCD} = \infty \ \Leftrightarrow\ \beta \le 2 \ \Leftrightarrow\ \alpha \le 1 + \dfrac{1}{i_0}$	(7.6)
Truncated Tails:		
	$q_{i_0} = b\,\dfrac{i_\Delta}{i_0}\,\mathrm{MBS} \qquad [\text{PT-Range for exponent } \beta]$	(7.7)
	$q_N = \mathrm{MBS}\left(1 - \dfrac{1-b}{\rho}\right) \quad [\text{PT-Range for all exponents}]$	(7.8)
	$\mathrm{mCD}(\mathrm{MBS}) \sim c_{\mathrm{mCD}}(\mathrm{MBS})^{2-\beta} \quad \text{for} \quad \beta < 2$	(7.9)
Transient Behavior:		
Time to first Loss:	$\mathrm{mFPT}(B) \sim c_{\mathrm{mFPT}}\, B^{\beta}$	(7.10)
Prob. of Loss:	$\gamma(t_0, B) \approx 1 - \exp\left(-\dfrac{t_0}{\mathrm{mFPT}(B)}\right)$	(7.11)
	$\sim 1 - \exp\left(-\dfrac{t_0}{c_{\mathrm{mFPT}}}\, B^{-\beta}\right)$	(7.12)

The background Poisson rate λ_0 does not appear in the list of the N-Burst parameters. All the formulas here assume that $\lambda_0 = 0$. If we want to compute the asymptotic behavior of N-Burst models with positive λ_0, we can use an N-Burst model without background Poisson rate instead, but reduce the service-rate ν, $\nu \to \nu - \lambda_0$, of the N-Burst/M/1 queue.

The Power-Law behavior of the performance parameters, $\mathrm{BOP}(B)$, $\mathrm{CLP}(B)$, $\mathrm{mFPT}(B)$ (and also for $\gamma(B)$, see (6.6)) of the N-Burst model is peculiar only for Power-Tailed ON periods and within the blow-up regions $i_0 = 1, ..., N$. The slow Power-Law decay (respectively increase for mFPT) has the effect that additional buffer-space is a very ineffective means for improving QoS. On the other hand, in traditional ON/OFF models with exponential (or other well-behaved) burst-length distributions, all overflow and loss probabilities drop off exponentially.

In addition to the knowledge of the Power-Law behavior and its PT exponent, approximations and bounds for the tail-constants were derived. Since the formulas for the worst blow-up region,

$i_0 = 1$, are less complicated, let us begin there:

$$i_0 = 1 \quad \Leftrightarrow \quad \beta = \alpha \quad \Leftrightarrow \quad \frac{N}{N + \frac{b}{1-b}} < \rho < 1 \ :$$

Tail Constants:

$$c_{\mathrm{BOP}} \approx \frac{1}{1-\rho} \, \frac{(i_\Delta \, b \, \overline{n_p})^\alpha}{i_\Delta \, \overline{n_p}} \, \frac{c_{\mathrm{PT}}^{(1)}(\alpha)}{\alpha - 1} \tag{7.13}$$

$$c_{\mathrm{CLP}} \approx \frac{(i_\Delta \, b \, \overline{n_p})^\alpha}{\overline{n_p}} \, \frac{c_{\mathrm{PT}}^{(1)}(\alpha)}{\alpha - 1} \tag{7.14}$$

$$c_{\mathrm{mFPT}} \approx \frac{1}{\lambda} \, \frac{\overline{n_p}}{(i_\Delta \, b \, \overline{n_p})^\alpha} \, \frac{1}{c_{\mathrm{PT}}^{(1)}(\alpha)} \tag{7.15}$$

$$c_{\mathrm{mCD}} < \frac{1}{\lambda} \, \frac{\rho}{1-\rho} \cdot \frac{i_\Delta \, b^2 \, \overline{n_p}^{\alpha-1}}{2 - \alpha} \, \frac{c_{\mathrm{PT}}^{(1)}(\alpha)}{\alpha - 1}, \quad \alpha < 2 \tag{7.16}$$

$$\gamma(t_0, B) \sim 1 - \exp\left(- \frac{(i_\Delta \, b \, \overline{n_p})^\alpha}{\overline{n_p}} \, c_{\mathrm{PT}}^{(1)}(\alpha) \, \frac{\lambda t_0}{B^\alpha} \right) \tag{7.17}$$

$$\mathrm{BOR}_c(t_0, B) \sim \frac{1}{\alpha - 1} \, \frac{1}{(1-\rho) \, i_\Delta} \cdot \frac{B}{\lambda \, t_0} \tag{7.18}$$

$$\mathrm{CLR}_c(t_0, B) \sim \frac{1}{\alpha - 1} \cdot \frac{B}{\lambda \, t_0} \tag{7.19}$$

$$B_{min} \approx i_\Delta \, b \, \overline{n_p} \, \sqrt[\alpha]{\frac{c_{\mathrm{PT}}^{(1)}(\alpha)}{\overline{n_p}} \cdot \frac{\lambda t_0}{x_0(\alpha)}} \tag{7.20}$$

Truncated Tails:

$$\mathrm{BOR}_c^{(TPT)}(t_0, B) \sim \frac{b}{1-\rho} \cdot \frac{\mathrm{MBS}}{\lambda \, t_0} \tag{7.21}$$

$$\mathrm{CLR}_c^{(TPT)}(t_0, B) \sim i_\Delta \, b \, \frac{\mathrm{MBS}}{\lambda \, t_0} \tag{7.22}$$

In the general case,

$$i_0 > 1 \ \Leftrightarrow \ 1 - b \ < \ \rho \ < \ \frac{N}{N + \frac{b}{1-b}} \, ,$$

there exist the following approximations and bounds:

$$c_{\mathrm{BOP}}: \quad c_{\mathrm{BOP}} \approx \frac{1}{1-\rho}\,\frac{(i_\Delta\, b\overline{n_p})^\beta}{i_\Delta\,\overline{n_p}}\,\frac{1-b}{\sqrt{1-i_\Delta}}\left[\frac{c_{\mathrm{PT}}^{(1)}(\alpha)}{\alpha-1}\right]^2, \quad \text{for } i_0 = 2 \tag{7.23}$$

$$c_{\mathrm{BOP}} \approx \frac{1}{1-\rho}\,\frac{(i_\Delta\, b\overline{n_p})^\beta}{i_\Delta\,\overline{n_p}}\,(1-b)^{i_0-1}\left[\frac{c_{\mathrm{PT}}^{(1)}(\alpha)}{\alpha-1}\right]^{i_0}, \quad \text{for } i_0 > 2 \tag{7.24}$$

$$c_{\mathrm{CLP}}: \quad \frac{\rho\, b\, i_\Delta\,(1-i_\Delta)}{N\,(1-b)}\,c_{\mathrm{BOP}} \le c_{\mathrm{CLP}}$$

$$c_{\mathrm{CLP}} \le (1-\rho)\,\frac{i_\Delta}{i_0}\,c_{\mathrm{BOP}} \tag{7.25}$$

$$c_{\mathrm{mFPT}}: \quad \frac{1}{c_{\mathrm{BOP}}}\cdot\frac{1}{\lambda}\cdot\frac{1}{(1-\rho)\,(\alpha-1)\,i_\Delta} \le c_{\mathrm{mFPT}} \tag{7.26}$$

$$c_{\mathrm{mFPT}} \le \frac{1}{c_{\mathrm{BOP}}}\cdot\frac{1}{\lambda}\cdot\frac{N\,(1-b)}{i_0\,\rho\,(\alpha-1)\,b\,i_\Delta\,(1-i_\Delta)}$$

$$c_{\mathrm{mCD}}: \quad c_{\mathrm{mCD}} < \frac{\rho\, c_{\mathrm{BOP}}}{\lambda\,(2-\beta)}\left[\frac{b\,i_\Delta}{i_0}\right]^{2-\beta} \tag{7.27}$$

$$\mathrm{BOR}_c: \quad \frac{1}{\alpha-1}\,\frac{1}{(1-\rho)\,i_\Delta}\,\frac{B}{\lambda\, t_0} < \mathrm{BOR}_c(t_0,B)$$

$$\mathrm{BOR}_c(t_0,B) < \frac{1}{\beta-1}\cdot\frac{1}{(1-\rho)\,i_\Delta}\cdot\frac{i_0-i_\Delta}{1-i_\Delta}\,\frac{B}{\lambda\, t_0} \tag{7.28}$$

$$\mathrm{CLR}_c: \quad \mathrm{CLR}_c(t_0,B) \sim \frac{1}{\beta-1}\cdot\frac{B}{\lambda\, t_0} \tag{7.29}$$

Truncated Tails:

$$\mathrm{BOR}_c^{(TPT)}: \quad \frac{b}{1-\rho}\,\frac{\mathrm{MBS}}{\lambda\, t_0} < \mathrm{BOR}_c^{(TPT)}(t_0,B)$$

$$\mathrm{BOR}_c^{(TPT)}(t_0,B) < \frac{b}{1-\rho}\cdot\frac{i_0-i_\Delta}{i_0\,(1-i_\Delta)}\,\frac{\mathrm{MBS}}{\lambda\, t_0} \tag{7.30}$$

$$\mathrm{CLR}_c^{(TPT)}: \quad \mathrm{CLR}_c^{(TPT)}(t_0,B) \sim \frac{b\,i_\Delta}{i_0}\cdot\frac{\mathrm{MBS}}{\lambda\, t_0} \tag{7.31}$$

When the asymptotic approximations are used in the network planning process, usually the upper bounds (7.25) and (7.28) for the tail-constants c_{CLP} and for the asymptotic behavior of BOR_c are used, since they represent the worst case. In case of the mFPT, the worst-case is determined by the lower bound (7.26).

7.2 Validation of the Asymptotic Results

Figure 7.1 compares the formulas listed in the previous section for the asymptotic behavior of the BOP and CLP with exact numerical results for a 3-Burst/M/1 model with truncated tails and a buffer-size $B = 10^4$. For the whole range of ρ and for all three blow-up regions, Eqs. (7.4) together with (7.13), (7.23) and (7.24), provide an excellent approximation for the true BOP. The true values for the BOP that are computed from the analytic 3-Burst/M/1 model with $\mathrm{MBS} = 1.3 \cdot 10^7$ cells are marked by '+' in Fig. 7.1. Only very close to the actual blow-up points

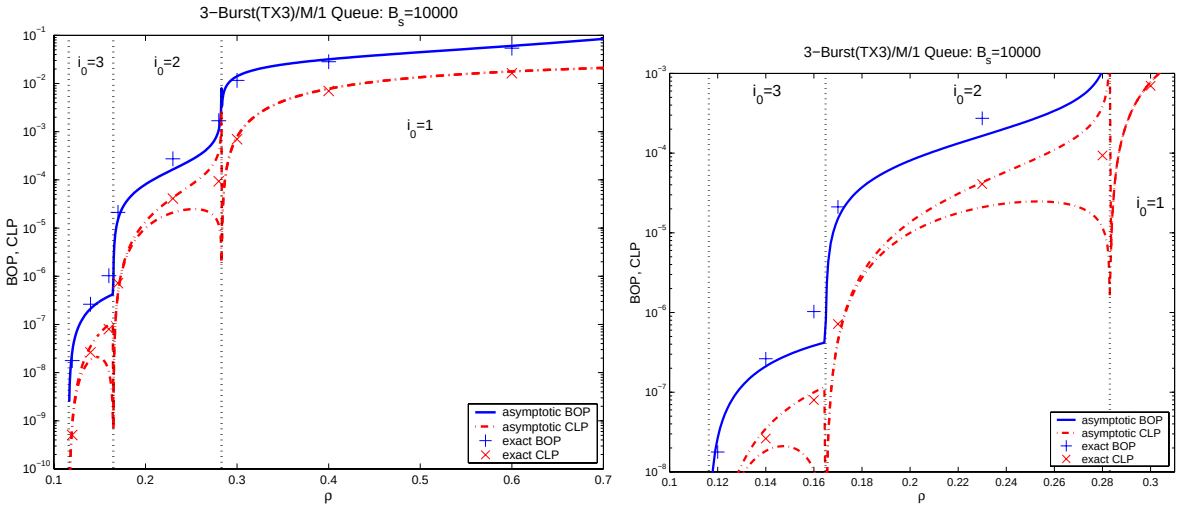

Figure 7.1: **Comparison of the asymptotic approximations with exact results for a 3-Burst(TX3)/M/1 Model:** The exact values of the BOP and CLP were computed for a truncated burst-length distribution with MBS $= 1.3 \cdot 10^7$ cells ($T = 12$ and $\theta = 0.15$). Both graphs show the same curves, but the right graph enlarges the interesting blow-up regions $i_0 = 3$ and $i_0 = 2$. The blow-up points are at $\rho_1 = 28.3\%$, $\rho_2 = 16.5\%$, and $\rho_3 = 1 - b = 11.6\%$. When $i_0 > 1$, the upper and lower bound for the asymptotic behavior of the CLP are shown by the dashed-dotted lines. The exact CLP of the 3-Burst/M/1/10^4 model turns out to be close to the upper bound of the CLP except for a close neighborhood around the blow-up point.

(when $i_\Delta \to 0$ or $i_\Delta \to 1$) does the approximation deviate from the exact results. However, note that in a small neighborhood around the blow-up points, the truncation of the burst-length distribution can show its impact.

The true values for the CLP (marked by '$\times$') in Fig. 7.1 are reasonably close to their asymptotic approximation, Eqs. (7.5) and (7.14), in the rightmost blow-up region $i_0 = 1$. For $i_0 > 1$, we have only been able to find an upper and lower bound for c_{CLP}, but the true values fall into that range. Those two bounds diverge considerably in the area close to the transition to the next lower blow-up region, when $i_\Delta \to 1$.

In summary, the knowledge of the asymptotic Power-Laws together with the approximations for the tail-constants turns out to be useful as long as the model does not operate too close to the blow-up points themselves ($i_\Delta > 5\%$ and $i_\Delta < 95\%$). For the mFPT and the transient overflow probability γ, the results of the comparison are not shown here, but the same conclusions hold.

7.3 Engineering Rules for Network Design

The asymptotic results of the previous sections for the N-Burst/M/1 queue can be used for the solution of practical problems in network design: for instance, for given traffic parameters, what choice of the switch parameters (ρ and B) is adequate? Alternatively, if the resources bandwidth and buffer-size at the switch are already given, what requirements have to be put on the incoming traffic, such that target values for the performance parameters (here i_0, mCD, CLP, BOP, γ) are achieved?

As the title of this section indicates, the numerical results of this section's formulas provide a

good guideline for a QoS-oriented choice of the model parameters. However, they should not be seen as a means of providing guaranteed QoS.

Within the set of the parameters that describe the traffic and the resources at the switch, the following parameters can be influenced at the network design stage (1, 2, and 3 in the list below) or by means of traffic policing (4 and 5):

1. *The output bandwidth ν of the switch:*
   With increasing ν, the utilization $\rho = \lambda/\nu$ decreases and the switch could be made to operate in a better blow-up region (i.e. larger i_0).

2. *The number N of multiplexed ON/OFF cell streams:*
   Multiplexing $N' > N$ traffic streams at the input port of a switch can also lead to a better blow-up region at the *same* utilization ρ. Note that this requires to increase the service-rate ν proportionally to N.

3. *The buffer-size B at the switch:*
   Buffer overflow or cell losses can also be reduced by increasing the buffer-size. However, the results in Chapters 5 and 6 show that in a 'bad' blow-up region (small i_0) with small effective PT exponent β, buffer-space is an inefficient means of improving QoS.

4. *The peak-rate λ_p of the cell streams:*
   So-called *traffic-shapers* reduce the peak-rate of the individual ON/OFF streams, thereby also reducing their burstiness, b.

5. *The Maximum Burst Size, MBS:*
   It is possible to impose an MBS on the incoming traffic stream as part of *traffic policing*. Cells within bursts that do not comply with the restriction will be discarded. A violation of the MBS restriction can usually be detected by some variation of a leaky bucket algorithm.

The remaining traffic parameters, the mean rate κ of the individual incoming link, the mean number of cells per burst $\overline{n_p}$, the Power-Tail exponent α of the burst-length distribution, and the tail-constant $c_{\mathrm{PT}}^{(1)}(\alpha)$ of the burst-length distribution cannot be changed easily in practice.

LRD ON/OFF traffic shows different blow-up regions. Frequently, if no exact target-values are specified for performance parameters, then operating in a sufficiently good (large i_0) blow-up region should provide satisfactory performance. In particular, blow-up region i has to be avoided ($i_0 > i$), where

$$i = \begin{cases} \left\lfloor \frac{1}{\alpha-1} \right\rfloor & \text{for bounded mean Delay}(\beta > 2). \\ \left\lfloor \frac{2}{\alpha-1} \right\rfloor & \text{for bounded variance of the Delay distribution}(\beta > 3). \end{cases}$$

Truncated Tails: When the burst-length distribution is truncated at MBS cells, the impact of any Power-Tails (with exponent $\beta(i)$, $i = i_0, ...,$) on the queue-length distribution has the PT Range

$$q_N = \frac{\rho - (1-b)}{\rho}\,\mathrm{MBS}\,.$$

Note that q_N is only meaningful in the blow-up regions, i.e. when $\rho > 1 - b$ (equivalently, $N\lambda_p > \nu$), since otherwise, the queue-length distribution decays geometrically. In the blow-up regions, $0 < q_N < \mathrm{MBS}$ holds due to the conditions $\rho > 1 - b$ and $0 < b < 1$.

The queue-length distribution and thus also the overflow probabilities drop off exponentially for $B > q_N$. An additional 'safety'-factor, $s > 1$ (e.g. $s = 10$), is advisable to ensure that the exponential drop-off has decayed far enough. Thus, a design goal could be

$$s \cdot q_N < B,$$

in which case the impact of the PT distribution only marginally affects buffer-overflow probabilities:

$$s \cdot q_N < B \quad \Longrightarrow \quad \mathrm{CLP}(B) < \mathrm{BOP}(B) \ll \exp(-s). \tag{7.32}$$

Capacity Planning

In this scenario, the task is to determine an adequate ρ (through an appropriate choice of ν) for given traffic parameters at the bottleneck switch.

Peak Rate Allocation:	$\rho < 1 - b$
Avoid blow-up region i :	$\rho < \dfrac{N}{N - (N-i)\,b}\,(1-b)$
Target BOP :	iterative algorithm using (7.4) and (7.13), (7.23), (7.24)
Target CLP :	iterative algorithm using (7.5) and (7.14), (7.25)
Target γ :	iterative algorithm using (7.12) and (7.15), (7.26)
Target mCD :	iterative algorithm using (7.9) and (7.16), (7.27)

If a truncation MBS of the Power-Tail distribution of the number of cells per burst exists, the utilization can also be kept low enough such that the BOP and CLP are in the region of their exponential drop-off (see Sect. 5.3) for the given buffer-size B:

$$\text{Restrict PT Impact, } s \cdot q_N < B: \qquad \rho < \frac{1-b}{1 - \frac{B}{s \cdot \mathrm{MBS}}}, \quad \text{if} \quad B < s \cdot \mathrm{MBS}$$

If $B > s \cdot \mathrm{MBS}$, the PT impact on the queue-length distribution is restricted to queue-lengths smaller than B for any ρ.

The asymptotic relationships for the BOP, CLP, γ and mCD do not present a continuous function of ρ due to the blow-up regions i_0. Therefore no closed-form expression can be given for the allowed utilization ρ, but instead an iterative algorithm has to be used, which for instance computes the asymptotic target performance parameter in the 'middle', of each blow-up region, i.e. for $i_\Delta = 0.5$, which corresponds to the following utilization values:

$$\tilde{\rho}_i = \left[1 + \frac{i - 0.5}{N}\,\frac{b}{1-b}\right]^{-1} \quad \text{for} \quad i = 1, ..., N\,.$$

At this point, the algorithm could just stop and use the largest $\tilde{\rho}_i$ at which the computed approximated performance parameter is below (better than) the target value. Such an approach is based on the argument that within the same blow-up region, performance does not change dramatically, so it is enough to pick the 'right' blow-up region in the planning algorithm.

Alternatively, another search for a slightly higher ρ value within the blow-up regions can be performed afterwards. Note that such a search algorithm should exclude the region close to the

blow-up points itself, since the asymptotic approximations are not very accurate there, or the upper bound can be too pessimistic, especially when i_Δ is small. Therefore, the search-algorithm should be restricted to $0.05 < i_\Delta < 0.95$, or even a smaller sub-interval.

Variations of this algorithm apply in most other scenarios below, except for the buffer-size planning.

Aggregation Level

Multiplexing a larger number N of traffic streams at the switch while keeping the utilization constant can result in great improvement of QoS due to a transition to a lower blow-up region. Note that a higher switch capacity is necessary with increasing N to keep the utilization constant.

$$\begin{aligned}
\text{Peak Rate Allocation:} &\quad \text{only if } \rho < 1-b \text{ regardless of } N \\
\text{Avoid blow-up region } i: &\quad N > i\,\frac{\rho}{1-\rho}\,\frac{b}{1-b} \\
\text{Target BOP, CLP, } \gamma\text{, mCD}: &\quad \text{iterative algorithm as above} \\
\text{Restrict PT Impact, } s\cdot q_{i_0} < B: &\quad \frac{N}{i_0(N)} > \left(1-\frac{B}{s\,b\,\mathrm{MBS}}\right)\frac{\rho}{1-\rho}\,\frac{b}{1-b} \text{ iteratively}
\end{aligned}$$

Since the the largest PT Range q_N does not depend on N, the PT-range q_{i_0} of the smallest, worst exponent β is restricted instead in this scenario.

Buffer-Size Dimensioning

Explicit asymptotic formulas exist if all parameters except for the buffer-size B are given, since the model operates in a fixed blow-up region i_0. Here, it is assumed that the switch does not operate in Peak Rate Allocation, $\rho > 1-b$ (equivalently, $nu < N\lambda_p$). First, use (7.1), (7.2), and (7.3) to compute the blow-up region i_0, the value of i_Δ, and the effective PT exponent β. Thereafter, (7.13), (7.23), or (7.24) provide an approximation for the tail-constant c_{BOP}. (7.14), (7.25), (7.15), and (7.26) provide approximations ($i_0 = 1$) or upper bounds ($i_0 > 1$) for the tail-constants c_{CLP} and c_{mFPT}.

Then the buffer-size B can be computed from the target performance parameters and the asymptotic relationships (7.4), (7.5), and (7.12):

$$\text{Target BOP}: \quad B > \left(\frac{c_{\mathrm{BOP}}}{\mathrm{BOP}}\right)^{\frac{1}{\beta-1}} \tag{7.33}$$

$$\text{Target CLP}: \quad B > \left(\frac{c_{\mathrm{CLP}}}{\mathrm{CLP}}\right)^{\frac{1}{\beta-1}} \tag{7.34}$$

$$\text{Target } \gamma(t_0): \quad B > \left(\frac{t_0}{c_{\mathrm{mFPT}}\,\ln\frac{1}{1-\gamma(t_0)}}\right)^{\frac{1}{\beta}} \tag{7.35}$$

For full Power-Tails in the burst-length distribution, Eqs. (7.33) to (7.35) have to be used. In case of truncated tails, the buffer-size can be chosen larger than the PT Range of the queue-length distribution:

$$\text{Restrict PT Impact:} \quad B > s\cdot q_N = s\cdot \mathrm{MBS}\,\frac{\rho-(1-b)}{\rho}$$

Since Equations (7.33) to (7.35) are applicable within the PT Range $q_{i_0} < q_N$ also for truncated tails, the resulting B of those formulas is the appropriate choice, if it is smaller than $s \cdot q_N$.

Shaping

Reducing the peak-rate λ_p of the bursts can result in a better blow-up region and thus in substantial performance gain. There exist devices – called traffic shapers – that perform such a task. However, one should be aware that there might be performance problems (delay and loss) that are caused by the shaper.

Since traffic shaping does not change the mean rate κ of the cell stream, a reduction of λ_p can be expressed in terms of the burstiness $b = 1 - \kappa/\lambda_p$.

$$\begin{aligned}
\text{Peak Rate Allocation:} \quad & b < 1 - \rho \\
\text{Avoid blow-up region } i: \quad & b < \frac{N}{N - (N-i)\,\rho}\,(1-\rho) \\
\text{Target BOP, CLP, } \gamma\text{, mCD:} \quad & \text{iterative algorithm as before} \\
\text{Restrict PT Impact, } s \cdot q_N < B: \quad & b < 1 - \rho\left(1 - \frac{B}{s \cdot \mathrm{MBS}}\right), \quad \text{if} \quad B < s \cdot \mathrm{MBS}
\end{aligned}$$

Traffic Policing via Maximum Burst Sizes

Finally, an MBS can be used to restrict the time-scales of the self-similar property of the network traffic. Thereby, the Power-Law behavior of several performance parameters is truncated. The goal in this scenario must be that the truncation occurs early enough, so that the target performance parameter is achieved. First, compute i_0, i_Δ, and β by Eqs. (7.1), (7.2), (7.3), respectively. Then, an approximation for an upper bound of the tail-constant c_{mCD} is obtained as in (7.27):

$$\overline{c_{\mathrm{mCD}}} \approx \frac{1}{\lambda} \cdot \frac{\rho}{1-\rho} \cdot \frac{i_\Delta}{i_0^{2-\beta}} \cdot \overline{n_p}^{\beta-1}\, b^2\, (1-b)^{i_0-1} \cdot \frac{\left[c_{\mathrm{PT}}^{(1)}(\alpha)\right]^{i_0}}{(2-\beta)\,(\alpha-1)^{i_0}}$$

Then, a given mean delay can be achieved by an appropriate choice of the MBS:

$$\text{Target mCD:} \qquad \mathrm{MBS} < \left(\frac{\mathrm{mCD}}{\overline{c_{\mathrm{mCD}}}}\right)^{\frac{1}{2-\beta}}, \quad \text{if} \quad \beta < 2$$

Alternatively, low BOP or CLP can be achieved by ensuring that the PT Range of the queue-length distribution (q_{i_0} for the PT exponent β, and q_N for the largest PT exponent $\beta_N = N\,[\alpha - 1] + 1$) is smaller than the buffer-size:

$$\begin{aligned}
\text{Restrict PT Impact, } s \cdot q_N < B: \quad & \mathrm{MBS} < \frac{\rho}{\rho + b - 1} \cdot \frac{B}{s} \\
s \cdot q_{i_0} < B: \quad & \mathrm{MBS} < \frac{i_0}{b\, i_\Delta}\,\frac{B}{s}
\end{aligned}$$

There is no explicit target BOP in the formulas. However, as mentioned before, the 'safety'-factor $s > 1$ implies the following upper bound:

$$s \cdot q_N < B \quad \Longrightarrow \quad \mathrm{CLP}(B) < \mathrm{BOP}(B) \ll \exp(-s).$$

Note that all the formulas hold asymptotically for network components with large buffers, so at least $B > 100$. Also, i_0 should be reasonably low ($i_0 \ll 50$), otherwise limiting models for highly multiplexed traffic might be more appropriate, see Sect. 10.2. Note that the latter restriction does not put bounds on the number of incoming links, N. For very bursty traffic (b close to 1), N can be very large but i_0 is still small, see Eq. (7.1).

7.4 Connection Admission Control

One of the major practical applications of stochastic modeling of network traffic is in the area of Connection Admission Control (CAC). The latter is necessary as soon as more sophisticated protocol mechanisms are developed which are based on a resource reservation approach, see also Sect. 10.7. Each connection announces its traffic parameters and its QoS requirements (e.g. upper bounds for delay or BOP) before it starts transmitting. Based on this knowledge, every switch along the data transmission path has to decide whether it has enough resources to handle that connection. If the answer is no, a different transmission path has to be found or the connections has to be rejected.

Current algorithms for CAC are frequently based on the *effective bandwidth approach*, see Sect. 7.7, or on experimentally derived tables that classify the scenarios in which connections have to be rejected.

Better suited algorithms – at least for connections that show ON/OFF behavior, e.g. persistent HTTP connections – can be developed from the results of the performance analysis of N-Burst/M/1 models in this thesis. Each connection has to specify its parameters κ, b, $\overline{n_p}$, α and MBS (more parameters are necessary if we allow for more general distributions than the TPT distributions of Sect. 3.4). In addition to its traffic parameters, the connection specifies its QoS requirements, here any combination of mCD, BOP (CLP), γ, and BOR_c (CLR_c).

Since the connection duration is finite, the transient parameter γ and the conditional overflow (loss) ratio BOR_c (CLR_c) are more appropriate, but hardly used in practice. In this section, we exemplarily discuss a CAC scenario with the target parameter γ, and we demonstrate that the wrong assumption of independent overflow events with probability BOP would lead to totally wrong results – in this case the results would be too pessimistic by far!

We discuss exemplarily the following scenario: An ATM-Switch with $\nu_1 = 155$ Mbit/sec has 2 incoming links each described by `TX3`-sources ($\kappa = 16.28$ cells/ms, $b = 0.884$, $\overline{n_p} = 9.1$, $\alpha = 1.4$). Those two incoming links cause a utilization, $\rho = 8.9\%$, of the switch.

Now, another Constant Bit Rate (CBR) connection with a large bandwidth requirement of $\tilde{\kappa} = 50$ Mbits/sec starts to use the switch, whose buffer is assumed to be empty at the start of the connection. The latter is reasonable, since the switch was not operating in any blow-up regions before the start of the connection. The duration of the connection is time t_0, and no priority queueing is used. During the duration of the connection, the utilization of the switch is raised to 41%.

We compute the probability that the connection is loss free, first using the steady-state $\mathrm{BOP}(B)$ and then compare that result with the transient probability $\gamma(t_0, B)$.

Since the connection has a constant bandwidth requirement, the number of cells in the buffer can be described by an approximating 2-Burst(TX3)/M/1 queue with reduced service-rate[1] $\nu = \nu_1 - \tilde{\kappa}$, although physically, not only cells from the incoming TX3-links are buffered, but also cells from the CBR connection. That 2-Burst/M/1 queue operates in blow-up region $i_0 = 2$ with $i_\Delta = 0.258$. The approximation (7.23) provides a numerical value for the tail-constant, $c_{\mathrm{BOP}} = 7.25\%$. Let the number of cells in the connection be $n(t_0)$. We will not look specifically at the cells of the CBR connection but we assume that if any overflows at all happen, cells of the CBR connection are also affected.

When making the (wrong) assumption that each cell independently experiences the same probability that it causes an overflow-event, the probability of having no overflows for all cells of the connection is:

$$p_0(t_0, B) = \left[1 - \frac{\lambda}{\lambda + \tilde{\kappa}} \mathrm{BOP}(B)\right]^{n(t_0)} \quad \text{where} \quad n(t_0) = \tilde{\kappa} \cdot t_0 \,.$$

Since $\mathrm{BOP}(B)$ is computed from the 2-Burst model with reduced service-rate ν, which has on average λt_0 arrivals in time t_0, it does not take into account the additional $\tilde{\kappa} t_0$ cells in the CBR connection. As a consequence, the computed value of BOP has to be scaled down by the factor $\lambda/(\lambda + \tilde{\kappa})$ to deliver an overflow probability for all cells.

$p_0(t_0, B)$ is the probability that no cells of the connection overflow; consequently, with probability $1 - p_0(t_0, B)$, at least one overflow occurs during the connection. p_0 uses the steady-state probability BOP, which does not take into consideration the correlation within overflow-events. Also, p_0 does not make use of the fact, that the buffer is empty at the start of the connection – which is of minor influence here, at least when t_0 is not very small.

On the other hand, the transient performance parameter $\gamma(t_0, B)$ captures both, the correlation of overflow events and the initially empty buffer. The following table computes the steady-state approximation $1 - p_0(t_0, B)$ and $\gamma(t_0, B)$ for several values of t_0 and various buffer-sizes B:

[1] Instead of deducting the bandwidth $\tilde{\kappa}$ from the service-rate, we could use a 2-Burst model with background Poisson rate $\lambda_0 = \tilde{\kappa}$. The asymptotic approximations are the same for both scenarios.

B	1000	5000	10^4	$5 \cdot 10^4$	10^5
BOP	$2.9 \cdot 10^{-4}$	$8.0 \cdot 10^{-5}$	$4.6 \cdot 10^{-5}$	$1.3 \cdot 10^{-5}$	$7.2 \cdot 10^{-6}$
$\frac{\lambda}{\lambda+\bar{\kappa}}$ BOP	$6.2 \cdot 10^{-5}$	$1.7 \cdot 10^{-5}$	$9.9 \cdot 10^{-6}$	$2.7 \cdot 10^{-6}$	$1.6 \cdot 10^{-6}$
$t_0 = 10\,\mathrm{s}$					
$1 - p_0(t_0, B)$	1	1	1	96.01%	84.27%
$\gamma(t_0)$	0.87%	$4.8 \cdot 10^{-4}$	$1.4 \cdot 10^{-4}$	$7.6 \cdot 10^{-6}$	$2.2 \cdot 10^{-6}$
$\mathrm{BOP}_c(t_0)$	3.3%	17%	33%	> 1	> 1
$t_0 = 1\,\mathrm{min}$					
$1 - p_0(t_0, B)$	1	1	1	1	1
$\gamma(t_0)$	5.1%	0.29%	$8.3 \cdot 10^{-4}$	$4.6 \cdot 10^{-5}$	$1.3 \cdot 10^{-5}$
$\mathrm{BOP}_c(t_0)$	0.57%	2.8%	5.5%	28%	55%
$t_0 = 5\,\mathrm{min}$					
$1 - p_0(t_0, B)$	1	1	1	1	1
$\gamma(t_0)$	23%	1.4%	0.41%	$2.3 \cdot 10^{-4}$	$6.6 \cdot 10^{-5}$
$\mathrm{BOP}_c(t_0)$	0.13%	0.56%	1.1%	5.5%	11%
$t_0 = 2\,\mathrm{h}$					
$1 - p_0(t_0, B)$	1	1	1	1	1
$\gamma(t_0)$	99.8%	29%	9.5%	0.55%	0.16%
$\mathrm{BOR}_c(t_0)$	$2.9 \cdot 10^{-4}$	$2.7 \cdot 10^{-4}$	$4.8 \cdot 10^{-4}$	0.23%	0.46%

The computation of BOP uses the asymptotic behavior (7.4), where $c_{\mathrm{BOP}} = 7.25\%$ is determined from the approximation (7.23). γ is computed by the exponential approximation (7.12). The lower bound in (7.26) is used as worst-case approximation for c_{mFPT}, $c_{\mathrm{mFPT}} = 4.6\,\mathrm{ms}$.

Obviously, the steady-state approach is of little value here, since it over-estimates the probability γ by far. If used for buffer sizing, the steady-state approximation leads to a large over-provision of buffer-space. For instance, for a target probability of 99% for an overflow-free connection of duration t_0, the necessary buffer-size can be computed by (7.33) and (7.35):

Connection length t_0	10 s	1 min	5 min	2 h
B for $p_0 = 99\%$	$6.8 \cdot 10^7$	$6.4 \cdot 10^8$	$4.8 \cdot 10^9$	$2.5 \cdot 10^{11}$
B for $1 - \gamma = 99\%$	923	$2.5 \cdot 10^3$	$6.1 \cdot 10^3$	$3.6 \cdot 10^4$

Finally, we compute the buffer-size $B_{min}(t_0)$ for which a minimal BOR_c results via Eq. (6.20):

Connection length t_0	10 s	1 min	5 min	2 h
$B_{min}(t_0)$	58	158	386	$2.3 \cdot 10^3$
$\mathrm{BOP}_c(B_{min}, t_0)$	0.37%	0.17%	$8.1 \cdot 10^{-4}$	$2.0 \cdot 10^{-4}$

Note that

$$\gamma(B_{min}(t_0), t_0) \approx 1 - \exp\left(\frac{t_0}{c_{\mathrm{mFPT}}\left[B_{min}(t_0)\right]^{\beta}}\right) = 1 - \exp\left(z_0(\beta)\right)$$

is constant (here 76.5%) for given β when (6.20) is used for the computation of the minimum. However, for smaller buffer-sizes, as obtained here for shorter t_0, the asymptotic approximations that are used here can fail.

7.5 Buffer Dimensioning

In this section, we discuss two specific scenarios in which optimal buffer-sizes are determined according to some QoS requirements. As in the last section, we compare the steady-state QoS parameter BOP with its transient counter-parts $\gamma(t_0, B)$ and $\mathrm{BOR}_c(t_0, B)$ for a whole busy-hour period of $t_0 = 5\,\mathrm{h}$. We assume full Power-Tails for the ON time distribution, thus no truncation impact is discussed.

The buffer-size planning is done via the asymptotic relationships as described in Eqs. (7.33) and (7.35) in Section 7.3. The necessary tail-constants are computed via the worst-case approximations in Sect. 7.1.

Traffic Shaper (1-Burst)

First, we discuss an application of the 1-Burst/M/1 model. Since there is no multiplexing involved, such a model corresponds either to a traffic shaper or to a hub whose bandwidth on the outgoing link is smaller than on the incoming link. Here, we assume a shaper that reduces the peak-rate λ_p of a `TX3` source by a factor of 4.3, such that the outgoing peak-rate is $\nu = 2\,\kappa$. At utilization $\rho = 0.5$, the shaper operates in blow-up region $i_0 = 1$:

N	λ, b, $\overline{n_p}$, α	ρ	i_0	i_Δ
1	`TX3`	0.5	1	0.869

The tail-constants c_{BOP} and c_{mFPT} of the asymptotic behavior of the BOP and mFPT can be determined either by Eqs. (7.13) and (7.15) or by direct numerical evaluation of the analytic 1-Burst/M/1 model with large truncation:

c_{BOP}	c_{mFPT}
2.057	0.172 ms

The necessary buffer-sizes that are obtained via (7.33) and (7.35) for different target values of the BOP and γ are listed in the following table:

Target	BOP			$\gamma(5\,\mathrm{h})$			$\min(\mathrm{BOR}_c)$
	10^{-6}	10^{-9}	10^{-12}	10%	1%	10^{-4}	
B	$6.1 \cdot 10^{15}$	$1.9 \cdot 10^{23}$	$6.1 \cdot 10^{30}$	$2.7 \cdot 10^{6}$	$1.4 \cdot 10^{7}$	$3.9 \cdot 10^{8}$	$3.1 \cdot 10^{5}$
mFPT	$6.8 \cdot 10^{10}$ y	$2.2 \cdot 10^{21}$ y	$6.8 \cdot 10^{31}$ y	47.5 h	20.7 d	5.7 y	2.3 h
$\gamma(5\,\mathrm{h})$	$8.4 \cdot 10^{-15}$	$2.7 \cdot 10^{-25}$	$8.4 \cdot 10^{-36}$	10%	1%	10^{-4}	88.21%
BOR_c	> 1	> 1	> 1	5.53%	28.24%	> 1	1.48%

Obviously, already a target BOP of 10^{-6} requires an unrealistically large buffer-size. Note that buffers larger than $B_{t_0} = t_0(\lambda_p - \nu) = 1.9 \cdot 10^9$ would not overflow during the busy hours, even if an over-saturation period lasts for the whole duration of the busy hours. For buffer-sizes $B = 3.9 \cdot 10^8$ and larger, the time interval of duration $t_0 = 5\,\mathrm{h}$ is not long enough to be able to observe the steady-state BOP, i.e. $\mathbb{E}\{\mathrm{BOR}\}(t_0)$ is substantially smaller than the BOP. This is indicated by the values of the computed approximation for BOR_c that are larger than 1. See Sect. 6.4.4 for a comparable scenario.

Multiplex Gain:

Next, we look at a switch in the network, whose $N = 8$ input ports are described by `TX3` sources. The switch capacity was chosen to yield the same utilization $\rho = 0.5$ as in the previous section.

N	λ, b, $\overline{n_p}$, α	ρ	i_0	i_Δ	c_{BOP}	c_{mFPT}
8	`TX3`	0.5	2	0.950	1.34	3.02 ms

Here c_{BOP} is computed with the approximation in formula (7.23). As usual, the lower bound of Eq. (7.26) is used as worst-case approximation of the c_{mFPT}.

Target	BOP			$\gamma(5\,\mathrm{h})$			$\min(\mathrm{BOR}_c)$
	10^{-6}	10^{-9}	10^{-12}	10%	1%	10^{-4}	
B	$4.5 \cdot 10^{7}$	$2.6 \cdot 10^{11}$	$1.4 \cdot 10^{15}$	$2.6 \cdot 10^{5}$	$9.7 \cdot 10^{5}$	$1.3 \cdot 10^{7}$	$6.1 \cdot 10^{4}$
BOP	$1.0 \cdot 10^{-6}$	$1.0 \cdot 10^{-9}$	$1.0 \cdot 10^{-12}$	$6.2 \cdot 10^{-5}$	$2.2 \cdot 10^{-5}$	$2.8 \cdot 10^{-6}$	$2.0 \cdot 10^{-4}$
mFPT	58 y	$3.3 \cdot 10^{8}$ y	$1.8 \cdot 10^{15}$ y	47.5 h	20.7 d	5.7 y	3.4 h
γ	$9.8 \cdot 10^{-6}$	$1.7 \cdot 10^{-12}$	$3.1 \cdot 10^{-19}$	10%	1.0%	$1.0 \cdot 10^{-4}$	76.5%
BOR_c	10%	> 1	> 1	$6.2 \cdot 10^{-4}$	0.22%	2.8%	$2.6 \cdot 10^{-4}$

Note that a multiplex gain is obvious: the necessary buffer-sizes are much smaller than for the single source model at the same utilization, see previous table. Such an enormous multiplex-gain (B smaller by a factor of at least 10^8 for the same target BOP values) is mainly a consequence of the better blow-up region $i_0 = 2$.

7.6 Steady-State and Transient Delay

The transient probability $\gamma(t_0, B)$ has a straightforward interpretation with respect to overflow-events. However, it can also be used to explain the fluctuations that are observed in measurements of per-cell delay in finite observation intervals: Large delay values are caused by long queue-lengths, and the probability that long queue-lengths n occur in an individual observation interval is again the transient probability $\gamma(t_0, n)$ (here, n is smaller than the buffer-size but still large). In the following, we illustrate the relationship between γ and transient per-cell delay by means of a practical example.

We look at the mean delay for a switch whose input is described by the aggregation of two `TX3` sources:

N	λ, b, $\overline{n_p}$, α	ρ	i_0	i_Δ	c_{BOP}	c_{mFPT}
2	`TX3`	0.15	2	0.509	15%	1.5 ms

As usual, the approximation for c_{BOP} is obtained by Eq. (7.23). However, this time we do not use the worst-case lower bound for c_{mFPT}, but an intermediate approximation between the two bounds ($1.15 \leq c_{\mathrm{mFPT}} \leq 1.75$). The deviation of the two bounds is not large, so any other value in that range does not change the obtained numerical results greatly: as already mentioned in the beginning of this chapter, the uncertainties in traffic forecasting outweigh the inaccuracy of the used approximations by far.

Note that the steady-state mean delay for our scenario is infinite when assuming infinite buffers, since $\alpha < 1 + 1/i_0 = 1.5$. In any finite observation interval however, only a finite range of queue-lengths is observed, which corresponds to finite mean delay values for truncated tails. With the help of the transient parameter γ, we discuss in the following how to relate steady-state mean delay values for truncated tails to observed mean delays in finite observation intervals.

If we truncate the burst-length distribution at MBS cells, the Power-Tail Range of the queue-length distribution is limited to

$$q_{i_0}(\mathrm{MBS}) = b \cdot \frac{i_\Delta}{i_0} \cdot \mathrm{MBS} = 0.227\,\mathrm{MBS},$$

see Sect. 5.3. So even in an infinite-buffer system, only queue-lengths up to

$$q_{i_0}(\lambda_p t_0) = 5.7 \cdot 10^8 \text{ cells}$$

would occur with non-negligible probability, if bursts longer than the busy-hour period, $t_0 = 5\,\mathrm{h}$, are neglected.

The following table computes amongst others the steady-state mCD of the 2-Burst model for different truncations T. Only the values in the column for $\gamma(t_0 = 5\,\mathrm{h})$ are computed by the asymptotic behavior (7.12) using $c_{\mathrm{mFPT}} = 1.5\,\mathrm{ms}$. All the other computed delay values are results from the exact Matrix-Analytic computation of 2-Burst/M/1 queues:

Truncation	MBS	$q_{i_0}(\mathrm{MBS})$	$\gamma(5\,\mathrm{h}, q_{i_0})$	mCD	$\mathrm{CD}_{90\%}$	$\mathrm{CD}_{99\%}$	$\mathrm{CD}_{1-10^{-5}}$
$T = 42$	$2.1 \cdot 10^9$	$4.8 \cdot 10^8$	$2.9 \cdot 10^{-9}$	$127.9\,\mu\mathrm{s}$	$35.0\,\mu\mathrm{s}$	$186.3\,\mu\mathrm{s}$	$723.1\,\mathrm{ms}$
$T = 40$	$8.0 \cdot 10^8$	$1.8 \cdot 10^8$	$1.6 \cdot 10^{-8}$	$107\,\mu\mathrm{s}$	$35.0\,\mu\mathrm{s}$	$186\,\mu\mathrm{s}$	$695\,\mathrm{ms}$
$T = 35$	$6.7 \cdot 10^7$	$1.5 \cdot 10^7$	$1.4 \cdot 10^{-6}$	$69.0\,\mu\mathrm{s}$	$35.0\,\mu\mathrm{s}$	$186\,\mu\mathrm{s}$	$575\,\mathrm{ms}$
$T = 20$	$4.1 \cdot 10^4$	$9.3 \cdot 10^3$	58%	$23.8\,\mu\mathrm{s}$	$34.8\,\mu\mathrm{s}$	$169\,\mu\mathrm{s}$	$41.4\,\mathrm{ms}$

As a comparison, note that the steady-state mCD in an M/M/1 model with same values of ρ and ν is $\mathrm{mCD}_{MM1} = 5.4\,\mu\mathrm{s}$.

Since $\alpha = 1.4$ and $i_0 = 2$, the queue-length distribution at cell-arrivals is Power-Tailed with tail-exponent $\beta = 1.8$ in its density function, see Eq. (7.3). Consequently, its steady-state mean is infinite when full PTs are used, and the asymptotic behavior of mCD(MBS) is given by Eq. (7.9):

$$\mathrm{mCD}(\mathrm{MBS}) \sim c_{\mathrm{mCD}}\,\mathrm{MBS}^{2-\beta}\,.$$

For this model, the Matrix-Analytic computation of the analytic 2-Burst model with truncation $T = 42$ provides a tail-constant of

$$c_{\mathrm{mCD}} = \frac{\mathrm{mCD}(T = 42)}{(\mathrm{MBS}_{T=42})^{0.2}} = 1.75\,\mu\mathrm{s}\,.$$

Note that Eq. (7.27) provides an upper bound of $2.57\,\mu\mathrm{s}$, which is reasonably close to the computed value above.

The steady-state mean delay can be made arbitrary large by increasing the MBS. A better measure for the delay are the p%-quantiles, $\mathrm{CD}_{p\%}$, i.e. the delay value such that only $1 - p\%$ of all cells experience a larger delay. The cell delay quantiles converge for any p, since the steady-state cell delay distribution is well defined for any T – also for $T = \infty$.

If one insists on the mCD as the performance parameter, the question is, what truncation is the right one to use with respect to the finite duration t_0 of the observation interval? Clearly, only

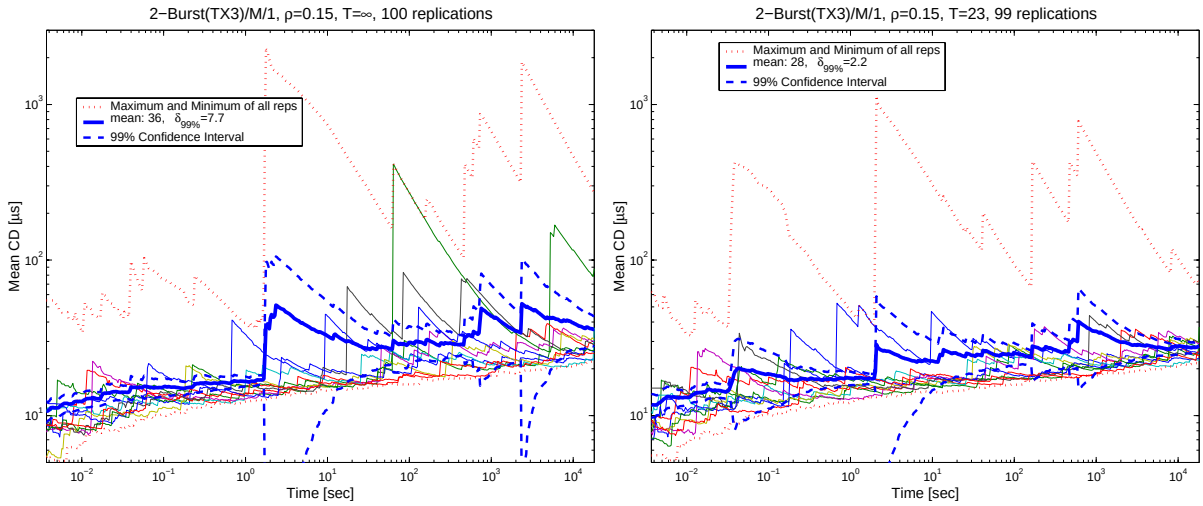

Figure 7.2: **Log-Log Plot of simulated mean per Cell Delay of 2-Burst/M/1 Model:** Plotted are 10 of the 100 independent replications (thin lines) together with the mean of all replications and a 99% confidence interval for the estimated mean. Within each replication the average per cell delay of all the cells that were served until time t is plotted.

burst up to length t_0 – or equivalently up to MBS $= \lambda_p t_0 = 2.5 \cdot 10^9$ cells – have to be considered in any case. The first row of the previous table where $T = 42$ corresponds to that worst case.

However, the simulation experiment that is shown in Fig. 7.2 suggests much lower values for the mean delay: 100 independent replication with each on average about $2 \cdot 10^9$ cells resulted in the estimate,

$$\widehat{\text{mCD}} = (36 \pm 7.7)\,\mu\text{s} \quad \text{at } 99\% \text{ confidence level,}$$

which is far below the $128\,\mu$s of the steady-state mCD for $T = 42$. So where is the catch?

A truncation of $T = 42$ affects buffer-sizes of up to $q_{i_0}(\text{MBS}_{T=42}) = 4.8 \cdot 10^8$ cells. However, the probability that the tail up to that size can show its impact in a single run is very low, namely $\gamma(t_0, q_{i_0}) \approx 3 \cdot 10^{-9}$. Consequently, a simulation experiment would have to cover more than 10^8 replications to have a good chance to also include a few potentially very bad runs that have a very strong impact on the mean delay.

Within the 100 runs, there is a high probability, that buffer-sizes 'only' up to a several 10^5 cells are actually used.

Via Eq. (7.7), such buffer-sizes correspond to MBS values for truncations between $T = 23$ and $T = 28$, computed in the next table:

Target	$\gamma(5\text{h}) \approx 10\%$	1%	10^{-3}	minimal BOR_c
B	$3.0 \cdot 10^4$	$1.1 \cdot 10^5$	$4.0 \cdot 10^5$	$3.0 \cdot 10^3$
T with $q_{i_0}(\text{MBS}_T)$ close to B	23	26	28	18
$\text{MBS}(T)$	$1.8 \cdot 10^5$	$7.8 \cdot 10^5$	$2.1 \cdot 10^6$	$1.5 \cdot 10^4$
$\gamma(t_0, q_{i_0}(\text{MBS}_T))$	6.0%	0.43%	$7.2 \cdot 10^{-4}$	
mCD	$28.3\,\mu$s	$34.4\,\mu$s	$39.6\,\mu$s	$21.5\,\mu$s
99%-quantile	$177\,\mu$s	$181\,\mu$s	$183\,\mu$s	$162\,\mu$s
$(1 - 10^{-5})$-quantile	96.6 ms	193 ms	277 ms	21.9 ms

The table shows that the models with a truncation in that range deliver steady-state values close to the observed delay in the simulation experiment.

In summary, the transient probability γ turns out to be of great value for understanding transient mean delay as well. Only with very small, but non-negligible probability, very bad delay values occur within a single observation period of length t_0. Unless a huge number of independent replications is performed, such bad runs are usually not obtained and mean delay is underestimated. However, the potential for such large delay is there. Thus simulation experiments have to be interpreted with care.

7.7 Effective Bandwidths

The concept of *effective bandwidths* is frequently used in practice for Connection-Admission Control (CAC), see [Kesidis et al. 93]. The underlying idea is that each individual

connection, which announces its mean cell-rate, κ, and peak cell-rate, λ_p, at connection setup-time, is assigned a so-called effective bandwidth, k, where $\kappa \leq k \leq \lambda_p$. If the sum of the effective bandwidths of all the incoming flows does not exceed the available bandwidth, the new connection is accepted.

The value of k depends on the required QoS-level. For $k = \lambda_p$ (called *peak-rate reservation*), no QoS problems are to be expected, but on the other hand, no multiplexing gain can be achieved. At the other end of the possible range for k, at $k = \kappa$ (*mean-rate reservation*), highest possible multiplex-gain will be achieved, but the provision of QoS will be difficult, since essentially the network is allowed to operate at $\rho = 1$ then.

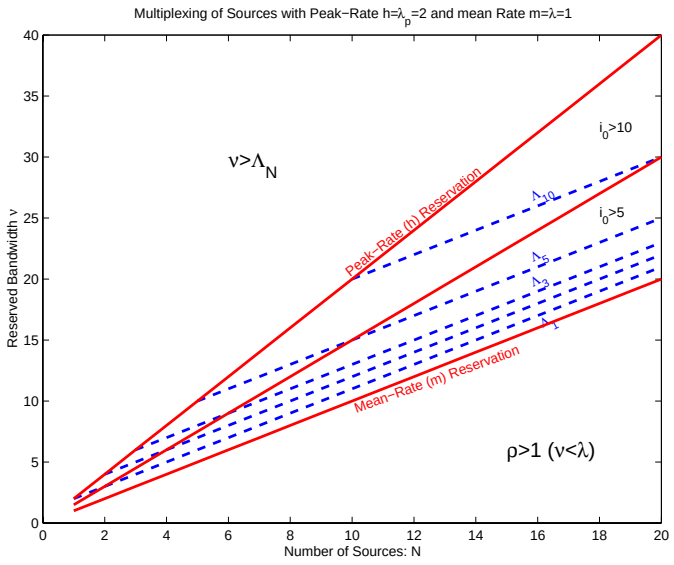

Figure 7.3: **Effective Bandwidth Reservation:** For each of th N sources in the N-Burst [IS] model an effective bandwidth k is reserved, i.e. the service-rate is guaranteed to be $\nu = N\,k$. The solid lines show the reserved bandwidth for a growing number of sources, N, for different choices of $\kappa \leq k \leq \lambda_p$. The dashed lines that are marked by Λ_i indicate the amount of bandwidth that is necessary to obtain a blow-up region $i_0 > i$.

Figure 7.3 investigates the impact of the effective-bandwidth reservation scheme on the blow-up regions in an N-Burst/M/1 model. The dashed lines indicate the amount of bandwidth that is

necessary to avoid ($i_0 > i$) a certain blow-up region i. To avoid any impact of Power-Tails in the LRD model, $\nu > \Lambda_N = N\lambda_p$ is necessary, i.e. PRA with the drawback of no multiplex gain.

As Figure 7.3 and Equation (5.1) clearly show, the separation lines between the blow-up regions is proportional to the number of sources, N:

$$\Lambda_i = i(\lambda_p - \kappa) + N\kappa.$$

Thus, the amount of reserved bandwidth that is necessary to avoid a certain blow-up region i grows increases linearly with slope κ in Fig. 7.3.

The effective bandwidth method with $k > \kappa$ leads to better blow-up regions (higher i_0) with growing number of sources, since the reserved bandwidth shows slope k, thus larger than κ. So any choice of $k > \kappa$ eventually leads to a large i_0 when enough of these sources are multiplexed. Therefore, asymptotically (for large enough N), the effective bandwidth approach works for LRD N-Burst traffic. However, it is not really the most suitable approach for a low number of connections/sources. A CAC algorithm that is based on asymptotic results for N-Burst/M/1 queues, as in Sect. 7.4 is definitely superior in such scenarios.

For given $\alpha = 1.4$, such as estimated for the Ethernet data of [Leland et al. 94], $i_0 > 5$ is sufficient to avoid the unbounded growth of both, the mean and the variance of the Cell-Delay distribution. So multiplexing more than 10 sources of Figure 7.3 with $k = (\kappa + \lambda_p)/2$ would be sufficient to avoid most negative effects of such LRD traffic in terms of the first two moments of Cell Delay (for Cell Loss Probabilities, the impact of the Power-Tails only vanishes for $\nu > N\lambda_p$, but increasing i_0 still helps since it increases the effective tail-exponent β).

Chapter 8

Summary of Results

This chapter presents a brief summary of the major results of the performance analysis together with a discussion of the practical relevance of those results. It is addressed to the network engineer and as such it only summarizes results that have a straightforward impact on the network planning or management task. Other discussions in Chapters 2 – 7 are left out, if they are of more theoretical interest or if they have not yet obtained a matureness that allows their straightforward practical application. Hence, this chapter only summarizes a subset of Chapters 2 – 7.

8.1 Motivation

Stochastic modeling and queueing theory have finally found a broad field of application: design, analysis, and management of telecommunication systems. However, performance models for packet or cell-switched data networks are far from trivial. So far, the simple M/M/1 model is one of the few models that is being used in algorithms for planning tools. However, it is now widely accepted that standard Poisson models do not adequately describe network traffic since they do not take into account its inherently bursty nature. This is accentuated by a series of measurements in the last decade (one of the earliest discussed by [Leland et al. 94]) that have shown that network traffic often exhibits a so-called *self-similar* property: It shows burstiness over a wide range of time-scales.

According to several measurements of network traffic, one such series done in [Gogl 98] at the entrance of the B-WiN network, multiplexed ON/OFF traffic with LRD properties appears to be a good model for real data traffic. This thesis develops techniques for and provides results from the analysis of queueing models with such traffic as input. Thereby, insights are gained into how QoS problems at single 'bottleneck' switches in the network occur and what remedies are possible.

The performance analysis at the bottleneck switches starts off with an investigation of the steady-state performance parameters *mean Cell Delay* (mCD), *Cell Loss Probability* (CLP), and *Buffer Overflow Probability* (BOP), see Chapter 5. The analysis showed that QoS does not reduce uniformly with, for instance, decreasing service rate. Instead, well-defined so-called *blow-up points* have been observed at which the QoS parameters deteriorate dramatically. The blow-up points were located and it is shown in Chapter 7 how their location can be used in network design.

The analysis of the steady-state behavior provides some fundamental insight into the peculiarities of multiplexed ON/OFF traffic with LRD properties. However, the long lasting correlation in the cell-arrival process, together with the finite duration of the 4-8 busy hours in the daily profile, could have the effect that steady-state behavior for the performance-critical time period might never be observed in practice. The results of the transient analysis in Chapter 6 confirm such

doubts. Furthermore, the steady-state BOP does not express the correlated nature of overflow events, which is especially accentuated in systems with large buffers and by traffic with LRD properties. Additional transient parameters are discussed in Chapter 6 in order to provide a better description of the switch's behavior.

The models in this thesis are discussed within the context and with the terminology (e.g. cells, switches) of ATM networks. However, the results are transferable to other packet based network technologies, including IP, as long as there is no feedback between the network behavior and the traffic sources. The latter requirement excludes protocols with flow-control mechanisms, such as TCP/IP. However, Chapter 9 discusses an approach to include such feedbacks.

8.2 Performance Model

8.2.1 N-Burst Arrival Process

The development of a realistic and flexible description for data traffic in modern telecommunication networks is the focus of Chapter 2. A brief summary of the N-Burst Independent Source model is given here, since most parts of the analysis are based on that particular model, although the results can be generalized to more general N-Burst models as well, see e.g. Section 5.1.

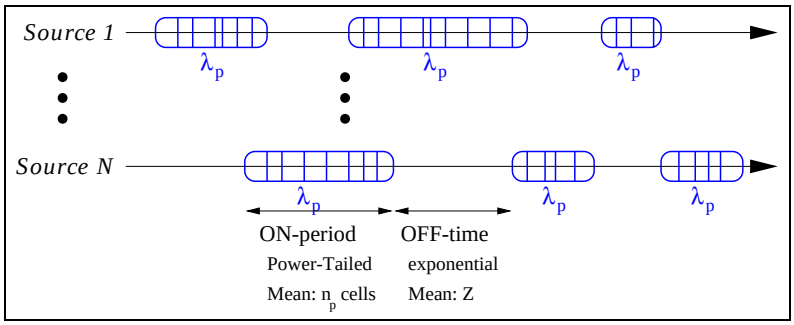

Figure 8.1: **The N-Burst Arrival Process:** Cells from N ON/OFF sources are multiplexed together.

The *N-Burst* arrival process is a superposition of traffic streams from N independent, identical sources of ON/OFF type, as shown in Fig. 8.1: each source emits cells at a Poisson-rate λ_p (peak-rate) during its ON-time (a *burst*), and then transmits nothing during its OFF-time. Let κ be the mean rate of the individual source (the average for the ON- and OFF-times together), then the N sources collectively generate cells at the mean rate $\lambda = N\kappa$. In Chapter 2, this arrival process is called an *N-Burst Independent Source* model, or more precisely, the *N-Burst* $[IS/G/N//N]$ model. Herein, we will refer to it simply as the 'N-Burst model' even though that name describes a more general class of models.

The 'degree of burstiness' of the individual sources can be expressed by the parameter

$$b := 1 - \frac{\kappa}{\lambda_p} \cdot \tag{8.1}$$

Consequently, $0 < b < 1$, and the higher b, the more bursty the traffic of the single source is. At the boundary value $b = 0$, the traffic has no ON/OFF structure and reduces to pure Poisson traffic, while at the other end, $b \to 1$, the peak-rate λ_p approaches infinity which results in a so-called bulk-arrival model, see Sect. 5.6.

Queueing Model

The output port of the bottleneck ATM-switch is modeled as an exponential server with rate ν. In many scenarios we look at buffer-overflow events for a primary buffer of size B and an infinite secondary buffer, i.e. the queueing model is of type N-Burst/M/1 where overflow-events are recognized by the threshold B in the infinite buffer.

Since in many cases, lost cells are retransmitted by higher protocol layers, the assumption of infinite back-up buffers is practically meaningful. However, note that some of the parameters used later on, in particular the *First Passage Times* and the *Probability of at least One Overflow* do not change when assuming a finite-buffer system.

Even when an infinite back-up buffer is assumed, the N-Burst/M/1 queue is stable as long as $\rho = \lambda/\nu < 1$. A well-defined steady-state queue-length distribution exists in the stable case, although its moments could well be infinite. In this chapter, we only consider queues with $\rho < 1$, or equivalently $\nu > \lambda$.

8.2.2 Truncated Power-Tail Distributions

The ON/OFF behavior of the individual sources causes bursty traffic at the cell-level. However, it should be emphasized that ON/OFF behavior by itself is not sufficient to produce LRD properties. The distribution of the duration of the ON periods is critical.

Measurements of files in large file-systems as well as measurements of files that were transfered over the Internet by the `http` protocol (see [Crovella & Bestavros 96]) both suggest that the size of such files is more adequately described by so-called Power-Tail (PT) distributions, which have the following asymptotic form of their *reliability function* $R(x)$:

$$R(x) := \mathbb{P}\{X > x\} \sim \frac{1}{x^\alpha} \quad \text{for large } x. \tag{8.2}$$

α is called the *Power-Tail (PT) exponent*. Thereby, $f(x) \sim g(x)$ is a short-hand notation for $\lim_{x\to\infty} f(x)/g(x) = c_{\mathrm{PT}}$ with $0 < c_{\mathrm{PT}} < \infty$. The log-log plot of such a function $R(x)$ looks linear with slope $-\alpha$, since $\log(x^\alpha) = \alpha \cdot \log(x)$.

In contrast to the widely used exponential distributions, the probability of very large events does not vanish quickly in PT distributions. One of the consequences is that such distributions can have infinite moments, in particular, their mean is infinite if $\alpha \leq 1$, and their variance is infinite if $\alpha \leq 2$. If Power-Tails with infinite variance are used for the ON-time distribution in the N-Burst model, asymptotically second order self-similarity in the inter-arrival process and in the counting process is the result, see Sect. 3.1.

Even if infinite Power-Tails in the burst-length distribution are real, for the analysis of the transient behavior, the tail matters at most up to the duration of the observation period. So naturally, truncated tails come into play, when finite time-intervals are considered. In our N-Burst model, we use *truncated* PT (TPT) distributions as introduced by [Greiner et al. 99]: their reliability function, $R_T(x)$, is a sum of T exponential distributions with geometrically increasing mean, weighted by geometrically decreasing factors (see Appendix B.2). $R_T(x)$ shows Power-Law behavior for a limited range of x until some $\overline{x_T}$ and drops off exponentially thereafter as shown in Fig. 8.2, left. That exponential drop-off is called truncation and physically corresponds to a *Maximum Burst Size* (MBS) in the N-Burst model. See Section 3.5 for a more rigorous definition of the *Power-Tail Range*, $\overline{x_T}$, which locates the exponential drop-off.

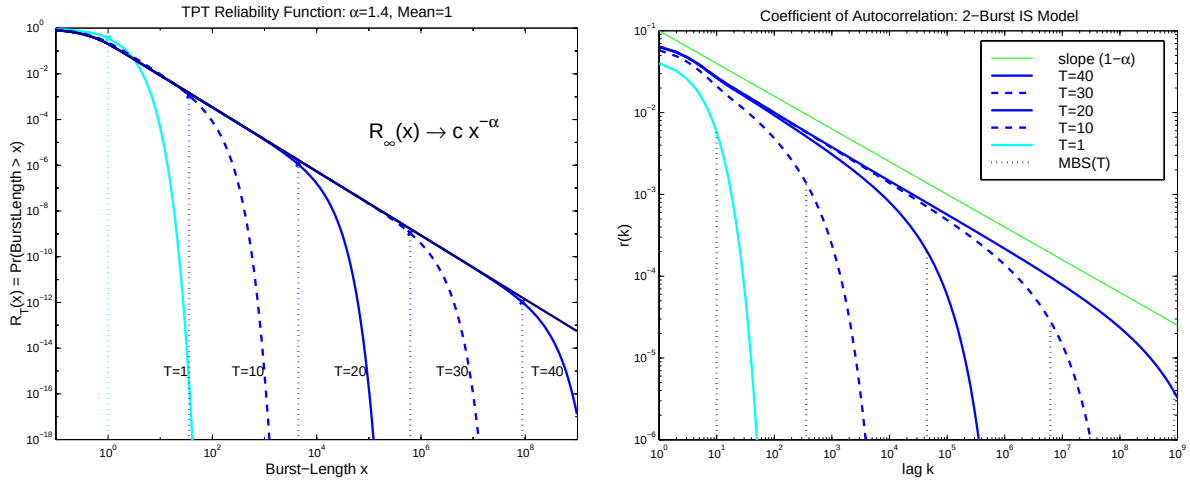

Figure 8.2: **Log-Log Plot of the Reliability Function, $R_T(x) = \mathbb{P}\{X > x\}$, for Truncated Power-Tail (TPT) Distributions with T Phases (Left) and the Resulting Autocorrelation Structure of a 2-Burst Model:** Both the TPT reliability function and the N-Burst autocorrelation function show a Power-Law behavior – appearing linear with negative slope α and $[\alpha - 1]$ respectively – for some range of x or k, before dropping off exponentially. The Power-Tail Range, $\overline{x_T}$, (dotted lines in left graph) can be extended arbitrarily by using more phases.

Using the TPT distribution for the burst-length distribution in the N-Burst model, the autocorrelation function $r(k) = \mathrm{Cov}(X, X_{+k})/\sigma^2$ of the inter-cell times also shows a truncated Power-Law behavior, see Fig. 8.2, right. Due to the finite PT-Range of the autocorrelation function, strictly speaking, the traffic is not LRD according to the mathematical definition in Sect. 3.1. Therefore, we talk of *traffic with LRD properties* instead.

8.2.3 Model Calibration

It is goal of this thesis to study the performance impact of the LRD properties in realistic settings. Therefore, methods for parameter estimation of a 1-Burst model were developed in Chapter 4 and applied to a set of measurements that is described in more detail in [Gogl 00]. Most parametric studies in this thesis use N-Burst processes that are calibrated on real measured network traffic: a particular measurement, called `TX3`, of IP over ATM traffic that was sent into the B-WiN backbone at the access point in Munich was basis for this calibration. Measurement `TX3` had the following statistical properties,

Measurement	`TX3`
Number of cells:	784 652
Duration:	48.2 s
Mean inter-cell time:	$\bar{X} = 61.43\,\mu$s
Variation of inter-cell times:	$C^2(X) = 13.67$

where $C^2(X)$ is the coefficient of variation, see Sect. 1.3. The analysis of the measurement `TX3` leads to the following parameters for a single-source 1-Burst model, see Sect. 4.3 and [Schwefel et al. 97]:

1-Burst Calibration	TX3
Average cell rate:	$\kappa = 16.3\,\text{cells/ms} = 6.90\,\text{Mbit/s}$
Cell-rate during burst:	$\lambda_p = 140\,\text{cells/ms} = 59.4\,\text{Mbit/s}$
Mean number of cells per burst:	$\overline{n_p} = 9.1$
Mean duration of OFF time:	$Z = 494\,\mu\text{s}$
Burstiness:	$b = 1 - \kappa/\lambda_p = 88.37\%$

This parameter set is the result of an earlier calibration method in [Schwefel et al. 97]. However, the parameters are close enough to the 'correct' parameter set that is obtained by the improved method in Sect. 4.3. Therefore, we refer to it as TX3 although this is not quite consistent with Sect. 4.3. The 1-Burst calibration for other measurements in Appendix I.2 resulted in comparable parameter sets, in particular the burstiness parameter showed a range of $0.84 < b < 0.93$. In that sense, TX3 can be seen as representative. However, we neglect the smooth background Poisson rate λ_0 here, since it contributes only a very small proportion of the overall traffic, see Sect. I.2.2. The formulas in this chapter assume $\lambda_0 = 0$, but any positive $\lambda_0 > 0$ can be deducted from the service-rate ν, when the asymptotic behavior of N-Burst/M/1 models in the blow-up regions is investigated.

The exponent of the reliability function of the burst-length distribution is set at $\alpha = 1.4$ (which corresponds to a Hurst parameter $H = 0.8$). Note that the cell-rate, λ_p, during a burst is more than eight times bigger than the average cell-rate, κ although the traffic at the measurement point is highly aggregated: in the order of 100 to 300 simultaneous IP connections contribute their ATM cells to the aggregated cell stream. Nevertheless, several properties of the measured data indicate that such an aggregated cell-stream in this scenario can be modeled by a single ON/OFF source.

8.3 Performance Impact

8.3.1 Steady-State Results

This section summarizes the performance results that were derived from the steady-state analysis of N-Burst/M/1 queues. See Chapter 5 and the publications [Lipsky et al. 99a], [Schwefel & Lipsky 99b], [Schwefel & Lipsky 99c], and [Schwefel & Lipsky 99d] for the details and additional discussions.

The analysis of the steady-state queue-length distribution of N-Burst/M/1 queues in Chapter 5 shows that the mean queue-length and thus the mean delay exhibit peculiar jumps (blow-ups) at N distinct utilization values ρ_i, $i = 1, ..., N$, called blow-up points. It turns out that the blow-ups of the mean delay occur, when the service-rate, ν, equals the average rate Λ_i when i of the N sources are simultaneously in a long-term ON period, while the remaining $(N - i)$ sources show average behavior, each with cell-rate κ:

$$\Lambda_i := i \cdot \lambda_p + (N - i) \cdot \kappa = \lambda + i\,(\lambda_p - \kappa) = \kappa \cdot \left(N + i\,\frac{b}{1-b} \right), \quad i = 0, ..., N\,, \tag{8.3}$$

Therefore, the blow-ups $i = 1, ..., N$ occur at utilization:

$$\rho_i = \frac{\lambda}{\Lambda_i} = \frac{1}{1 + \frac{i}{N} \cdot \frac{b}{1-b}}\,. \tag{8.4}$$

The larger the PT Range of the burst-length distribution is, the more pronounced are those blow-up effects, reproduced in the left graph of Fig. 8.3.

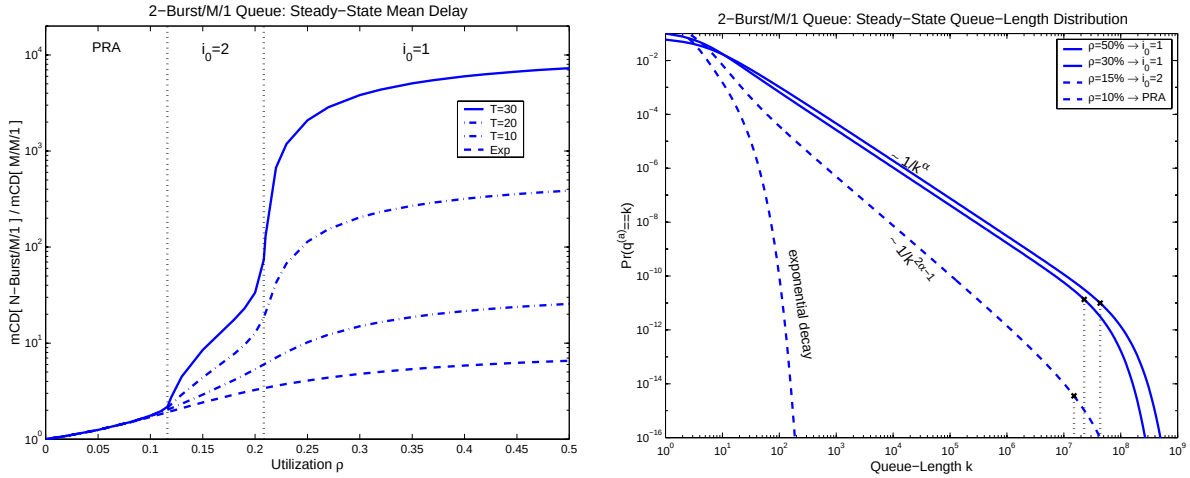

Figure 8.3: **Blow-up Regions in the Parameter Space of the 2-Burst/M/1 Queue:** For exponential burst-lengths, the mean delay of the 2-Burst/M/1 queue relatively to the delay of an M/M/1 queue with same ρ increases gradually with ρ. For TPTs, the relative mean delay blows up for two particular values of ρ, see left graph. The cause of this behavior is shown in the right graph: the queue-length distribution is Power-Tailed if $N\lambda_p > \nu$, but the PT exponent changes at the blow-up points.

The cause of the blow-up effects is discussed in Sect. 5.2, here illustrated again in the right graph of Fig. 8.3: Let

$$i_0 := \min\{i \mid \Lambda_i > \nu\} \quad [\Longrightarrow \quad \Lambda_{i_0-1} \leq \nu < \Lambda_{i_0}] \tag{8.5}$$

be the minimum number of sources that are sufficient to over-saturate the server if they are in a long-lasting ON period. In that scenario, the queue-length distribution turns out to be Power-Tailed with the exponent

$$\beta := \beta(i_0, \alpha) = i_0(\alpha - 1) + 1\,, \tag{8.6}$$

in its *distribution* function (hence exponent $\beta - 1$ in its reliability function). The important property is that the exponent β changes at the blow-up points. The reason for such behavior is that the long queue-lengths are mainly due to long over-saturation periods, which require at least i_0 sources to be permanently active, see Sect. 5.1. The duration of such over-saturation periods is shown in Section 3.3 to be Power-Tailed with the exponent β as in (8.6).

The equivalent effect of changing PT-exponents β is proven for fluid-flow ON/OFF models in [Dumas & Simonian 00].

When $\nu > N\lambda_p$ or equivalently $\rho < 1 - b$, over-saturation of the server cannot happen even if all sources are active simultaneously. We call that region of the parameter space *Peak-Rate Allocation* (PRA). However, we focus on the more interesting blow-up regions herein, numbered by $i_0 = 1, ..., N$ where

$$i_0 = \left\lceil N \cdot \frac{1-\rho}{\rho} \cdot \frac{1-b}{b} \right\rceil\,, \tag{8.7}$$

which follows from (8.5) and (8.3). Furthermore, we introduce the parameter i_Δ:

$$i_\Delta = i_0 - N \cdot \frac{1-\rho}{\rho} \cdot \frac{1-b}{b}\,, \tag{8.8}$$

which has a range $0 \leq i_\Delta < 1$. The larger i_Δ the closer does the queueing model operate to the blow-up region $i_0 - 1$, where performance is worse.

If the product within the ceiling operator in Eq. (8.7) is larger than N, the model operates in Peak-Rate Allocation. Also, if that product is an integer ($i_\Delta = 0$), the model operates exactly at one of the blow-up points, see Sect. 10.1. Both cases will be excluded from the following discussion, i.e. the N-Burst/M/1 queue is assumed to operate in one of the blow-up regions $i_0 = 1, ..., N$ but not exactly at one of the N blow-up points itself. The blow-up points itself represent singularities that have to be treated separately.

As a consequence of the Power-Tailed queue-length distribution with exponent $\beta = i_0(\alpha - 1) + 1$, the steady-state *Buffer-Overflow Probability*

$$\mathrm{BOP}(B) = \mathbb{P}\left\{\text{at least } B \text{ buffer slots are occupied at cell arrival}\right\},$$

decays very slowly with increasing buffer-size, namely by a Power-Law with exponent $\beta - 1$:

$$\mathrm{BOP}(B) \sim c_{\mathrm{BOP}}\, B^{1-\beta}, \tag{8.9}$$

The tail-constant c_{BOP} can be approximated in the first blow-up region,

$$c_{\mathrm{BOP}} \approx \frac{b^\alpha}{1-\rho}\, \frac{(i_\Delta\, \overline{n_p})^{\alpha-1}}{\alpha - 1}\, c_{\mathrm{PT}}^{(1)}(\alpha), \quad \text{for} \quad i_0 = 1, \tag{8.10}$$

which follows from results of [Jelenkovic & Lazar 99] for fluid-flow models. Thereby, $c_{\mathrm{PT}}^{(1)}(\alpha)$ is the tail-constant of the reliability function of the burst-length distribution when normalized to have mean 1. See Appendix I.1 for numerical values.

In all other blow-up regions $i_0 > 1$, only the following approximation is obtained for the tail-constant c_{BOP}:

$$c_{\mathrm{BOP}} \approx \frac{b^\beta\, (1-b)^{i_0 - 1}}{1-\rho} \left[\frac{(i_\Delta\, \overline{n_p})^{\alpha-1}}{\alpha - 1}\, c_{\mathrm{PT}}^{(1)}(\alpha) \right]^{i_0}. \tag{8.11}$$

Numerical experiments and comparisons to computed tail-constants for the analytic N-Burst/M/1 model indicate that an additional factor $1/\sqrt{1 - i_\Delta}$ should be included when $i_0 = 2$, otherwise Eq. (8.11) underestimates the true value of c_{BOP}. Even with that correction factor, the approximation (8.11) is not as accurate for $i_0 > 1$ as (8.11) is in the first blow-up region. But the numerical values are still sufficient for practical applications.

The *Cell Loss Probability* (CLP) in a finite-buffer queue shows the same tail-behavior as the BOP, yet with smaller tail-constant:

$$\mathrm{CLP}(B) \sim c_{\mathrm{CLP}}\, B^{1-\beta}. \tag{8.12}$$

Results from the transient analysis in Chapter 6 allow to come up with a relationship between the tail-constants $c_{\mathrm{BOP}} > c_{\mathrm{CLP}}$:

$$(1-\rho)\, i_\Delta\, \frac{1 - i_\Delta}{i_0 - i_\Delta}\, c_{\mathrm{BOP}} \leq c_{\mathrm{CLP}} \leq (1-\rho)\, \frac{i_\Delta}{i_0}\, c_{\mathrm{BOP}}. \tag{8.13}$$

The upper and lower bound are identical for $i_0 = 1$, thus

$$c_{\mathrm{CLP}} = (1-\rho)\, i_\Delta\, c_{\mathrm{BOP}} \quad \text{for} \quad i_0 = 1. \tag{8.14}$$

Truncated Tails:

For truncated tails, the Power-Laws for the queue-length distribution and consequently also for the BOP or the CLP are truncated as well. It turns out that the Power-Law of exponent β in blow-up region i_0 only holds for queue-lengths up to

$$q_{i_0} = \frac{\text{MBS}}{\lambda_p\, i_0} \cdot (\Lambda_{i_0} - \nu) = \frac{b\, i_\Delta}{i_0}\,\text{MBS}\,. \tag{8.15}$$

However, q_{i_0} is the PT Range of the Power-Law with exponent β in the queue-length distribution. Other Power-Laws with larger exponents $\beta(i) = i\,(\alpha - 1) + 1$, $i = i_0 + 1, ..., N$ come into effect for longer queues. Thus, a geometric drop-off of the BOP is only observed for buffer-sizes larger than

$$q_N = \frac{\text{MBS}}{\lambda_p\, N} \cdot (\Lambda_N - \nu) = \text{MBS}\,\left(1 - \frac{1-b}{\rho}\right). \tag{8.16}$$

Note that $q_{i_0} < q_N < \text{MBS}$ when $i_0 < N$.

Finally, an asymptotic Power-Law behavior is derived for the mean Cell Delay (mCD), see Sect. 5.3.2:

$$\text{mCD(MBS)} \sim c_{\text{mCD}}\;(\text{MBS})^{2-\beta}\,, \quad \text{for} \quad \beta < 2\,. \tag{8.17}$$

Consequently, the mCD for full Power-Tails is infinite, if and only if

$$\beta \le 2 \quad \Leftrightarrow \quad \alpha \le 1 + \frac{1}{i_0}\,\cdot \tag{8.18}$$

The tail-constant in Eq. (8.17) can be bounded from above asymptotically for large MBS:

$$\text{mCD} < \frac{\rho\, c_{\text{BOP}}}{\lambda\,(2-\beta)}\,\left[\frac{b\, i_\Delta}{i_0}\right]^{2-\beta}\,\text{MBS}^{2-\beta} \quad \text{for} \quad \beta < 2 \tag{8.19}$$

Together with the approximation (8.11), an upper bound is obtained for $\beta < 2$:

$$c_{\text{mCD}} < \frac{1}{\lambda}\cdot\frac{\rho}{1-\rho}\cdot\frac{i_\Delta}{i_0^{2-\beta}}\cdot\overline{n_p}^{\beta-1}\,b^2\,(1-b)^{i_0-1}\cdot\frac{\left[c_{\text{PT}}^{(1)}(\alpha)\right]^{i_0}}{(2-\beta)\,(\alpha-1)^{i_0}}\,\cdot \tag{8.20}$$

Note however, that the additional factor $1/\sqrt{1-i_\Delta}$ that is required in Eq. (8.11) when $i_0 = 2$ should also be included in this bound for c_{mCD}.

8.3.2 Transient Behavior

The steady-state results are useful to understand the behavior of N-Burst/M/1 queues with LRD properties in the arrival process. Engineering rules for network dimensioning can be derived from those results. As it is shown in Sect. 6.1, steady-state behavior can be observed within rather short time intervals for traditional models without LRD properties. However, LRD models can behave differently: first, steady-state values might well be infinite, especially if parameters like mean delay are considered. Furthermore, large fluctuations of the QoS parameters can be observed even when averaging over long observation periods.

Cell losses or buffer overflows[1] on the other hand turn out to be highly correlated. As a consequence, the steady-state Loss or Overflow Probabilities might be misleading: It is possible that no losses at all are observed on many days, but on the few 'bad' days, a larger proportion of the cells is lost. To be able to describe such behavior, the analysis of *First Passage Times* (FPT) turns out to be very fruitful. The FPT is the random variable that expresses the time until the buffer becomes fully occupied for the first time when some initial condition is given[2]. Again, a Power-Law relationship can be observed for the mean First Passage Time (mFPT), see Chapter 6:

$$\mathrm{mFPT}(B) \sim c_{\mathrm{mFPT}}\, B^{\beta}\,. \tag{8.21}$$

Since (8.21) contains a Power-Law with exponent β, and β depends on the blow-up region, it follows that the blow-up effects also show up in the transient performance parameter mFPT (and also in the behavior of γ, see below).

The probability that no cell loss at all will be observed in a certain time-interval can be derived from the FPT distribution:

$$1-\gamma(t_0,B) := \mathbb{P}\left\{\text{“no overflow during time } t_0\text{ ”}\right\} = \mathbb{P}\left\{\mathrm{FPT}(B) > t_0\right\}\,. \tag{8.22}$$

Furthermore, the simulation experiments in Chapter 6 show that the distribution of $\mathrm{FPT}(B)$ can be approximated for large buffer-sizes B by an exponential distribution, provided that t_0 is not too small:

$$\gamma(t_0,B) = \mathbb{P}\left\{\mathrm{FPT}(B) \le t_0\right\} \approx 1-\exp\left(-\frac{t_0}{\mathrm{mFPT}(B)}\right) \quad \text{for} \quad t_0 \gg \frac{B}{\Lambda_{i_0}-\nu}\,. \tag{8.23}$$

Such an exponential approximation follows also from the asymptotic theory that is discussed in [Asmussen et al. 00a].

If the ratio $t_0/\mathrm{mFPT}(B)$ is very small, the exponential function in Eq. (8.23) can be approximated by the linear term of its Taylor series:

$$\gamma(t_0,B) \approx \frac{t_0}{\mathrm{mFPT}(B)} \quad \text{for} \quad \frac{B}{\Lambda_{i_0}-\nu} \ll t_0 \ll \mathrm{mFPT}(B)\,. \tag{8.24}$$

Using the asymptotic behavior (8.21) of the mFPT, another Power-Law relationship is obtained for the transient overflow probability:

$$\gamma(t_0,B) \sim \frac{t_0}{c_{\mathrm{mFPT}}}\, B^{-\beta}, \quad \text{for} \quad 0 \ll \frac{B}{\Lambda_{i_0}-\nu} \ll t_0 \ll \mathrm{mFPT}(B)\,. \tag{8.25}$$

In order to account for the potentially huge fluctuations in the fraction of overflowed cells in the individual observation intervals, the *conditional Buffer-Overflow Ratio*, $\mathrm{BOR}_c(t_0,B)$, is defined in Sect. 6.4 as the expected fraction of overflowed cells conditioned on at least one occurring overflow event in the observation interval. For PT distributed burst-lengths, it can be shown that the $\mathrm{BOR}_c(t_0,B)$ for large buffers grows asymptotically linearly as:

$$\frac{1}{\alpha-1}\,\frac{1}{(1-\rho)\,i_\Delta}\cdot\frac{B}{\lambda\, t_0} \le \mathrm{BOR}_c(t_0,B) \le \frac{1}{\beta-1}\,\frac{1}{(1-\rho)\,i_\Delta}\cdot\frac{i_0-i_\Delta}{1-i_\Delta}\cdot\frac{B}{\lambda\, t_0} \tag{8.26}$$

[1] Cell losses happen in a finite-buffer loss model, while we talk about buffer-overflows when an infinite secondary buffer exists, see also Sect. 2.6.

[2] In the experiments, an empty buffer at time $t=0$ is assumed. However, for large buffers B, an initial buffer occupation of only a few cells shows practically no impact.

The two bounds are identical in the worst blow-up region $i_0 = 1$.

The asymptotic Power-Law growth Eq. (8.21) of the mFPT(B) follows from the asymptotic behavior of $\mathrm{BOR}_c(t_0, B)$ and the knowledge of the Power-Law (8.9) for the steady-state BOP(B). Furthermore, the respective tail-constants have the following relationship:

$$\frac{1}{c_{\mathrm{BOP}}} \cdot \frac{1}{\lambda} \cdot \frac{1}{\alpha - 1} \cdot \frac{1}{(1-\rho)\, i_\Delta} \leq c_{\mathrm{mFPT}} \leq \frac{1}{c_{\mathrm{BOP}}} \cdot \frac{1}{\lambda} \cdot \frac{1}{\beta - 1} \cdot \frac{1}{(1-\rho)\, i_\Delta} \cdot \frac{i_0 - i_\Delta}{1 - i_\Delta} \cdot \tag{8.27}$$

Section 6.4.3 shows that the BOR_c has a minimum close to

$$B_{min} = \sqrt[\beta]{\frac{t_0}{c_{\mathrm{mFPT}}\, z_0(\beta)}}, \tag{8.28}$$

where $z_0(\beta)$ is the solution of the non-linear equation

$$\exp(-z_0) = \frac{1}{1 + \frac{\beta}{\beta - 1}\, z_0} \cdot \tag{8.29}$$

Numerical values for $z_0(\beta)$ can be found in Appendix I.3.

Analogous arguments lead to the asymptotic behavior of the finite-buffer *conditional Cell-Loss Ratio*, $\mathrm{CLR}_c(t_0, B)$:

$$\mathrm{CLR}_c(t_0, B) \sim \frac{1}{\beta - 1} \cdot \frac{B}{\lambda\, t_0} \cdot \tag{8.30}$$

As a corollary, it follows for the tail-constant of the steady-state CLP:

$$c_{\mathrm{CLP}} = \frac{1}{\beta - 1} \cdot \frac{1}{\lambda\, c_{\mathrm{mFPT}}} \cdot \tag{8.31}$$

Truncated Tails: For exponentially truncated Power-tails in the burst-length distribution, the $\mathrm{BOR}_c(t_0, B)$ converges with the following asymptotic bounds (for large MBS and even larger B):

$$\frac{b}{1-\rho} \cdot \frac{\mathrm{MBS}}{\lambda\, t_0} \leq \mathrm{BOR}_c^{(TPT)}(t_0, B) \leq \frac{b}{1-\rho} \cdot \frac{i_0 - i_\Delta}{1 - i_\Delta} \cdot \frac{1}{i_0} \cdot \frac{\mathrm{MBS}}{\lambda\, t_0} \cdot \tag{8.32}$$

Analogously, in the finite-buffer loss system:

$$\mathrm{CLR}_c^{(TPT)}(t_0, B) \longrightarrow \frac{b\, i_\Delta}{i_0} \cdot \frac{\mathrm{MBS}}{\lambda\, t_0} \cdot \tag{8.33}$$

Consequence of the asymptotic growth of BOR_c for LRD N-Burst traffic: With increasing buffer-space the overall fraction ($\approx$ BOP) of buffer-overflows is reduced while the overflow-events become more and more unevenly distributed over the individual observation periods (which could be for instance the typically 4-7 busy hours in the daily profile of the network utilization). As a consequence, the number of intervals with at least one overflow decreases with increasing buffer-space, but the expected fraction of overflowed cells during those overflow intervals even becomes larger. Thus, additional buffer-space seemingly increases the fluctuations of the number of buffer-overflow events in a number of observation intervals.

8.4 Engineering Rules for Network Planning

A collection of the most important relationships that are summarized in the previous two sections can be found in Section 7.1. In Section 7.3 of the same chapter, a number of scenarios are discussed and methods are presented how various network or traffic parameters can be computed when certain QoS requirements are proposed. The following table provides an overview on the parameters and the references to the corresponding part of Section 7.3:

Parameters	N	b	κ	$\overline{n_p}$	α	$c_{\mathrm{PT}}^{(1)}$	MBS	ν	B
Controllable...									
... at design phase	√							√	√
... via traffic policing		√					√		
Impact on i_0	√	√	√					√	
Impact on $\overline{q}[i_0]$	√	√	√				√	√	
Impact on $\overline{q}[N]$	√	√	√				√	√	
Choice discussed on	p. 120	p. 121					p. 121	p. 119	p. 120

8.5 Avoiding the Negative Impact of LRD

The negative performance behavior due to the LRD properties of ON/OFF traffic can be avoided or at least diminished by appropriate network design or by additional restrictions on the incoming traffic. In addition, the insights about how congestion occurs for such traffic can be used in the future network protocols. The following list summarizes all the practical remedies against poor QoS caused by ON/OFF traffic with LRD properties:

- Since an outstanding feature of the behavior of the performance parameters is the existence of blow-up points, at which performance deteriorates dramatically, one of the major goals must be to make the switch operate in a sufficiently well behaved blow-up region (where the effective PT exponent β is sufficiently large). In the extreme case of PRA ($\rho < 1 - b$ or equivalently, $\nu > N\lambda_p$), no Power-Law behavior at all exists.

  There are practically three possibilities:

  1. Run at *lower Utilization*: Add capacity ν.
  2. *Traffic shaping*: Reduce the burstiness b.
  3. Use *higher Degree of Multiplexing*: Increase N at same utilization ρ.

- *Traffic Policing:* Restrict the time-scales of the LRD properties by introducing a Maximum Burst Size for the incoming traffic. Cells of bursts that violate the imposed restriction will be discarded. That way the network is protected from 'bad' congestion, but additional responsibility may be imposed on the applications/user.

- *Throttle the cell-rate of large bursts:* In some scenarios, e.g. at `http` or `ftp` servers, the size of a requested transfer is known in advance. Since it is shown here, that the poor performance for LRD ON/OFF models is caused by the large bursts, the negative effects can be alleviated by reducing the cell-rate in those large bursts.

- Throttle the cell-rate during time-intervals with more than i_0 long-term active bursts: The problem here is the prediction of how long a cell stream is going to send at its peak-rate, when such information is not provided by the sources. Transmission protocols would have

to be extended by dynamic flow-control mechanisms, that use probabilistic arguments to decide whether the incoming traffic produces a long-lasting over-saturation period, in which case the individual cell-rates have to be throttled.

On the other hand, additional buffer-space does not lead to a substantial reduction of the probabilities BOP, CLP, and γ, in the blow-up regions. However, if there exists a finite MBS, enough buffer-space could be added, such that the Power-Tail Range q_{i_0} (or even q_N) of the queue-length distribution is exceeded. Only then, overflow probabilities are reduced strongly.

Note that for very bursty traffic (b close to 1), the blow-ups can occur at relatively low utilization ρ, see Eq. (8.7). Hence it can be dangerous to design a network only based on utilization values while neglecting to look at the resulting blow-up regions.

8.6 Conclusion

It is becoming increasingly clear that more accurate traffic models have to be developed in order to be able to obtain any realistic evaluation of QoS provision in telecommunication networks. As recent (and also not so recent) measurements show, multiplexed ON/OFF models can describe the bursty nature of network traffic quite well, and thus they reproduce one of the most important properties of such traffic. However, burstiness by itself is not sufficient, since additionally so-called Long-Range Dependent properties of network traffic became evident for traffic from several applications and also in different transmission protocols. One major cause of the LRD properties is the occasional occurrence of very large transmissions, called bursts herein. Classical ON/OFF models with exponentially distributed ON times do not show such properties. Power-Tail distributions have to be used instead for the ON times. Other measurements of e.g. file-sizes confirm the necessity of such distributions in realistic traffic models.

This thesis has investigated the impact of the LRD properties of multiplexed ON/OFF traffic on the performance behavior of network components with large buffers. The steady-state analysis as well as the transient analysis showed a very peculiar behavior: With increasing utilization of the switch, both steady-state and transient performance parameters do not worsen uniformly, but at several well-identified utilization values (called blow-up points), performance deteriorates rapidly. The awareness of those blow-up effects and the location of the blow-up points is one of the main results of the thesis, which has straightforward impact on network design procedures. For such traffic, the design of a network solely based on utilization values can be a fatal mistake with respect to QoS.

For large buffers, the asymptotic behavior of several performance parameters could be derived: In many scenarios, the performance is mainly determined by the tail of the queue-length distribution, whose exponent and tail-constant could be derived. That knowledge can be successfully applied to practical problems of network design, see Chapter 7.

Equally important to the derivation of asymptotic formulas for the performance parameters is the insight into how congestion for such traffic can occur even when the overall utilization of the switch is low: Large delay and overflow events are mainly due to multiple, simultaneous long bursts which cause a long-lasting over-saturation period for the switch. The probability of such an event is negligible for models without Power-Tail distributed burst-lengths. However, when the traffic has LRD properties due to Power-Tail distributions, that probability is small, but never negligible.

Furthermore, the choice of performance parameters is critical: cell delay and also overflow events are highly correlated for such models with LRD ON/OFF traffic. As a consequence, steady state

performance parameters can be misleading and transient analysis has to be applied, see Chapter 6.

It is a big advantage to be able to obtain numerical values for many of the performance parameters. However, in our opinion it is much more the method of analysis, the choice of the performance parameters, the fundamental principles for QoS oriented network design in Sect. 8.5, and finally the understanding of the behavior of such traffic in queueing models that must be added to the engineer's set of tools. In that context, the key findings with respect to LRD network traffic (as described by the N-Burst model) are:

- Existence of potential 'discontinuities':
  Changing one network parameter (e.g. the switch capacity, ν) does not necessarily affect performance uniformly. Instead, so-called blow-up points do exist, at which the performance changes radically.

- Cause of poor performance:
  Assuming large buffers, overflow events as well as large delays are mainly caused by (simultaneous) long ON periods of the sources. Such behavior is not rare when Power-Tail distributions are involved, which is known to be the case at least for file transmissions and `http` transfers.

- Potential for infinite mean delay:
  Although the actual distribution of per-cell delay is well-defined, its mean can be infinite (since it is Power-Tailed in most scenarios). Measurements of such mean delay would be very erratic even for a large number of samples. Properties of the distribution other than the mean should be considered instead, e.g. Quantiles.

- Knowledge of the behavior of performance parameters:
  A very important conclusion is that buffer-space might be of little help in reducing overflow probabilities. This is due to the Power-Law behavior Eq. (8.9) of $\mathrm{BOP}(B)$. Traditional models with exponential or other well-behaved ON times behave differently in that respect.

- Choice of Performance Parameters:
  Buffer Overflow or Cell Loss Probabilities cannot be considered to be independent for each cell in the traffic, but such events are highly correlated. The substitution of the single parameter, Buffer-Overflow Probability, by the parameter-pair $\langle \gamma(t_0, B), \mathrm{BOR}_c(t_0, B) \rangle$ is recommended.

- Fluctuations between observation intervals:
  Increasing buffer-space reduces the overall number of overflow-events. However, for N-Burst traffic with LRD properties, it even increases the fluctuations between the individual observation intervals. Many intervals might not show any overflows at all, but if there are any, then a huge fraction of the cells is lost. This discrepancy between individual observation periods increases for larger buffers when LRD properties of the traffic are present.

Carried a few steps further, the understanding and the validation of poor performance behavior of networks forms the basis for evaluating the effectiveness of more advanced mechanisms for QoS provisioning, see Sect. 10.7. The individual numerical results are useful, but only a more comprehensive understanding of the underlying performance problems can contribute to an objective debate about future QoS developments in Telecommunication Systems.

Chapter 9

Extension: Investigation of Elastic Traffic

In this chapter we develop a modified version of the N-Burst model for so-called *elastic traffic*, i.e. traffic whose transmission rates adjust in reaction to some condition outside of the source – here the congestion level of the network. Such a model can be used to investigate the behavior of the *Transmission Control Protocol* (TCP) which is used to transmit most of today's Internet traffic. TCP in turn uses services of the Internet Protocol (IP), whose fundamental data-units for transmission are called *packets*. Hence, we change our terminology for the models in this chapter: what used to be (ATM) *cells* are now called (IP) *packets*, and *switches* turn into *routers*. Again, the model is not exclusively applicable for TCP flow-control, but it can be used for any packet based transmission protocol whose flow-control mechanisms reacts to congestion by throttling the packet-rates at the sources.

9.1 Motivation

As we have seen in the previous chapters of this thesis, traffic modeling for data networks is already difficult even if the incoming traffic is assumed to be independent of the events that happen in the network, see also [Paxson & Floyd 95]. However, most of today's Internet traffic is transmitted using TCP (*Transmission Control Protocol*), whose built-in flow-control mechanism introduces a dependence between the network status (its congestion level) and the offered packet traffic, see [Arvidson & Karlsson 99]. As a consequence, large delays and high buffer-overflow probabilities in the network components themselves can be avoided, but at the cost of slowed down transmission rates at the sources.

This chapter introduces and discusses an analytic queueing model that captures the essence of the TCP flow-control mechanism while still remaining tractable. Preliminary performance results for this model and their practical implications are discussed. In particular, the use of truncated Power-Tail distributions for the ON periods leads to conclusions about the behavior of Long-Range Dependent traffic under the influence of the TCP flow-control mechanism.

Frequently, the behavior of TCP traffic is investigated in the scenario of TCP connections with infinite amount of data to transmit. For instance, [Padhye et al. 98] derive an estimate for the throughput for such an ever-lasting TCP connection, if its packets are subject to some loss probability p along their transmission path.

[Heyman et al. 97] introduce and discuss a model that includes some kind of ON/OFF behavior of the users. Their model works on flow-level, i.e. individual packets are not considered. One of the consequences of that model is that the throughput is insensitive to the actual distribution of the size of the connections; only its expected value matters. The model in this chapter is an extension of the N-Burst Independent Source model that is based on the same idea as the model

of [Heyman et al. 97], but it is packet-based, which allows to include a queueing model and a more realistic feedback mechanism. One of the implications is that the insensitivity towards the distribution of the connection size does not hold any more, if some buffer-space is available in the bottleneck router, see Sect. 9.4.2.

9.2 A Model for Dynamic TCP Traffic

In this chapter, the interest is particularly in dynamic TCP connections, i.e. connections start and end as described by some stochastic process. When neglecting the flow control mechanism, those dynamics can be captured by the model that we were focusing on so far, the N-Burst. The N-Burst model is used as a base model and it is shown in the remaining part of this section, how it can be modified to account for the flow-control mechanism of TCP.

9.2.1 Modification 1: Shared Bottleneck Bandwidth

The assumption of ON/OFF traffic with constant Poisson rate λ_p during the ON periods is reasonable for many real-time applications and for protocols without flow-control mechanisms. However, TCP works differently. After a certain number of packets (the so-called *congestion window*) is sent out, the sender awaits acknowledgment packets from the receiver. The size of the congestion window is dynamically adjusted when congestion is detected. We omit the details here, since they will not show up in the model in any case. Also, they depend on the actual TCP implementation. The interested reader be referred to [Stevens 94]. The important feature of the flow-control mechanism is that through the adjustment of the size of the congestion window, the effective sending rate of packets can be throttled from its maximum λ_p (which is determined by the speed of the access line) to a sufficiently low rate such that congestion is (hopefully) avoided.

The first modification of the N-Burst model is made along those lines of throttling the packet-rate of each individual source, if the bottleneck router in the transmission path cannot handle all the active connections any more. If we assume that the output bandwidth of the bottleneck router corresponds to a packet-rate ν, and i sources (connections) are active, the maximal sending rate λ_p is only used if $i\lambda_p < \nu$, i.e. when no overload situation at the bottleneck is created. If $i\lambda_p > \nu$, all sources equally throttle down their packet-rate by a factor β_i. One approach is to throttle the sources by a factor $\nu/(i\lambda_p)$, so that the i active sources share the bandwidth of the router. See [Schwefel 00b] for results in that case. When the active sources share exactly the output bandwidth, congestion would be avoided while still keeping the router fully utilized. In practice however, the flow-control mechanism of TCP cannot be expected to achieve that optimal high utilization. Instead, the regulation mechanism throttles the active sources slightly stronger, so that they collectively use slightly less than the bandwidth ν. [Heyman et al. 97] derive a formula that allows to quantify such an 'over-reaction' of the TCP flow-control. In general, such a formula depends on several TCP parameters. Here, we assume for simplicity, that while being throttled, the sources only use 80% of the bottleneck router's bandwidth ν, i.e. the individual packet-rate λ_p is throttled by the factor

$$\beta_i := \frac{0.8\,\nu}{i\lambda_p},$$

when i sources are active and $i\lambda_p > \nu$. We call the resulting aggregated ON/OFF model the SHARED model.

Since the actual number of packets in the connection must not be changed, a throttling of the packet-rate also requires a change in the duration of the connection: In the unthrottled N-Burst model with exponential ON time distribution, the rate of a transition to the $(i-1)$ active state is $i/\overline{x_p}$. This rate is changed in the SHARED model to $\beta_i\, i/\overline{x_p}$, i.e. the state holding time is extended. See Appendix E.1.3 for the general case of Matrix-Exponential ON times. This modification is also mentioned in Sect. 2.5.1.

The rates for transitions that correspond to starting connections remain unchanged however, since idle sources are not affected by the throttling. The modification of the MMPP to obtain the SHARED model from the N-Burst model are far from trivial. It has to be investigated whether the distribution of the number of packets per connection is not changed by that modification. While an individual source is active, the throttling that is caused by other active sources only leads to a different scaling of the clock for this and all the other active sources. When additional sources become active, or active sources finish, that scaling of the clock for the remaining active sources is changed. However, the phase transitions within the individual ON time remain in the same order, and also the number of packets that are generated while in a certain phase is not affected. Hence, the distribution of the number of packets in an ON period is not changed. Numerical computations of the distribution of the number of packets in the SHARED model confirm that there is no impact of the throttling.

Note that the extension of the ON periods is not made up for by a reduction of the length of the subsequent OFF period of the same source. Therefore, the throttling not only decreases the observed packet rates during the connection, but also the long-term average packet rate of the individual source.

9.2.2 Modification 2: React to Existing Congestion

In contrast to the assumptions in the SHARED model, the individual real TCP source does not have the knowledge about any other, newly starting TCP connections, but it only reacts to existing congestion situations. In that sense, the control mechanism of the SHARED model is too good, since it adjusts the sending rates of the sources instantaneously when new connections become active.

A second modification of the traffic model accounts for that behavior: As long as no congestion is present, each source generates packets during ON-periods with Poisson rate λ_p as in the basic N-Burst model. Whenever the buffer-occupancy at the bottleneck router reaches B_1 or more packets, TCP is assumed to recognize the congestion situation, and the arrival process switches to the SHARED model. As an approximation, the buffer itself is assumed to be infinite, so that we do not have to worry about retransmissions of lost packets (they are stored in the 'backup' buffer-space beyond level B_1 instead). This model will be called TCP_{B_1}. It is easy to see that the N-Burst/M/1 queue is the limit $B_1 \to \infty$, where the throttling of packet-rates is never performed. At the other end, SHARED is the other limiting model for $B_1 = 0$.

The computational methods to solve for the steady-state solution of the TCP_{B_1} queueing model are described in Appendix F.2. They follow in spirit the model of [Krieger & Naumov 99].

The remainder of this chapter discusses the impact of the throttling on the packet-stream of the individual source as well as on the performance behavior of the TCP_{B_1} model.

9.3 Impact on Packet Flows

The numerical examples in this chapter look at the scenario that N fast (10Mb/s) LANs are connected via the bottleneck router to a slow (1Mb/s) access line. Each LAN is assumed to be used by only one ON/OFF source. If we assume an average packet-size of 1kB and an average connection size of 50kB, the parameters for the TCP_{B_1} model follow as: $\overline{n_p} = 50$ packets, $\lambda_p = 1250$ packets/s, $\nu = 125$ packets/s, $\overline{x_p} = 40$ ms in the unthrottled N-Burst. Furthermore, we assume that the exponentially distributed OFF periods have mean $Z = 5$ s. Consequently, the average packet rate in the unthrottled N-Burst comes out to be $\kappa = 9.92$ packets/s and somewhat lower when $B_1 < \infty$, see next section. Whenever the truncated Power-Tail distributions of Section 3.4 are used for the ON-periods, the tail-exponent $\alpha = 1.4$ is used. Note that since $\lambda_p > \nu$, even a single TCP connection by itself gets throttled as soon as the buffer-occupancy reaches B_1 packets.

Before we investigate the performance of the TCP_{B_1} model, we look at some properties of the traffic of an individual source that is subject to the impact of the throttling in the SHARED model.

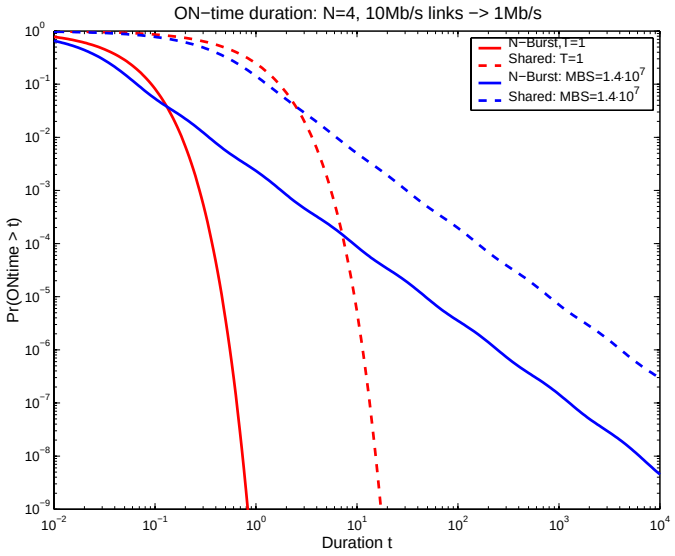

Figure 9.1: **The complementary distribution function of the duration of the ON periods in the N-Burst and SHARED model:** Power-Tail behavior (appearing as straight line) is observed for both models when MBS $= 1.4 \cdot 10^7$.

Connection Duration: As already mentioned in the previous section, the distribution of the number of packets per connection is identical in the N-Burst and in the SHARED model. However, since the packet-rate is at times (in our scenario always) reduced in the SHARED model, the duration of the connection is prolonged. The distribution of the connection duration in the SHARED model is a complicated Matrix-Exponential distribution even when the original ON-time distribution in the N-Burst model is exponential. Figure 9.1 plots the numerical values of the tail-probabilities of the connection duration in the N-Burst model and the SHARED model. The throttling does not affect the shape of the tails, exponential remains exponential, and so do Power-Tails. For the given parameters and $N = 4$ ON/OFF sources, the throttling reduces the utilization of the bottleneck router from 31.75% in the N-Burst model to 28.27% in the SHARED model .

Autocorrelation: Since Long-Range Dependent (LRD) properties of network traffic are dis-

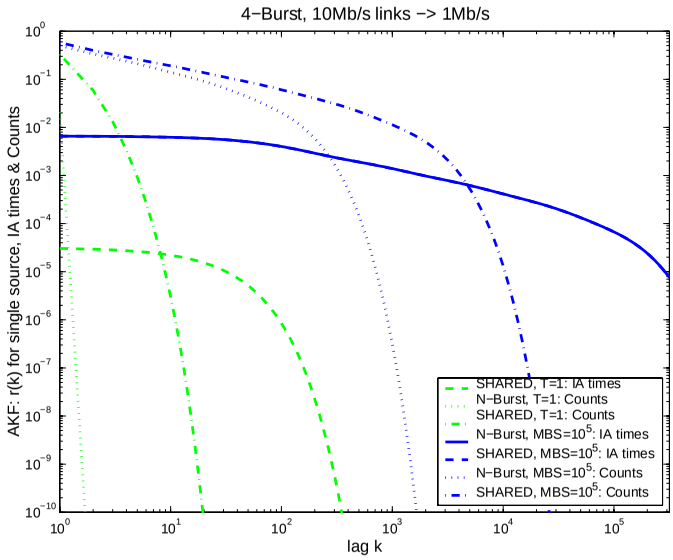

Figure 9.2: **The autocorrelation function of the inter-packet times and the counting process:** The correlation of the inter-packet times of the SHARED model is only marginally larger than in the N-Burst model when MBS $= 10^5$: the two curves are indistinguishable. For $T = 1$, an exponential ON/OFF source in the N-Burst model shows no correlation in its inter-packet times, that curve does not appear on the logarithmic scale.

cussed in detail in this thesis, it is interesting to investigate whether such slowly decaying correlation structure could be a result of the throttling of the individual connections. Therefore, we look at the coefficient of autocorrelation for both, inter-packet times and number of packets in an interval of size $\Delta = 1$ s, for the traffic of a single ON/OFF source, which is subject to various degrees of throttling due to other TCP connections in the SHARED model. Figure 9.2 shows the numerical values of the coefficient of correlation on log-log scale: A single ON/OFF source with exponential ($T = 1$) ON and OFF periods generates inter-packets times that are in fact a renewal process. Thus their correlation is zero in the N-Burst model with $T = 1$. The throttling in the SHARED model introduces some correlation (lower dashed line) yet the values of $r(k)$ are so small that they would not be noticeable in real measurements (due to the unavoidable noise). However, this is different in the counting process (dashed-dotted curve for N-Burst and dotted curve for SHARED). There, the value of $r(k)$ is raised substantially for small lags k, but still no LRD properties can be observed.

If truncated Power-Tail distributions are used for the ON periods, the correlation of the inter-packet times shows the LRD properties already for the N-Burst model. The throttling in the SHARED model increases those values, but only marginally so that the curves cannot be distinguished in Figure 9.2. In the counting process, the impact of throttling can be seen much more clearly.

9.4 Performance Results

When talking about performance of the TCP$_{B_1}$ model, it is important to keep in mind that now the delay is not only caused at the bottleneck router (Sect. 9.4.1), but packets can be delayed already at the source (Sect. 9.4.2).

9.4.1 Buffer-Occupancy

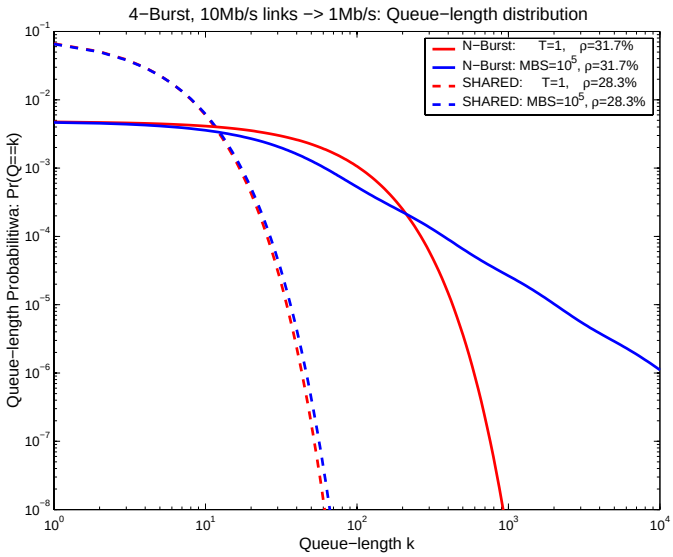

Figure 9.3: **The Queue-Length Distribution of the N-Burst and SHARED Model**: the throttling in the SHARED model (dashed lines) reduces the probabilities of very long queues even in the LRD case.

It is shown in Sect. 5.2 that the queue-length probabilities for an unthrottled N-Burst model with Power-Tailed ON periods also show a Power-Tail, whose exponent depends on the original tail-exponent α and the number i_0 of sources that are sufficient to over-saturate the router when they are simultaneously in a long ON period. In our scenario, $\lambda_p > \nu$, a single source already creates an over-saturation period for the router, and the queue-length distribution of the N-Burst/M/1 model decays with a Power-Law with exponent α, see Fig. 9.3 and Sect. 5.2. If we now turn to the SHARED model, then the arrival rate never exceeds the service-rate due to the throttling. What used to be an over-saturation period (temporarily $\rho > 1$) in the N-Burst model, is now only a stretched period of relatively high utilization, temporarily $\rho_t = 80\%$ in our case. Within this time-period of length Δ, the behavior of the SHARED/M/1 queue corresponds to a transient M/M/1 queue with same service-rate ν and utilization ρ_t. From the results of the First Passage Time analysis for M/M/1 queues in Appendix A.2 it follows that the queue grows approximately to length[1]

$$q_\Delta = \frac{\log \frac{\Delta}{\nu\,(1-\rho)^2}}{\log \rho}$$

in this time-interval ($\mathrm{mFPT}(q_\Delta) \approx \Delta$). In the original N-Burst model, the queue grows approximately to length $(\Lambda_{i_0} - \nu)\Delta$ during an over-saturation period of duration Δ with average arrival rate $\Lambda_{i_0} > \nu$. Therefore, the throttling leads to much slower queue-growth, as demonstrated by the dashed lines in Fig. 9.3. This corresponds to the observations of [Arvidson & Karlsson 99] that overflow-probabilities in simulation experiments of TCP traffic are not nearly as dramatic for large buffers, as conventional ON/OFF models without flow-control predict.

Let us now turn to the more realistic TCP_{B_1} model, which takes into account the buffer-space at

[1] If the active sources share exactly the bandwidth ν, $\beta_i = \nu/(i\lambda_p)$, a transient M/M/1 queue with $\rho_t = 1$ describes the model behavior during the over-saturation periods. In that case, the queue-length grows to values of approximately $\sqrt{2\nu\Delta}$ in time Δ, yet still much slower than in the original over-saturation period. See [Schwefel 00b].

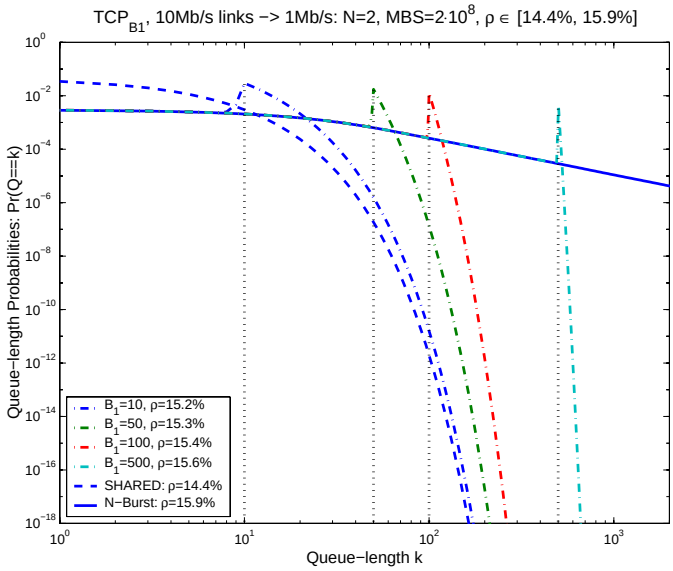

Figure 9.4: **The Queue-Length Distribution of the TCP$_{B_1}$ Model for $N = 2$ Sources and Long Power-Tails:** Below the peak at $k = B_1$, the distributions is close to the N-Burst model. After the peak, the queue-length of the TCP$_{B_1}$ model approaches the distribution of the SHARED model.

the router. Figure 9.4 illustrates, that the queue-length distributions of the two limiting models, N-Burst and SHARED, provide an excellent description of the behavior of the more complicated TCP$_{B_1}$ model: For buffer-occupancies below the threshold B_1, the unthrottled N-Burst arrival process is predominant, and the queue-length probabilities in the TCP$_{B_1}$ model follow the ones of the N-Burst model. The analog observation holds for buffer-occupancies above B_1, when the SHARED model takes over. The probability-mass that is taken away from the N-Burst curve for large buffers is now mainly concentrated in a peak around the buffer threshold B_1. Thus, Fig. 9.4 illustrates nicely the effectiveness of the flow-control mechanism.

9.4.2 Average Packet-Rate

The results in the previous section showed, that the flow-control mechanism prevents the built-up of huge queues. After all this is not too surprising, since that is its goal in the first place. However, there is of course a price to be payed for the improvement of the performance at the bottleneck router: The packet-rates during the connections are reduced, so packets are held back at the source instead of at the router. This additional delay must not be neglected in a fair discussion of the effectiveness of TCP flow control.

The steady-state solution for the TCP$_{B_1}$ queueing model in Appendix F.2 allows to determine an average packet-rate $\overline{\lambda_p}(B_1)$ that each connection achieves. Obviously, in the limit $B_1 \to \infty$, the N-Burst model is obtained, so $\overline{\lambda_p}(B_1)$ converges monotonically from below to λ_p. The numerical computations in Fig. 9.5 show clearly that at the other end $B_1 = 0$, the average packet-rate in the connection is independent of the actual type of the burst-length distribution. This is in principle the insensitivity result that was already pointed out in [Heyman et al. 97] for their flow-level model.

However, Fig. 9.5 also shows that the insensitivity is abolished, as soon as buffer-space $B_1 > 0$ is available. Exponential connection sizes are better off in those scenarios, because most of the

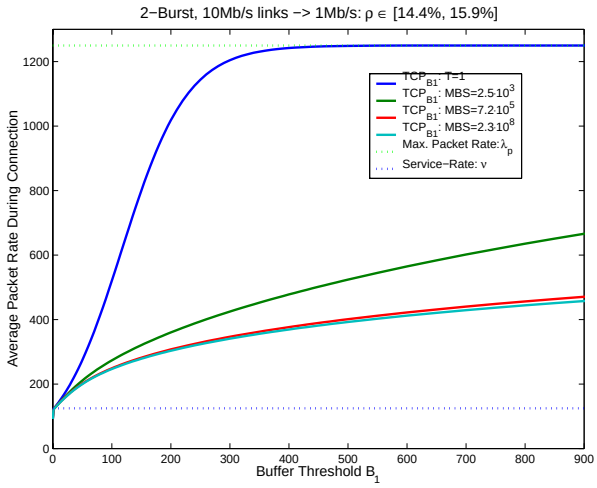

Figure 9.5: **The Average Packet-rate during Connections of the TCP$_{B_1}$ Model for $N = 2$ Sources:** An increased buffer-size B_1 also leads to less throttling. However, the gain is less pronounced for Power-Tailed connection sizes (with large T). The uppermost solid curve belongs to the exponential TCP$_{B_1}$ model with $T = 1$. The other solid curves are labeled according to the order in that is used in the legend.

time, the buffer can absorb a large part of the connection without throttling the source. In case of Power-tails on the other hand, the occasional huge connections increase the probability of throttling taking place.

Finally, we investigate the impact of adding additional sources (LANs) at the access of the router, see Fig. 9.6. With increasing number of sources N, the overall utilization of the TCP$_{B_1}$ model increases, shown by the ratio of the throughput (dashed lines) to the constant service rate ν (dotted line). Note that due to the throttling, the TCP$_{B_1}$ queueing model always remains stable ($\rho < 1$), yet the overall throughput shows a horizontal asymptote for $N \geq 15$. The average-packet rate $\overline{\lambda_p}(B_1)$ during the connections benefits most from larger buffer-space, when the utilization is low (small N). For very high utilization ($N \geq 13$), the buffer is almost always filled up to level B_1 regardless of the actual value of B_1. Therefore, the connections are almost always throttled. The curves in Fig. 9.6 show the scenario of exponential connection sizes. If we use Power-Tails instead, the curve for the SHARED model is not affected at all (insensitivity!), while the benefit of $\overline{\lambda_p}(B_1)$ for small utilization values becomes less pronounced.

Note that the overall throughput for the SHARED model in Fig. 9.6 only approaches 80% of the service-rate ν for a large number of sources $N > 15$, since any arrival rate larger than that is throttled immediately. As soon as some small amount of buffer-space $B_1 > 0$ is present, the utilization can approach 100%, since the throttling is restricted to time-periods where the queue-length is equal to or exceeds B_1.

In conclusion: For bottlenecks with high overall utilization, buffer-space and also the actual distribution of the connection sizes has little impact. But in lowly utilized routers (e.g. $\rho < 0.5$), buffer-space can lead to substantial improvement of the packet-rates during the connections. However, then there is a strong impact of the actual distribution of the connection sizes.

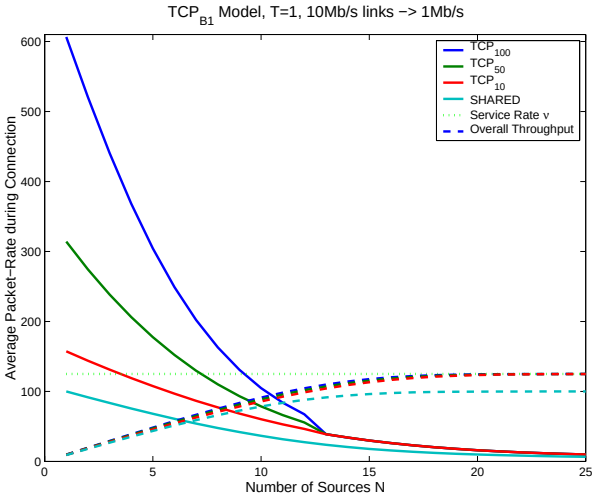

Figure 9.6: **The decay of the average packet-rate during connections due to throttling for an increasing number of sources N:** For $N > 12$ sources, the router is almost saturated, so that buffer-size B_1 is almost always exceeded and throttling is hardly ever turned off. Again, the labeling of the solid curves corresponds to the order that is used in the legend: the uppermost solid curve belongs to the TCP_{B_1} model with largest buffer $B_1 = 100$, while the lowest solid curve describes the results of the SHARED model where $B_1 = 0$.

9.5 Conclusions

When modeling performance for TCP traffic it is important to capture the feedback between the network and the offered input traffic. However, the dynamics of newly starting or ending TCP connections should not be neglected, i.e. the fundamental ON/OFF source behavior is important.

The model that is introduced in this chapter presents an extension of the N-Burst/M/1 queue that is investigated in Chapters 2 – 8 of this thesis. The extension captures the essential features of TCP flow control. These are:

- The sending rate of packets is adjusted in order to avoid congestion. Optimally, all active connections share the output bandwidth of a bottleneck router. This is implemented in the SHARED model.
- Rather than keeping track of the bandwidth requirements of the active connections, TCP only reacts to an existing congestion situation. Therefore, its flow-control mechanism does not come into play, before actual congestion has occurred. Model TCP_{B_1} takes this feature into account by conditioning the throttling of the individual packet rates not only on the number of active sources but also on buffer-occupancy of the bottleneck switch.

The developed TCP_{B_1} model still remains tractable, so it provides exact numerical results for several performance parameters. Here, we investigate the buffer-occupancy at the bottleneck router and the average packet-rate that a single connection can achieve. The results from the analysis of the model are the following:

- The autocorrelation of inter-packet times is increased by the throttling, yet no LRD like behavior could be observed that is caused purely by the throttling mechanism. Further-

more, the dominating part of the autocorrelation function of the inter-packet times is contributed by non-exponential ON period distributions rather than by the shaping due to the flow control. The autocorrelation function of the counting process is more strongly affected.

- The throttling prevents the build-up of large queues at the router. However, the prize is that the packets are delayed already at the sources. The numerical results for the average packet rate during a TCP connection illustrate that the delay shifts from the router to the source.

- In lowly utilized router, additional buffer-space can dramatically improve the packet-rates during connections, and thus it reduces the time that is necessary to send out a certain number of packets. The improvement is most pronounced if the distribution of the connection size is well-behaved (exponential). If Power-Tail distributions are involved, the improvements are still observed but only to a smaller extent.

The results in this chapter provide quite a few insights about the behavior of such regulated traffic in networks. However, it also raises a number of questions which have to be considered in future work

- Theoretical Issues: It is mentioned in Sect. 9.2.1 that the distribution of the number of packets in a connection is not affected by the way, the throttling is implemented. This is intuitively reasonable and was validated numerically for several different parameter sets. However, a rigorous mathematical proof is still missing.

- Queueing Delay: The investigation of average packet rates during connections can be used to derive conclusion about the time, it takes to send out the packets of a connection. However, when $B_1 > 0$, the last packet of the connection can still be stuck in the buffer. That particular queueing delay of the last packet of a connection is of interest for more detailed performance evaluation.

- Model modification: The TCP_{B_1} is still only an approximation of the true TCP behavior in the following features:

  - Delayed Reaction to Congestion: TCP needs at least a round-trip time or a timeout interval to recognize a packet loss and react to it. In the TCP_{B_1} model, a buffer-occupancy of at least B_1 results in instantaneous action (replacement of the N-Burst arrival process by the SHARED process).
  - Slow-Start: Current TCP implementations start off with a congestion window of size 1 which is only gradually increased. As a consequence, short connections only achieve a packet-rate which is lower than λ_p even when no congestion arises during the connection. That feature is currently not implemented in the model.

  The question is of course, in which scenarios do the current model simplifications matter? A feeling for what the answer probably looks like can be derived by simulation. However, the experiments have to be evaluated carefully, since recent investigations in [Veres & Boda 00] have shown that generalizations from TCP simulation results can be questionable.

  Simulation experiments in [Heyman et al. 97] showed that their model (which is a flow-based model of the SHARED model) provided a good match to the TCP behavior in

their simulated scenario. Therefore, there is hope that the TCP_{B_1} model is able to describe the TCP behavior at buffered bottleneck routers, even though the actual simulation experiments need to be conducted yet.

- Applicability of steady-state analysis: It is known that steady-state analysis for Long-Range Dependent traffic can be misleading, since the performance parameters rarely reflect the large fluctuations that will be observed in practice for such traffic. Transient analysis as performed for the N-Burst model in Chapter 6 provides a better description of such behavior. In this chapter, we only present steady-state analysis. It has to be investigated, whether transient analysis provides additional insight into the flow-control mechanisms as well.

Chapter 10

Open Problems and Future Directions

In this thesis, a very flexible class of realistic traffic models, called N-Burst, is developed. By an extensive analysis of analytic queueing models with N-Burst traffic with LRD properties, a detailed understanding is obtained how congestion can occur for such traffic even at overall lowly utilized bottleneck switches. Due to the flexibility of the model, not every possible aspect of the performance behavior could be covered in detail for every possible model variation. Instead, the thesis concentrates on the problems that seem to be essential for the network planning task.

However, quite a few questions of interest had to be left out or could only be touched briefly. Some of them are straightforward extensions of the investigations of Chapters 5, 6, and 9. Other problems are directed at future applications and require more elaboration and probably also new approaches. Own sections are dedicated to those more complex problems below. However, we start with a list of immediate continuations of this research:

- Inhomogeneous Bursts:
  The performance analysis in this thesis is mostly focused on the N-Burst Independent Source model, which assumes all the sources to be identical. This represents the base case, but in reality, scenarios with traffic mixes from different applications can occur. Such scenarios can be easily modeled by modifications of the N-Burst model, see Section 2.5.2. Preliminary results of the performance analysis are also discussed in Sect. 5.8: the existence of the blow-up points in the case of bursts with different cell-rates $\lambda_p^{(j)}$ is discussed and their location $\Lambda_i^{(j)}$ is derived. However, an in-depth discussion of the performance behavior (e.g. Power-Laws and Tail-constants for BOP, CLP and γ) is still missing for those scenarios.

- Tail-Constants:
  In order to use the asymptotic Power-Laws of BOP, CLP and mFPT for quantitative analysis, the values of the tail-constants c_{BOP}, c_{CLP}, and c_{mFPT} are necessary. In Chapters 5 and 6, approximations and bounds for these tail-constants as well as bounds for the behavior of BOR_c are discussed. Although the comparison with the exact numerical results for N-Burst/M/1 queues in Sect. 7.2 shows that those approximations and bounds are useful, there is still room for improvement. In particular in the region close to the blow-up points, tighter bounds for c_{CLP} and c_{mFPT} would be useful. Alternatively, the discussion of the model behavior at the blow-up points might be useful, see Sect. 10.1.

- Waiting Bursts and Different Burst-Start Processes:
  The performance analysis in Chapters 5 and 6 is almost exclusively based on the N-Burst Independent Source model. However, the N-Burst model class is more general, and allows for additional, physically meaningful burst-start processes, see Sects. 2.4 and 2.5. The qualitative behavior of performance models with such more general N-Burst processes with LRD properties can be expected to be comparable, since the major cause of congestion – the over-saturation periods of Power-tailed duration – still occur. However, the location of the blow-up points changes and also the tail-constants for the behavior of BOP, CLP,

and mFPT will be different. See Sect. 5.1 for remarks on the location of the blow-up points for such models.

- Parameter Extrapolation for Higher Load:
  In Chapter 4, 1-Burst models are used successfully to describe the measured traffic on a single incoming link of a switch. That traffic is in fact an aggregation of several hundreds of individual end-to-end connections. The 1-Burst parameters κ, λ_p, $\overline{n_p}$, Z, and λ_0 can be estimated from actual measured inter-cell times by the method that is developed in Sect. 4.3. For the purpose of network planning however, these parameters have to be extrapolated to some extent to the future, since traffic volumes tend to grow rapidly in today's networks (by a yearly factor between 2 and 3 in the B-WiN). The immediate question arises, in what way the traffic parameters have to be modified to account for the expected increase in traffic volume. Reasonable scenarios are a decrease of Z, or an increase of $\overline{n_p}$ or a combination of both.

  If the N-Burst model is used at source-level (as in Chapter 9) the seemingly obvious way of extrapolation is an increase in the number of sources N. However, adequate techniques of forecasting for this model class are still an open research field.

- Performance Models for TCP/IP:
  Much of today's Internet traffic is transmitted via TCP/IP, in which case the offered traffic from the source depends on the congestion level along the transmission path via TCP's sliding window flow-control mechanism. Chapter 9 introduces a model that captures the essence of the TCP flow-control mechanism – the back-off behavior of the sources. However, that model abstracts from many implementation details of TCP: for instance, the traffic that results from the acknowledgment packets is not modeled explicitly. Also, the connection set-up phase with its three-way handshake as well as its initial slow-start is not considered. Much more research can be done in this area. See also Sect. 9.5.

  Furthermore, the packet-size in IP is variable, and the sizes of subsequent packets in the same connection are very likely highly correlated. Also, the packet headers introduce overhead which needs to be taken into account when discussing packet sizes. Preliminary results for such a discussion can be found in [Lipsky & Schwefel 00], but more research is necessary.

10.1 Behavior at Blow-up Points

The performance analysis in this thesis concentrates on N-Burst/M/1 queues that operate in one of the blow-up regions $i_0 = 1, ..., N$ but not at the blow-up points itself (where $i_\Delta = 0$). In order to understand the behavior in the transition region close to the blow-up points better, it is necessary to discuss the behavior at those singular points.

Within blow-up region i_0, over-saturation periods are created by a minimum of i_0 long-term active bursts. During those over-saturation periods, the average arrival rate $\Lambda_{i_0} = i_0\lambda_p + (N - i_0)\kappa$ temporarily exceeds the service-rate of the switch, $\Lambda_{i_0} > \nu$, and the queue grows on average with rate $\Lambda_{i_0} - \nu$. The transient analysis of such models during the over-saturation period in Sect. 6.5 provides a more exact description and it confirms the growth with average rate $\Lambda_{i_0} - \nu$. Note that since the argumentation uses an average excess rate, the specific details of the arrival process during the ON periods are not important: Whether those arrivals are described by a Poisson process as in our N-Burst model, or by a continuous fluid flow process as in [Dumas & Simonian 00], does not have an impact on the average growth rate.

Such an 'insensitivity' to specific model details cannot be expected at the blow-up points themselves when $i_\Delta = 0$. In that case, the average arrival rate during a period with i_0 long-term active bursts is exactly equal to the service rate, $\Lambda_{i_0} = \nu$. Thus, transient models with $\rho = 1$ have to be investigated to describe the behavior during 'over-saturation' periods with i_0 long-term active bursts for N-Burst/M/1 models that operate at blow-up point i_0 ($i_\Delta = 0$).

At the transition from the last blow-up region $i_0 = N$ to PRA, the behavior of an N-Burst/M/1 queue during an over-saturation period of length Δ is identical to an M/M/1 queue with $\rho = 1$. From results of the transient analysis of M/M/1 queues (see Appendix A.2), it is known that the queue-length of an M/M/1 queue with utilization $\rho = 1$ grows to values of approximately $\sqrt{2\nu\Delta}$ in time Δ. Within the blow-up region $i_0 = N$, the queue grows much faster, as $(\Lambda_N - \nu)\Delta$, during an over-saturation period of length Δ. Consequently, although long queue-lengths can still be observed, the speed of growth is much slower at the blow-up points itself, and different exponents in the Power-Law behavior of the queue-length distribution can be expected. When changing the assumptions for the arrival process during ON periods and for the service process to a fluid-flow model, no accumulation of workload in the fluid queue occurs at all when the arrival and service rates are equal. Thus, at the blow-up points themselves, those specific model details show a major impact. The latter statement reduces the practical value of the models in those scenarios, since a lot of effort would have to be spent in order to use the 'correct' process for the inter-cell times during bursts.

10.2 Highly Multiplexed Traffic

The blow-up effects are particularly pronounced when a relatively small number of sources can lead to over-saturation periods, i.e. i_0 is relatively small (say $i_0 < 10$). Large values of i_0 can result from scenarios with a large number of sources that are not too bursty (b not close to 1), since

$$i_0 = \left\lceil N \frac{1-b}{b} \frac{1-\rho}{\rho} \right\rceil .$$

An example for such a scenario would be the entrance point of a backbone at which hundreds of slow (1Mbit/sec) links merge into a thick trunk (155 Mbit/sec). This is different from the scenario in Chapter 4, where the access networks were almost as fast as the backbone links.

If i_0 is large, the effective Power-Tail exponents, $\beta = i_0(\alpha - 1) + 1$, of the over-saturation periods with i_0 long-term active sources becomes rather large. As a consequence, the probability that the same i_0 sources cause a long over-saturation period becomes so small that it is not the dominant cause of congestion any more. This conjecture is supported by the performance analysis of limiting Fractional Gaussian Noise (FGN) models, see [Norros 95] and [Krishnan 96]. As pointed out in Sect. 2.7.5, the N-Burst model converges to such FGN models for large N and for appropriate scaling. Since essentially always $i_0 = \infty$ in the limiting FGN model , no blow-up effects of the performance parameters are observed.

In addition to the obvious theoretical relevance of understanding that limit process for large N, there is also a practical relevance: It is necessary to know how large i_0 (and thus $N > i_0$) has to be, such that FGN models provide an adequate traffic description. Also, the parameters of the individual sources have to be linked to the parameters of the FGN model. However, note that for very bursty sources (b close to 1), a high multiplex degree N does not automatically imply a high i_0: it is i_0 and not N that determines whether the results in this thesis are applicable.

10.3 Calibration Methods & Measurement Based Reconfiguration

The analysis of measurement data is not the focus of this thesis. Nevertheless, the parameter estimation from real data is necessary for the practical application of the model and also to judge the practical relevance of the obtained performance results. Hence, Chapter 4 develops a viable approach for the 1-Burst calibration. However, the attitude in that chapter is more an engineering approach than a mathematically rigorous derivation. The calibration method satisfies the purpose for this thesis: it works. A rigorous statistical analysis still needs to be performed. Furthermore, a more feasible approach to calibrate N-Burst models with $N > 1$ than the very laborious and heuristic method presented in [Schwefel et al. 97] is necessary. Therefore, calibration methods for N-Burst models are still waiting to be developed further.

Furthermore, the calibration method in Chapter 4 is based on the inter-arrival times of cells. Especially, if continuous online measurements of traffic at various locations in the network are desired, it is in practice not very feasible to collect all individual inter-cell times since this creates an immense amount of data – which has to be collected somewhere and therefore its transmission adds substantial load on the network as well. Therefore, calibration methods are necessary that rely on the number of cells in some time-interval (i.e. the counting process). In order to reduce the amount of measurement data that has to be transmitted through the network, the length of those time-intervals cannot be arbitrarily small (in practice, about 1 second as lower bound).

If such a calibration method was developed, it would be possible to use continuous on-line measurements to recompute the traffic parameters dynamically and thereby adjust to fluctuations in the load that are caused by unpredictable events (such as the sudden appearances of very popular documents on the Web that are accessed by many users). Based on the measured traffic parameters, congestion could be predicted and avoided by dynamic reconfiguration of the network, e.g. through the update of routing tables, or by more elaborate concepts as they are brought up in the context of *active* or *self-sizing* networks.

10.4 Output Process Study

The analysis in this thesis concentrates on a single switch in the network. Frequently, single bottlenecks can be identified in a network, therefore the performance results that are obtained in this thesis can be very useful for design decisions at those bottleneck switches. However, measurements have to be performed at the incoming links in order to parameterize the traffic at the switch entrance. Whenever the topology (or the routing table) of the network is changed, the traffic parameters at the individual switches can be affected and new measurements are necessary. A more flexible approach can be obtained if the traffic at the entrance points of the network is characterized by a parameterization of 1-Burst (or N-Burst) models, and and a method is developed to derive an approximation for the output traffic at each node.

If the traffic on each of the N incoming links is described by 1-Burst processes with a peak rate that is comparable or higher than the service-rate ν, or if the switch is highly utilized, then the output process of the switch can be expected to look like a 1-Burst process, whose ON periods are the Busy Periods of the switch. The Busy Period analysis of the N-Burst/M/1 model can then be used to obtain the parameters of that output 1-Burst process. Preliminary results of the Busy-Period analysis are discussed in [Schwefel 99e]. The scenario of the measurements of ATM cells at the B-WiN entrance in Sect. 4.1 falls into that category – which is the reason, why the 1-Burst model can be fitted successfully to the measured data, although the traffic is highly aggregated at the measurement point.

If the peak-rates on the incoming links are much smaller than the service-rate ν and the switch works at low utilization, then the 1-Burst model cannot be expected to provide an adequate approximation for the output process, but the N-Burst model may be more appropriate. The detailed parameters and the type of burst-start process (the load dependence factors in the burst-start rate) cannot be predicted without an in-depth investigation of such switch models. The Busy-Period analysis is again helpful to obtain the distribution of the time-intervals, during which the switch-output generates cells at maximal rate ν. However, the individual Busy Periods can be correlated in this scenario, since the N-Burst process that describes the switch input is not memoryless. That is in contrast to the above scenario with high peak-rates on the incoming links: if the Busy Period ends, all incoming links are very likely in the OFF state, which is the only memoryless state of the N-Burst input process.

10.5 Integration in Global Network Planning Process

Once the step of the previous section is accomplished, i.e. a function is developed that computes the N-Burst parameters of the approximate output process of a switch from the input N-Burst processes, the results of the performance analysis of this thesis can be integrated in network topology planning tools, such as described in [Fischer 00]. At this time, current planning tools use either no QoS constraints at all, or only simple M/M/1 models. At best, a topology is designed, and later on it is verified whether it meets certain QoS constraints.

The input to such sophisticated QoS oriented topology planning tools would then have to be a more detailed description of the incoming traffic at each access node than the currently used traffic matrices. As a consequence, the traffic forecasting becomes more complicated as well. This is very likely one of the major obstacles for the application of our more realistic N-Burst models in the global network planning process: there do not exist any satisfying methods for forecasting a simple traffic matrix, but expert intuition together with some heuristics are commonly used instead. Therefore chances are very small that such forecasting methods are developed quickly for the higher dimensional traffic matrices that are necessary for the more realistic N-Burst models. Nevertheless, the integration of better traffic models in the network planning process is a worthwhile long-term research goal.

10.6 Impact of Different Service Strategies

The models in this thesis use buffers with the FIFO (First In First Out) strategy. One of the consequence is that it is possible that a large number of cells from a single burst occupy the buffer and delay all other incoming bursts. However, other service strategies are possible and also partially implemented in modern switches. Certain priority levels can be assigned to cells of the different sources or even to cells of individual bursts. Cells of a lower priority level have to wait until all cells of higher levels have left the queue. In order to prevent starvation of the lower priority levels, more sophisticated service strategies, such as weighted fair queueing, can be implemented.

Such methods can be especially effective for Power-Tail distributed burst-lengths, since the occasional huge bursts cause the performance problems in that case. Already the approximation of the N-Burst/M/1 model by an M/G/1 queue demonstrates clearly that intelligent service strategies can alleviate performance problems greatly for Power-Tail distributed burst-lengths: In the M/G/1 queue with FIFO service strategy, the P-K formula (D.1) states that the mean

queue-length becomes infinite for service time distributions with infinite variance, which is the case for Power-Tailed distributed burst-lengths with $\alpha < 2$. If the service discipline in the M/G/1 queue is changed to processor sharing, which corresponds approximately to a round-robin service discipline for the cells in different bursts, the mean queue-length is finite.

The insight into the impact of different service strategies on the performance behavior of N-Burst/M/1 queues is in particular important for an evaluation of QoS oriented protocol mechanisms as they are discussed in the next section.

10.7 QoS-Oriented Protocol Mechanisms

The technology of telecommunication system has to take into account the growing demand of the users for a certain transmission quality (QoS), in particular for real-time application like voice or video traffic. One method to satisfy those QoS requirements is an over-provisioning of resources: links and switches or routers are designed overly 'large' so that any potential upcoming QoS requirement can be met without special treatment of the critical connections. How much over-provisioning is necessary in that case? The answer requires obviously a forecasting of the future traffic and an estimation of the QoS demands that future applications request. Thereafter the network has to be designed appropriately, see Sect. 10.5.

Another option is the implementation of protocol mechanisms that either prioritize the QoS critical traffic or alternatively reserve the necessary resources at the beginning of each of the critical connections. Both approaches are currently being discussed as possible improvements of the current Internet protocols (e.g. DiffServ and MPLS, respectively RSVP). However, in order to evaluate the effectiveness of such protocol mechanisms and also to obtain a good parameterization of the methods, performance evaluation with the help of stochastic models is necessary. Furthermore, any reservation based protocol mechanism requires Connection Admission Control (CAC), which is a straightforward application of the results of this thesis, see Sect. 7.4.

Appendix A
M/M/1 Queue

Although it is known that network traffic is not adequately described by a Poisson process, nevertheless the simple M/M/1 model is often used as a performance model in practical capacity design problems, see e.g. [Cahn 98]. The reason for that is twofold: First, a set of only two parameters, the service-rate ν and the utilization ρ (alternatively, the arrival rate λ), are sufficient to describe the M/M/1 model. Secondly, closed-form analytic expressions exist for many of its performance parameters.

Even though the M/M/1 model is not really realistic, it can be used as a base model for comparison purposes. Also, the transient M/M/1 queue with $\rho > 1$ describes the behavior of N-Burst/M/1 queues in some scenarios, in particular during over-saturation periods with N active sources. Therefore, we summarize some performance results for M/M/1 models in this chapter.

A.1 Steady-State Behavior

The M/M/1 queue can be represented as simple Birth-Death Markov Process, see e.g. [Kleinrock 75]. The steady-state queue-length probabilities follow as

$$\pi_i = \frac{\rho^i}{1-\rho} \quad \text{when} \quad \rho < 1 \cdot$$

By evaluation of the infinite sum $\sum_i i\pi_i$ the mean queue-length is obtained, and the mean delay follows from Little's law:

$$\mathbb{E}\{Q\} = \frac{\rho}{1-\rho}, \quad \text{mCD} = \frac{\mathbb{E}\{Q\}}{\lambda} = \frac{1}{\nu - \lambda} \cdot$$

The queue-length probabilities of finite M/M/1/B queues also show a geometric decay with the factor ρ:

$$\pi_i^{(B)} = \pi_0^{(B)} \rho^i, \quad \text{with} \quad \pi_0^{(B)} = \frac{1-\rho}{1-\rho^{B+1}} \cdot$$

A.2 Transient Behavior

The behavior of the transient M/M/1 queue with $\rho < 1$ is interesting as a base model for the comparison with the transient behavior of N-Burst/M/1 queues. In addition, the unstable cases $\rho = 1$ and $\rho > 1$ are interesting as well: the case $\rho > 1$ corresponds to the temporary behavior of N-Burst/M/1 queues in the blow-up regions $i_0 = N$ during over-saturation periods with N long-term active bursts. The case $\rho = 1$ corresponds to the behavior of the N-Burst/M/1 model right at the last blow-up point when $\Lambda_N = \nu$, see Sect. 10.1.

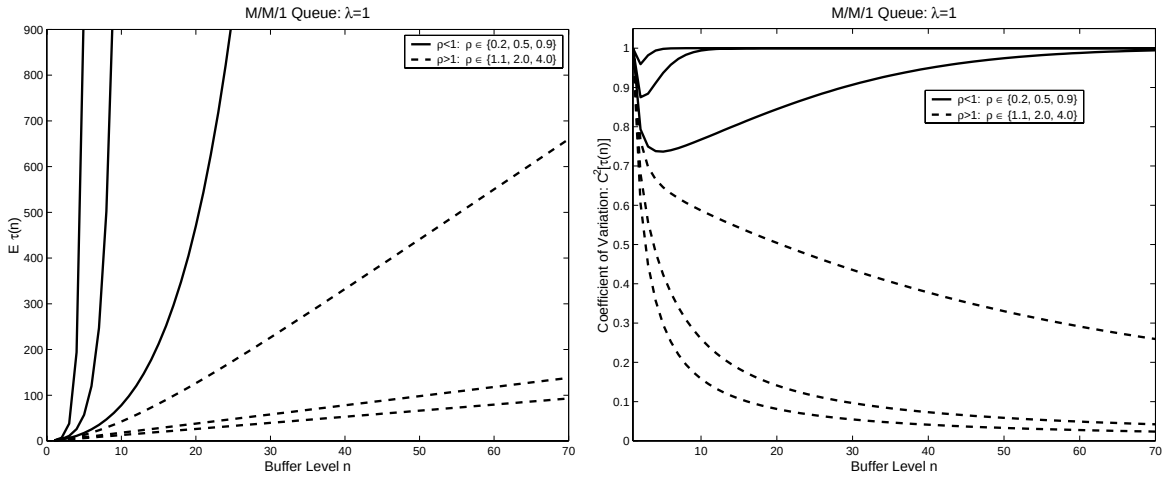

Figure A.1: **Expected Value and Coefficient of Variation of the First Passage Time τ_n for M/M/1 Queues:** The solid lines correspond to $\rho < 1$, while the dashed ones show the results for $\rho > 1$.

First Passage Process

Asymptotically for large n, the mean First Passage Time of an M/M/1 queue grows geometrically when $\rho < 1$, but only linearly when $\rho > 1$, see also the left graph of Fig. A.1:

$$\mathbb{E}\{\tau_n\} = \frac{1}{\nu - \lambda}\left(\frac{1-\rho^n}{\rho^n(1-\rho)} - n\right) \sim \begin{cases} \frac{1}{\nu(1-\rho)^2}\,\rho^{-n} & \text{when } \rho < 1 \\ \frac{1}{\lambda-\nu}\,n & \text{when } \rho > 1 \end{cases}$$

It is shown in [Asmussen et al. 00a] that the coefficient of variation of the FPT τ_n behaves asymptotically as:

$$C^2[\tau_n] = \frac{\mathbb{E}\{\tau_n^2\}}{(\mathbb{E}\{\tau_n\})^2} - 1 \sim \begin{cases} 1 & \text{when } \rho < 1 \\ \frac{\lambda+\nu}{\lambda-\nu}\,n^{-1} & \text{when } \rho > 1 \end{cases}$$

The right graph of Figure A.1 illustrates that asymptotic behavior: In the unstable case $\rho > 1$, the coefficient of variation converges to zero, which implies that the distribution of τ_n appears more and more deterministic, the larger the buffer-size n is.

At the border-line case $\rho = 1$, the mean First Passage Time grows quadratically, see e.g. [Lipsky 92]:

$$\mathbb{E}\{\tau_n\} = \frac{n(n+1)}{2\nu}, \quad \text{for} \quad \rho = 1.$$

The distribution of the FPT is shown in Figure A.2. For large n and $\rho < 1$, the density function is very close to an exponential density with same mean. However, the exponential density overestimates the true First-Passage Time distribution for very small value of t. In the unstable case $\rho > 1$, the asymptotic First-Passage Time density converges to a Gamma or Normal distribution. See [Asmussen et al. 00a] for a mathematically rigorous discussion of the asymptotic theory.

Busy Period Analysis

The Busy Period is defined as the time-interval that starts with an arrival at the idle server and ends, when the server becomes idle again for the first time. Let B be the duration of the busy

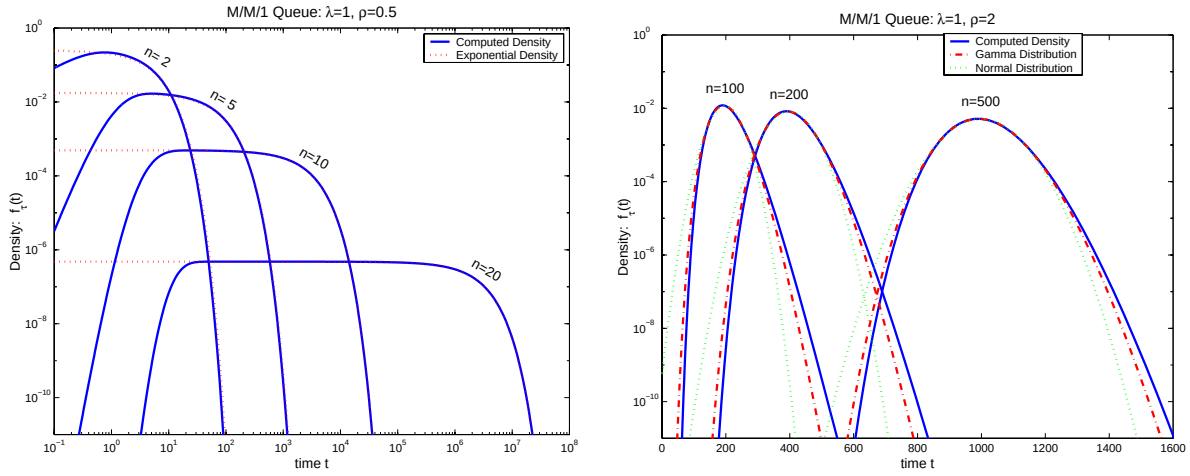

Figure A.2: **The distribution of the First Passage Time τ_n of M/M/1 queues:** The left graph shows curve on log-log scale for the stable case $\rho < 1$ in comparison to an exponential distribution. The right graph illustrates the unstable case $\rho > 1$. Fitted Gamma and Normal distributions can be used as an asymptotic approximation in the latter case.

period. It can be shown (see e.g. [Lipsky 92]) that:

$$\mathbb{E}\{B\} = \frac{1}{\nu - \lambda} \quad \text{for} \quad \rho < 1.$$

A Busy Cycle of the M/M/1 queue can be defined by the Busy Period and the preceding (alternatively, the succeeding) idle period. For $\rho < 1$, the duration of the Busy Cycle has the expected value

$$\mathbb{E}\{C\} = \frac{1}{\lambda} + \frac{1}{\nu - \lambda},$$

and the probability that the queue length reaches level n during the Busy Cycle is:

$$f_n = \mathbb{P}(\tau_n \leq C) = \frac{1-\rho}{1-\rho^n}\,\rho^{n-1}.$$

Appendix B

Matrix-Exponential Distributions

The fundamental modeling approach in this thesis uses modulated Poisson processes for the traffic description. With the concept of *Matrix Exponential* distributions (see Sect. B.1), the source models reduce to *Markov Modulated Poisson Processes (MMPPs)*, whose queueing behavior can be analyzed analytically by *Matrix-Analytic* methods. This and the next chapter contain a fundamental introduction into such a modeling approach.

B.1 Definition

Tractable analytic models are frequently based on *Markov Processes/Chains*. A stochastic process is called Markovian if its behavior only depends on its current state. For continuous time Markov Chains that property can only be achieved if the distribution of the state-holding times is memoryless, i.e. exponential.

When modeling network traffic, exponential state times (e.g. holding times) have turned out to be inadequate, see Chapter 3. Therefore an extension of the pure Markovian concept is necessary.

To get rid of the restrictions of exponential state times, so-called *Matrix-Exponential* (ME) distributions are used[1]. When replacing a single exponential state by a box that contains a network of states (called phases) with exponentially distributed state-times within the box, any distribution for the state time can be approximated arbitrarily closely. Figure B.1 demonstrates that process. This approach was first brought up as Phase-type distributions by Neuts in 1975 (see [Neuts 81]), and is generalized to ME distributions by [Lipsky 92]. The following paragraph contains a brief introduction and mentions the necessary formulas in reference to the latter.

The distribution of the time between a single customer entering a subsystem (the box) such as in Figure B.1 and leaving it again can be determined elegantly by using matrix algebra. Let $\mathbf{P} = (\mathbf{P}_{ij})$ be the matrix of the *transition probabilities* within the subsystem. It is sub-stochastic (row sums ≤ 1), since the probabilities of leaving the box do not show up in this matrix.

Next we define the *completion rate matrix* $\mathbf{M}$, which is a diagonal matrix of the single state leaving rates μ_k. Furthermore, let $\boldsymbol{\tau}$ be a vector whose components τ_i are the mean times to leave the system, given that the customer started at phase i. Then the following equation holds:

$$\boldsymbol{\tau}' = \mathbf{M}^{-1}\boldsymbol{\varepsilon}' + \mathbf{P}\boldsymbol{\tau}'.$$

The first summand represents the mean time at the current phase, while the second summand is the mean time in the subsystem after the next transition. Solving for $\boldsymbol{\tau}$ gives:

$$\boldsymbol{\tau}' = [\mathbf{M}(\mathbf{I} - \mathbf{P})]^{-1}\boldsymbol{\varepsilon}' =: \mathbf{V}\boldsymbol{\varepsilon}'.$$

[1] Related concepts are Coxian Servers, Kendall Distributions, and Phase Distributions.

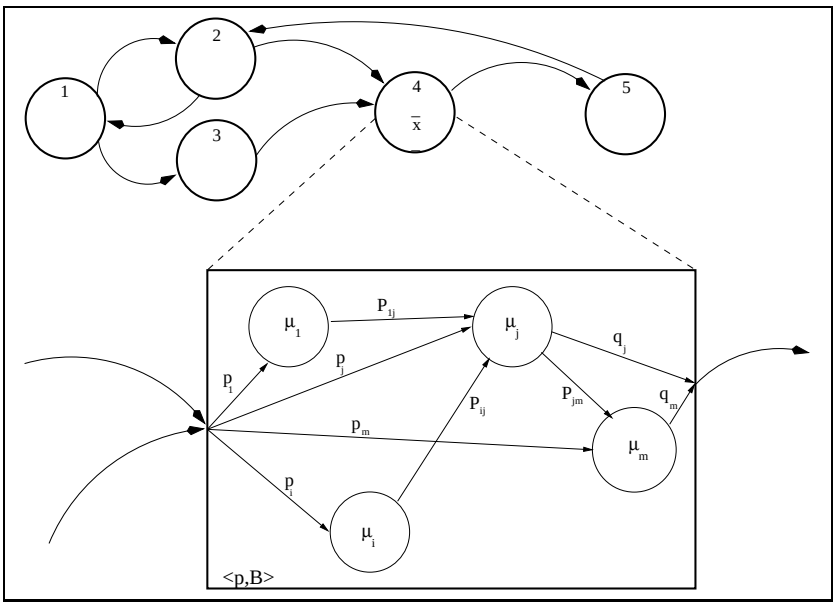

Figure B.1: **Generalization of Markov Processes by Matrix-Exponential State Times:** The exponential state times are generalized by replacing the single state (here #4 in the upper Markov-Process) by a bigger box, containing a network of phases. The structure of the replacement box is described by its entry vector $\mathbf{p} = [p_1, p_2, ...]$ and the rate-matrix $\mathbf{B} := \mathbf{M}(\mathbf{I} - \mathbf{P})$. $\mathbf{M}$ is a diagonal matrix of the leaving rates of the individual states within the box, and $\mathbf{P} = (\mathbf{P}_{kl})$ is a sub-stochastic matrix, describing the probabilities of state changes within the box. The distribution of the time spent in the box (which is now the generalized state #4) turns out to be $R(x) = \mathbf{p}\exp(-x\mathbf{B})\,\boldsymbol{\varepsilon}'$, therefore the name Matrix-Exponential (ME) distribution.

Therefore, $\mathbf{V} := [\mathbf{M}(\mathbf{I} - \mathbf{P})]^{-1}$ is called the *service time matrix*. Its elements V_{ij} are the overall mean times spent at phase j until the customer leaves the system, given that it started at phase i.

The inverse of $\mathbf{V}$, $\mathbf{B} := \mathbf{M}(\mathbf{I} - \mathbf{P})$, has the name *service rate matrix*. The distribution of the time that the customer spends in the box follows from the Kolmogorov differential equations:

$$\frac{d\mathbf{R}(t)}{dt} = -\mathbf{R}(t)\mathbf{B} \implies \mathbf{R}(t) = \exp(-t\mathbf{B}),$$

where $R_{ij}(t)$ is the probability that the customer is in phase j at time t, given that it started in i.

The components p_i of the *entrance vector* $\mathbf{p}$ are the probabilities that a new customer upon entering the system will directly go to phase i. Since this vector $\mathbf{p}$ is constant, the new customer enters the system independently of any previous customer. Thus, the process that describes the time in the subsystem for a series of customers is a renewal process.

The pair $\langle \mathbf{p}, \mathbf{B} \rangle$ completely describes the subsystem and gives rise to an elegant formula for the reliability function of the time that the customer spends in the box:

$$R(x) = \mathbf{p} \cdot \exp(-x\mathbf{B}) \cdot \boldsymbol{\varepsilon}'. \tag{B.1}$$

Thus the ME representation, $\langle \mathbf{p}, \mathbf{B} \rangle$, defines the distribution (B.1). If the underlying probabilistic structure is given, $\langle \mathbf{p}, \mathbf{B} \rangle$ is exactly the class of Phase-type distributions of [Neuts 81].

However, the concept can be generalized: Any vector-matrix pair $\langle \mathbf{p}, \mathbf{B} \rangle$ that generates a valid reliability function via (B.1) can be used, e.g. $\mathbf{p}$ and $\mathbf{B}$ could contain negative or complex elements. Such distributions are called ME distributions. The class of distributions that is covered by the ME concept is broader than the Phase-type class, see Appendix B.2 for an example.

Differentiation of the reliability function yields the pdf:

$$f(x) = -\frac{dR(x)}{dx} = \mathbf{pB}\exp(-x\mathbf{B})\boldsymbol{\varepsilon}'.$$

Also, the nth moments of these distributions follow from (1.1):

$$\mathbb{E}\{X^n\} = n!\mathbf{p}\mathbf{V}^n\boldsymbol{\varepsilon}'. \tag{B.2}$$

B.2 Special ME Distributions

There are a few basic types of distributions that are essential for modeling network traffic. They are described in this section together with their ME representation $\langle \mathbf{p}, \mathbf{B} \rangle$.

Exponential Distribution

The Exponential distribution (1.4) is the trivial ME distribution. The replacement box of Figure B.1 just contains one state, thus the matrices have dimension $T = 1$, i.e. they are scalars:

$$\mathbf{p} = 1, \quad \mathbf{B} = \mu, \quad \mathbf{V} = \frac{1}{\mu} =: \bar{x}.$$

The exponential distribution has as its only parameter its rate μ. Therefore it is completely defined by its mean $\mathbb{E}\{X\} = 1/\mu$. The coefficient of variation is always $C^2 = 1$.

Hyperexponential-2 Distribution

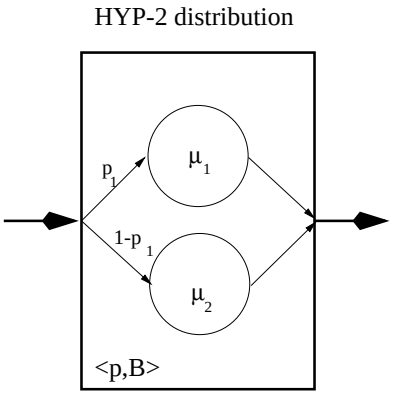

Figure B.2: **ME Representation of a HYP-2 Distribution:** The simplest case of a distribution with higher coefficient of Variation than the exponential distribution is a Hyperexponential distribution with 2 states. The $\mathbf{P}$-matrix is zero, since there are no transitions within the box. Thus $\mathbf{B} = \mathbf{M}$.

Figure B.2 shows the ME-representation of a two-state hyperexponential distribution (HYP-2). It has the reliability function:

$$R(x) = p_1\, e^{-\mu_1\, x} + (1 - p_1)\, e^{-\mu_2\, x}. \tag{B.3}$$

This function has three free parameters (p_1, μ_1, μ_2), two of which can be fixed by matching the first two moments:

$$\mathbb{E}\{X\} = \frac{p_1}{\mu_1} + \frac{1-p_1}{\mu_2} =: \bar{x}, \quad \mathbb{E}\left\{X^2\right\} = 2\left(\frac{p_1}{\mu_1^2} + \frac{1-p_1}{\mu_2^2}\right) \quad (\Rightarrow \quad C^2 > 1).$$

Any distribution with $C^2 > 1$ can be matched in its first 2 moments by a HYP-2 distribution in the following way:

$$\mu_2 = \frac{1}{\mathbb{E}\{X\}\left[1 - \sqrt{\frac{p_1}{1-p_1} \cdot \frac{C^2-1}{2}}\right]}, \quad \mu_1 = \frac{p_1}{\mathbb{E}\{X\} - \frac{(1-p_1)}{\mu_2}}.$$

The probability p_1 can still be chosen freely with the restriction that:

$$p_1 < p_{max} = 2\,\frac{\mathbb{E}\{X\}^2}{\mathbb{E}\{X^2\}} = \frac{2}{C^2+1}.$$

In fact, the smaller p_1, the larger becomes the third moment $\mathbb{E}\left\{X^3\right\}$. Consequently, p_1 can be used to match the third moment of a given distribution within some range.

LAQT-Matrices:

$$\mathbf{p} = [p_1, (1-p_1)], \quad \mathbf{B} = \begin{bmatrix} \mu_1 & 0 \\ 0 & \mu_2 \end{bmatrix}, \quad \mathbf{V} = \begin{bmatrix} \frac{1}{\mu_1} & 0 \\ 0 & \frac{1}{\mu_2} \end{bmatrix}.$$

Erlangian-T Distribution

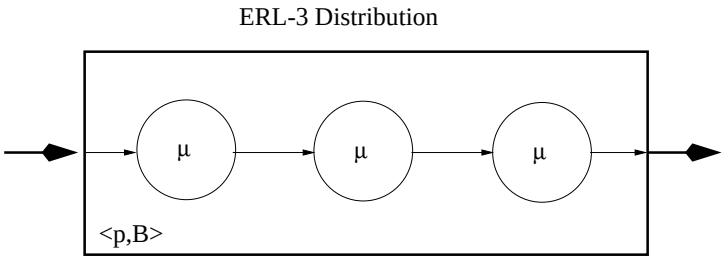

Figure B.3: **ME Representation of Erlangian-3 Distribution:** Distributions with small coefficient of variation ($C^2 < 1$) can be approximated by Erlangian-T distributions. The Figure shows the case $T = 3$.

An Erlangian-T distribution is the convolution of T identical exponentials, i.e. the distribution of the sum of T iid exponential random variables. Its probability density function is:

$$f(x) = \mu\frac{(\mu x)^{T-1}}{(T-1)!}e^{-\mu x}. \tag{B.4}$$

Note that the density at the origin $x = 0$ is $f(0) = 0$ as opposed to Hyper-exponential distributions, whose density has its maximum at the origin.

It follows for the first two moments of the Erlangian-T distribution:

$$\mathbb{E}\{X\} = \frac{T}{\mu} =: \bar{x}, \quad \mathbb{E}\left\{X^2\right\} = \frac{T^2+T}{\mu^2}. \quad \Rightarrow \quad C^2 = \frac{1}{T} < 1.$$

An Erlangian-T distribution with high T can be used as approximation for a deterministic distribution (e.g. for service times within ATM-switches), which has variation $C^2 = 0$.

LAQT-Matrices:

$$\mathbf{p} = [1, 0, \ldots, 0], \qquad \mathbf{B} = \mu \begin{bmatrix} 1 & -1 & & & \\ & 1 & -1 & & \\ & & \ddots & \ddots & \\ & & & 1 & -1 \\ & & & & 1 \end{bmatrix},$$

$$\mathbf{V} = \frac{1}{\mu} \begin{bmatrix} 1 & 1 & \ldots & 1 \\ & \ddots & \ddots & \vdots \\ & & 1 & 1 \\ & & & 1 \end{bmatrix}.$$

Truncated Power-Tail Distributions

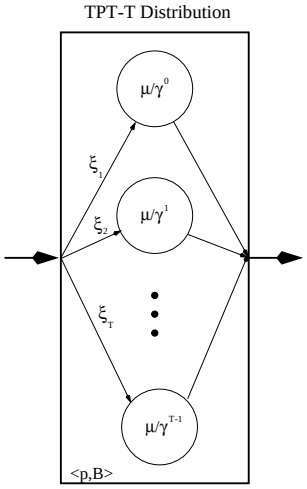

Figure B.4: **Phase diagram for the TPT-T distribution:** The probabilities $\zeta_i = \left[\theta^{i-1}(1-\theta)\right]/(1-\theta^T)$ of entering phase i decay geometrically by a factor $\theta < 1$. Furthermore, the state-holding times grow geometrically by the factor $\gamma = 1/\theta^{1/\alpha}$.

When using special hyper-exponentials such as in Figure B.4, the reliability function of those distributions,

$$R_{Y_T}(x) = \frac{1-\theta}{1-\theta^T} \sum_{i=0}^{T-1} \theta^i \exp\left[\frac{-\mu_T}{\gamma^i} x\right], \tag{B.5}$$

show Power-law behavior, $R(x) \sim x^{-\alpha}$ for some orders of magnitude before they drop off exponentially, see [Greiner et al. 99] and Sect. 3.4. The higher the number of phases, T, the later the drop-off occurs. The exponential drop-off is characterized in more detail in Sect. 3.5 by the so-called Power-Tail Range.

The variable θ can be chosen freely in the range $0 < \theta < 1$. For larger value of θ, more phases are necessary to obtain the same PT Range as for lower θ. In order to show Power-Law behavior with exponent α, and to have mean $\bar{x}$, the other constants in (B.5) have to be (see [Greiner et al. 99]):

$$\gamma = \left(\frac{1}{\theta}\right)^{1/\alpha},$$
$$\mu_T = \frac{1-\theta}{1-\theta^T}\,\frac{1-(\theta\gamma)^T}{1-\theta\gamma}\,\frac{1}{\bar{x}}.$$

LAQT-Matrices:

$$\mathbf{p}_T = \frac{1-\theta}{1-\theta^T}\left[\theta^0, \ldots, \theta^{T-1}\right].$$
$$\mathbf{B}_T = \mu_T \begin{bmatrix} 1/\gamma^0 & & \mathbf{0} \\ & \ddots & \\ \mathbf{0} & & 1/\gamma^{T-1} \end{bmatrix}.$$
$$\mathbf{V}_T = \frac{1}{\mu_T} \begin{bmatrix} \gamma^0 & & \mathbf{0} \\ & \ddots & \\ \mathbf{0} & & \gamma^{T-1} \end{bmatrix}.$$

Example of a non Phase-Type distribution

So far, all distributions in this section are of Phase-Type. As an example of a distribution which has an ME representation but is not Phase-Type, we discuss briefly the following density function:

$$f(x) = \frac{1}{c}(1-\sin x)\,e^{-x} = \frac{1}{c}\left[e^{-x} - \frac{1}{2i}\,e^{(i-1)\,x} + \frac{1}{2i}e^{-(i+1)\,x}\right],$$

where $c := \int_0^\infty (1-\sin x)e^{-x}\,dx$ and $i := \sqrt{-1}$ is the imaginary unit. This distribution has the following ME representation:

$$\mathbf{p} = \frac{1}{c}\left[1,\ -\frac{1}{2\,(i+1)},\ \frac{1}{2\,(i-1)}\right], \quad \mathbf{B} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & -(i-1) & 0 \\ 0 & 0 & i+1 \end{bmatrix}.$$

From $\mathbf{p}\boldsymbol{\varepsilon}' = 1$, it follows that $c = 1/2$.

The density function $f(x_k) = 0$ for infinitely many

$$x_k = (2k+1)\,\frac{\pi}{2}, \quad k = 0, 1, \ldots$$

It can be shown that no Phase-Type distribution can have such a property. Consequently, this distribution provides an example of an ME distribution which is not Phase-Type.

Mixtures of Distributions

In the N-Burst model, the ME distributions are used for the burst-length distribution, i.e. the length of a single ON period during which cells are transmitted. Different applications result in different distributions of the burst-lengths, i.e. file sizes are known to be TPT-distributed ([Garg et al. 92]) thus FTP-data transfers are as well. On the other hand, the frame-sizes

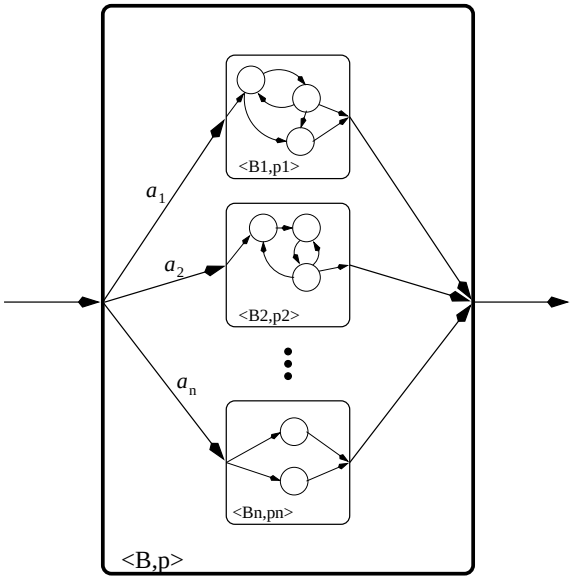

Figure B.5: **ME Representation of Burst-Length Distribution for Traffic Mixes:** The notion of *traffic mixes* is implemented by mixing different distributions $< \mathbf{p_1}, \mathbf{B_1} >$, ..., $< \mathbf{p_n}, \mathbf{B_n} >$ according to their proportions $a_1, ..., a_n$ ($\sum_i a_i = 1$). The resulting matrix $\mathbf{B}$ is block-diagonal.

of a video transmission are closer to Erlangian distributions. In a heterogeneous network, such applications usually coexist, often called *traffic-mix*.

ME distributions allow for a simple incorporation of this idea: knowing the proportions $a_1, ...a_n$, ($\sum_i a_i = 1$) of each application and their distributions $< \mathbf{p_n}, \mathbf{B_n} >$, the overall burst-length distribution $\langle \mathbf{p}, \mathbf{B} \rangle$ can be constructed in a way shown in Figure B.5. One obtains a mixture of the individual ME distributions with ME representation:

$$\mathbf{p} = [a_1\mathbf{p_1}, \ \ldots, \ a_n\mathbf{p_n}], \quad \text{where} \quad \sum_{i=1}^{n} a_i = 1.$$

$$\mathbf{B} = \begin{bmatrix} \mathbf{B_1} & \ldots & \mathbf{O} \\ \mathbf{O} & \ddots & \mathbf{O} \\ \mathbf{O} & \ldots & \mathbf{B_n} \end{bmatrix}, \quad \mathbf{V} = \begin{bmatrix} \mathbf{B_1}^{-1} & \ldots & \mathbf{O} \\ \mathbf{O} & \ddots & \mathbf{O} \\ \mathbf{O} & \ldots & \mathbf{B_n}^{-1} \end{bmatrix}.$$

Hyper-Erlangian Distributions: A special case of a distribution mix is obtained when the individual distributions are all identical Erlangians. The distribution mix is called *Hyper-Erlangian* in this case. The relevance of such distributions is that they can have a high variance, $C^2 \gg 1$, but zero density at the origin $[f(0) = 0]$.

B.3 Power-Tail Range of ME and Matrix Geometric Distributions

In the following, we look at ME distributions $\langle \mathbf{p}, \mathbf{B} \rangle$ whose rate-matrix $\mathbf{B}$ is diagonalizable,

$$\mathbf{B} = \mathbf{U_B}^{-1} \cdot \mathbf{diag}\left([\lambda_1, ..., \lambda_n]\right) \cdot \mathbf{U_B}$$

where λ_i are the eigenvalues of $\mathbf{B}$. The spectral decomposition of the ME reliability function (B.1),

$$R(x) = \sum_i \left[\left(\mathbf{p}\mathbf{U}_\mathbf{B}{}^{-1}\right)_i \exp\left(-x\,\lambda_i(\mathbf{B})\right) \left(\mathbf{U}_\mathbf{B}\boldsymbol{\epsilon}'\right)_i \right] ,$$

shows that the last exponential drop-off of the reliability function is determined by the minimum real-part of the eigenvalues $\lambda_i(\mathbf{B})$ of $\mathbf{B}$. Here, we assume that $\left(\mathbf{p}\mathbf{U}_\mathbf{B}{}^{-1}\right)_i$ and $\left(\mathbf{U}_\mathbf{B}\boldsymbol{\epsilon}'\right)_i$ are non-zero, otherwise a reduction of the state-space could be performed.

Using the spectral decomposition, the PT Range as defined in Eq. (3.9), can be computed for ME distributions by:

$$\mathrm{Rng}\left(< \mathbf{p}, \mathbf{B} >\right) = \max_i \left\{ \frac{1}{Re\left\{\lambda_i(\mathbf{B})\right\}} \right\} . \tag{B.6}$$

Note that the PT Range is independent of $\mathbf{p}$. Therefore, we can abbreviate, $\mathrm{Rng}\left(\mathbf{B}\right) := \mathrm{Rng}\left(\langle \mathbf{p}, \mathbf{B} \rangle\right)$.

In case of two independent ME distributions with rate-matrices $\mathbf{B}_1$ and $\mathbf{B}_2$, the distribution of the time until one of them finishes has the rate matrix $\mathbf{B}_1 \oplus \mathbf{B}_2$. Thus, its PT Range is

$$\mathrm{Rng}\left(\mathbf{B}_1 \oplus \mathbf{B}_2\right) = \max_i \left\{ \frac{1}{Re\left\{\lambda_i(\mathbf{B}_1 \oplus \mathbf{B}_2)\right\}} \right\} = \frac{1}{\frac{1}{\mathrm{Rng}(\mathbf{B}_1)} + \frac{1}{\mathrm{Rng}(\mathbf{B}_2)}} . \tag{B.7}$$

For i independent ME distributions with the same rate-matrix $\mathbf{B}$, the PT Range of the distribution of the time until one of them finishes is given by:

$$\mathrm{Rng}\left(\mathbf{B}^{\oplus i}\right) = \frac{\mathrm{Rng}\left(\mathbf{B}\right)}{i} . \tag{B.8}$$

Discrete matrix-geometric distributions: The discrete equivalent of a ME distribution is a Matrix-Geometric distribution:

$$p_k = \mathbf{a}\mathbf{A}^k\mathbf{b}' .$$

The two major examples for Matrix-Geometric distributions are the queue-length distribution of an SM/M/1 queue in Appendix D.5 and the distribution of the number of cells during a burst in Appendix C.4.

The PT range of p_k is determined by the Eigenvalue λ_{max} of $\mathbf{A}$ with largest real part:

$$p_k \sim \lambda_{max}^k = \exp\left[-(-\ln\lambda_{max})\,k\right] .$$

Consequently, the PT Range of p_k follows as

$$\mathrm{Rng}\left(\mathbf{A}\right) := \mathrm{Rng}\left(p_k\right) = \frac{1}{-\ln\lambda_{max}(\mathbf{A})} . \tag{B.9}$$

B.4 Residual Time of ME Distributions

[Lipsky 92] showed that the residual time of a ME distribution $< \mathbf{p}, \mathbf{B} >$ with mean $\bar{x} = \mathbf{p}\mathbf{B}^{-1}\boldsymbol{\epsilon}'$ for a randomly arriving observer has the ME representation

$$< \mathbf{p}\mathbf{B}^{-1}/\bar{x}, \mathbf{B} > . \tag{B.10}$$

Therefore, the density function of the residual time,

$$f_{\text{res}}(x) = \frac{\mathbf{p}\mathbf{B}^{-1}}{\bar{x}}\,\mathbf{B}\,\exp\left(-x\mathbf{B}\right)\boldsymbol{\epsilon}' = \frac{1}{\bar{x}}\,R(x)\,, \tag{B.11}$$

is the appropriately scaled reliability function of the original ME distribution.

This relationship holds in general, not only for ME distributions, see e.g. [Kleinrock 75].

Appendix C
Semi-Markov Processes

C.1 Semi-Markov Processes

The ME distributions of Section B.1 are very useful for describing renewal processes. However, as already mentioned, network traffic cannot possibly be modeled by a renewal process, since the individual inter-cell times are highly correlated. Therefore, [Fiorini et al. 95] extended the concepts of Section B.1 to LAQT representations of correlated *Semi-Markov* (SM) processes. If only the sub-class of distributions with Phase-Type representation are used within the SM processes, they are identical to the Markovian Arrival Processes (MAPs) in [Lucantoni 93].

Notation: Later on, SM processes are introduced that use embedded ME distributions. Matrices that refer to the SM-process will always be represented by caligraphic letters (e.g. $\boldsymbol{\mathcal{B}}$, $\boldsymbol{\mathcal{V}}$, $\boldsymbol{\mathcal{L}}$, $\boldsymbol{\mathcal{Y}}$, $\boldsymbol{\mathcal{I}}$) Be careful not to confuse those with the simultaneously used same Roman letters ($\mathbf{B}$, $\mathbf{V}$, $\mathbf{I}$, $\mathbf{p}$) for the embedded ME distribution process. $\mathbf{I}$ and $\boldsymbol{\mathcal{I}}$ are both identity matrices but possibly of different dimensions. Except for $\boldsymbol{\wp}$, vectors that refer to the SM system are usually also written Roman, since they always appear together with other matrices, and thus confusion should not arise.

The LAQT representation of correlated arrivals also uses a network of phases such as in Figure B.1. However, to introduce dependence, after a customer leaves the subsystem the next customer enters in a phase that depends on the leaving phase of the previous customer.

Similar to $\mathbf{B}$ for the renewal-process, $\boldsymbol{\mathcal{B}} := \boldsymbol{\mathcal{M}}(\boldsymbol{\mathcal{I}} - \boldsymbol{\mathcal{P}})$ describes the behavior of the customer in the box until leaving the box. Then, another matrix $\boldsymbol{\mathcal{L}}$ is introduced. The components $\boldsymbol{\mathcal{L}}_{ij}$ are the departure rates from phase i, whereafter the next customer starts in phase j, i.e. for small time intervals Δ, $(\boldsymbol{\mathcal{L}}_{ij} \cdot \Delta)$ is the probability that a customer at phase i leaves within Δ and the next one starts in phase j. $\boldsymbol{\mathcal{L}}$ and $\boldsymbol{\mathcal{B}}$ must be consistent, i.e. $\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}' = \boldsymbol{\mathcal{B}}\boldsymbol{\varepsilon}'$, because both matrices describe the departure rate of the customers. However, $\boldsymbol{\mathcal{L}}$ does not capture what is going on before the departure, while $\boldsymbol{\mathcal{B}}$ does not affect the next customer. Analogous to the matrix representation of ME distributions, we define $\boldsymbol{\mathcal{V}} := \boldsymbol{\mathcal{B}}^{-1}$.

The entrance vector $\mathbf{p}$ of the renewal processes is replaced by the $\boldsymbol{\mathcal{L}}$-matrix of the SM-Process. However, an initial entry vector is important when investigating the transient behavior of SM processes. The equivalent of the $\mathbf{p}$-vector for renewal processes is the steady-state entry vector, $\boldsymbol{\wp}$, whose i-th component is the probability that a newly arriving customer enters in phase i, when the influence of the startup of the process has ceased (i.e. in steady state). [Fiorini et al. 95] showed that $\boldsymbol{\wp}$ is the left eigenvector of the matrix $\boldsymbol{\mathcal{Y}} := \boldsymbol{\mathcal{V}}\boldsymbol{\mathcal{L}}$ with eigenvalue 1, i.e.

$$\boldsymbol{\wp} = \boldsymbol{\wp}\boldsymbol{\mathcal{Y}}.$$

Analogous to (B.2), the moments of the SM-Process are:

$$\mathbb{E}\left\{X^k\right\} = k!\,\boldsymbol{\wp}\boldsymbol{\mathcal{V}}^k\boldsymbol{\varepsilon}'. \tag{C.1}$$

C.2 Markov Modulated Poisson Processes

The SM processes of the last section supply a means of modeling correlated and bursty traffic on a network line. The N-Burst model belongs to a smaller family of processes within the SM processes, called *Markov Modulated Poisson Processes (MMPPs)*, for which the $\mathcal{L}$-matrix is diagonal, i.e. generated cells (departures from the subsystem) do not change the state of the subsystem.

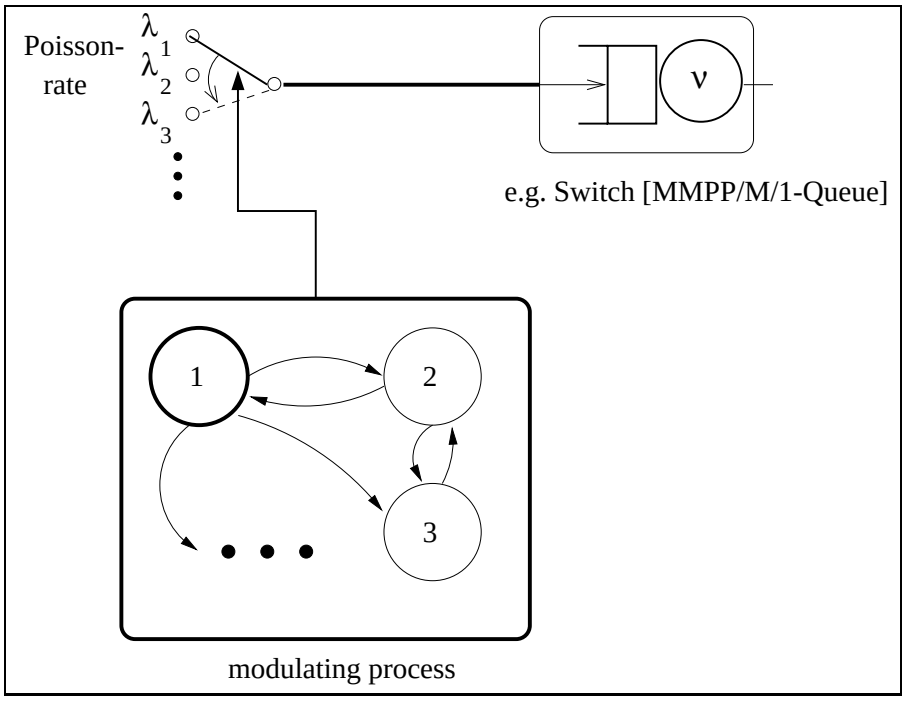

Figure C.1: **Illustration of MMPP Arrival Process:** A Markov Modulated Poisson Process (MMPP) consists of a Markov Process (lower box) which describes the modulating process: A token moves around in the box but never leaves it. The actual arrivals of e.g. cells to a system (e.g. Switch) are generated by a Poisson process, whose rate is determined by the position of the token within the modulating process. In the N-Burst Model of Chapter 2 the modulating process is a closed queueing system of M/G/N//K-type which describes the arrivals and the duration of the individual bursts.

The idea of SM processes is that when a customer leaves the system to go somewhere (e.g. as arrival to the exponential server that models the ATM-switch) the next one comes in according to the $\mathcal{L}$-Matrix. For MMPPs $\mathcal{L}$ is diagonal, and the mathematically equivalent abstraction is possible that a token moves around in a closed system (describing the number of active bursts) and the actual arrivals (cells) to the switch are separate entities. Both, the SM-view and the MMPP-view describe exactly the same process. The only difference is that the first one uses the customer of the modulating system as cell-arrival to the SM/M/1 queue and replaces it immediately with a new one, while in the MMPP-view, the modulating process is a closed system.

Figure C.1 illustrates the MMPP-view: The cells are always generated by a Poisson Process. However, the Poisson rate depends on the state of a closed Markov Process. Let $\mathcal{Q}$ be the generator[1] of that modulating Markov process.

[1]The generator of a Markov Process is a matrix containing the transition rates $\mathcal{Q}_{ij}$ from state i to state j in its non-diagonal elements. The diagonal elements are such that $\mathcal{Q}\boldsymbol{\varepsilon}' = \mathbf{0}$.

The stationary probability vector[2] $\boldsymbol{\pi}$ of the process has the property:

$$\boldsymbol{\pi}\boldsymbol{\mathcal{Q}} = \mathbf{0}.$$

As mentioned before, the $\boldsymbol{\mathcal{L}}$ matrix is diagonal and contains the corresponding Poisson rates:

$$\boldsymbol{\mathcal{L}} = \begin{bmatrix} \lambda_1 & 0 & \dots \\ 0 & \lambda_2 & \ddots \\ 0 & \dots & \ddots \end{bmatrix}.$$

The $\boldsymbol{\mathcal{B}}$-Matrix which describes the SM-view of the process (i.e. the customer moves around in the box and eventually leaves to go to the switch) must add these additional departures of $\boldsymbol{\mathcal{L}}$ to the description $\boldsymbol{\mathcal{Q}}$ of the modulating process:

$$\boldsymbol{\mathcal{B}} = (-\boldsymbol{\mathcal{Q}}) + \boldsymbol{\mathcal{L}}.$$

Note that by definition the main diagonal of $\boldsymbol{\mathcal{B}}$ contains positive elements while the elements on the main diagonal of $\boldsymbol{\mathcal{Q}}$ are negative. Hence, the minus in the formula.

As a consequence of $\boldsymbol{\pi}\boldsymbol{\mathcal{Q}} = \mathbf{0}$, it follows that

$$\boldsymbol{\pi}\boldsymbol{\mathcal{B}} = \boldsymbol{\pi}\boldsymbol{\mathcal{L}},$$

and, consistent with our original definition,

$$\boldsymbol{\mathcal{B}}\boldsymbol{\varepsilon}' = \boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}'.$$

Therefore the vector

$$\boldsymbol{\wp} = \frac{\boldsymbol{\pi}\boldsymbol{\mathcal{L}}}{\boldsymbol{\pi}\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}'} = \frac{\boldsymbol{\pi}\boldsymbol{\mathcal{B}}}{\boldsymbol{\pi}\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}'}$$

has the property $\boldsymbol{\wp}\boldsymbol{\mathcal{Y}} = \boldsymbol{\wp}$, where $\boldsymbol{\mathcal{Y}} = \boldsymbol{\mathcal{V}}\boldsymbol{\mathcal{L}}$, i.e. $\boldsymbol{\wp}$ is left eigenvector of $\boldsymbol{\mathcal{Y}}$ for the eigenvalue 1 and thus is the steady-state entry vector in the SM view of the MMPP.

Since

$$\boldsymbol{\wp}\boldsymbol{\mathcal{V}} = \frac{\boldsymbol{\pi}}{\boldsymbol{\pi}\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}'},$$

the mean inter-cell time of the MMPP comes out to be

$$\mathbb{E}\{X\} = \boldsymbol{\wp}\boldsymbol{\mathcal{V}}\boldsymbol{\varepsilon}' = \frac{1}{\boldsymbol{\pi}\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}'},$$

which makes sense intuitively, since $\boldsymbol{\pi}\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}' = \sum_i \pi_i\lambda_i$ is the average cell rate of the MMPP.

The modulating process of an MMPP is a Markov process, i.e. the state times are exponential. Generalized modulated Poisson processes are used in Chapter 2, but the non-exponential states are replaced by the ME representation of the general distribution (as in Figure B.1). After opening up the ME boxes, the modulating process is a Markov process again, if the ME distribution is of Phase-Type – which is the case almost always in this thesis.

MMPPs are a subclass of SM processes. The inter-cell times that are generated by MMPPs are correlated, except for models with $\mathrm{rank}(\boldsymbol{\mathcal{Y}}) = 1$, in which case the MMPP reduces to a renewal process, see the next section. The marginal distribution of the generated inter-cell times in steady-state has the ME representation $< \boldsymbol{\wp}, \boldsymbol{\mathcal{B}} >$.

The N-Burst model of Chapter 2 belongs to the class MMPP. However, variations such as models that generate cells right at burst-starts require the broader SM concept (with non-diagonal $\boldsymbol{\mathcal{L}}$).

[2] i.e. its components $\boldsymbol{\pi}_i$ are the probabilities that the modulating process is in state i at arbitrary points of time.

C.3 Autocorrelation of Inter-Cell Times and Counts

LRD properties manifest themselves in the correlation structure of a stochastic process. For the inter-cell times of a SM process, [Fiorini et al. 95] derive a formula for the covariance of the process in steady state:

$$\mathrm{Cov}(X, X_{+k}) = \boldsymbol{\wp}\boldsymbol{\mathcal{V}}(\boldsymbol{\mathcal{Y}}^k - \boldsymbol{\varepsilon}'\boldsymbol{\wp})\boldsymbol{\mathcal{V}}\boldsymbol{\varepsilon}'.$$

From there, the coefficient of autocorrelation follows by definition:

$$r(k) = \frac{\mathrm{Cov}(X, X_{+k})}{\mathrm{Var}(X)} = \frac{\boldsymbol{\wp}\boldsymbol{\mathcal{V}}(\boldsymbol{\mathcal{Y}}^k - \boldsymbol{\varepsilon}'\boldsymbol{\wp})\boldsymbol{\mathcal{V}}\boldsymbol{\varepsilon}'}{2\boldsymbol{\wp}\boldsymbol{\mathcal{V}}^2\boldsymbol{\varepsilon}' - (\boldsymbol{\wp}\boldsymbol{\mathcal{V}}\boldsymbol{\varepsilon}')^2}.$$

The covariance of the counting process $N_i(\Delta)$ of an MMPP for an interval size $\Delta > 0$ is obtained in [Neuts 89]:

$$\mathrm{Cov}(N, N_{+k}) = \boldsymbol{\pi}\boldsymbol{\mathcal{L}}\,[\boldsymbol{\mathcal{I}} - \exp(\Delta\boldsymbol{\mathcal{Q}})]^2\,\exp[(k-1)\Delta\boldsymbol{\mathcal{Q}}]\,(\boldsymbol{\varepsilon}'\boldsymbol{\pi} - \boldsymbol{\mathcal{Q}})^{-2}\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}'\,.$$

Let $\lambda := \pi\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}'$ be the average cell-rate of the MMPP. Then, the expected value of N in steady state is

$$\mathbb{E}\{N\} = \frac{\Delta}{\mathbb{E}\{X\}} = \lambda\Delta\,.$$

The second moment of N is (see [Meier-Hellstern & Fischer 92]):

$$\begin{aligned}\mathbb{E}\left\{N^2\right\} \;=\; & \lambda\Delta + \lambda^2\Delta^2 + 2\Delta\,\left[\lambda^2 - \boldsymbol{\pi}\boldsymbol{\mathcal{L}}(\boldsymbol{\varepsilon}'\boldsymbol{\pi} + \boldsymbol{\mathcal{Q}})^{-1}\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}'\right] + \\ & +2\boldsymbol{\pi}\boldsymbol{\mathcal{L}}\,[\exp(\Delta\boldsymbol{\mathcal{Q}}) - \boldsymbol{\mathcal{I}}]\,(\boldsymbol{\varepsilon}'\boldsymbol{\pi} + \boldsymbol{\mathcal{Q}})^{-2}\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}'\,.\end{aligned}$$

The coefficient of autocorrelation for the counting process follows as

$$r_N(k) = \frac{\mathrm{Cov}(N, N_{+k})}{\mathbb{E}\{N^2\} - (\mathbb{E}\{N\})^2}\,.$$

C.4 Distribution of Number of Cells in MMPP Subspace

The problem of computing the distribution of the number of cells per burst can be formulated in the following abstract form: Let $\langle \mathbf{p}, \mathbf{B}\rangle$ be a Phase-type distribution with its probabilities of phase changes in $\mathbf{P}$, i.e. $\mathbf{B} = \mathbf{M}(\mathbf{I} - \mathbf{P})$. While in phase i, Poisson events with rate λ_i are generated. Let $\boldsymbol{\Lambda}$ be the diagonal matrix that contains the Poisson rates λ_i. Call $\boldsymbol{\xi}_j(i)$ the probability that exactly i events are generated during the $\langle \mathbf{e}_j, \mathbf{B}\rangle$ distributed time-interval, where $\mathbf{e}_j$ is a unit vector with all components except for the j-th being zero. Formulas for those probabilities in the special case of constant $\lambda_i \equiv \lambda$ are derived in [Latouche & Ramaswami 99]. However, we need the general case in Chapter 9 since the Poisson rate during a burst can vary due to the throttling in that model.

In the general scenario, the following recursive equation holds for $\boldsymbol{\xi}(0)$:

$$\boldsymbol{\xi}(0) = (\mathbf{M} + \boldsymbol{\Lambda})^{-1}\,\left[\mathbf{B}\boldsymbol{\epsilon}' + \mathbf{M}\mathbf{P}\boldsymbol{\xi}(0)\right]\,.$$

This equation follows from the two possibilities for generating no Poisson events given that the process is in state j: Either the Phase-type distribution ends directly from state j with probability

$$\frac{(\mathbf{B}\boldsymbol{\epsilon}')_j}{\mathbf{M}_{jj} + \lambda_j},$$

or a transition to another state k occurs, which contributes the probability

$$\frac{\mathbf{M}_{jj}}{\mathbf{M}_{jj}+\lambda_j}\sum_k \mathbf{P}_{jk}\boldsymbol{\xi}_k(0)\,.$$

The analog argumentation leads to the recursive equations

$$\boldsymbol{\xi}(n) = (\mathbf{M}+\boldsymbol{\Lambda})^{-1}\,[\boldsymbol{\Lambda}\boldsymbol{\xi}(n-1)+\mathbf{MP}\boldsymbol{\xi}(n)] \quad \text{for} \quad n=1,2,\ldots,$$

for a positive number n of Poisson arrivals.

The two recursive equations together with the entrance vector $\mathbf{p}$ combine into the solution for the scalar probabilities

$$\xi(n) = \mathbf{p}\,\left[(\mathbf{B}+\boldsymbol{\Lambda})^{-1}\boldsymbol{\Lambda}\right]^n\,(\mathbf{B}+\boldsymbol{\Lambda})^{-1}\mathbf{B}\boldsymbol{\epsilon}'\,. \tag{C.2}$$

Thus the distribution of the number of Poisson events has a Matrix-Geometric form with the factor matrix $(\mathbf{B}+\boldsymbol{\Lambda})^{-1}\boldsymbol{\Lambda}$.

If the Poisson rates do not depend on the current state, $\boldsymbol{\Lambda}=\lambda_p\mathbf{I}$, (C.2) simplifies to

$$\xi(n) = \mathbf{p}\left[\lambda_p\mathbf{V}(\mathbf{I}+\lambda_p\mathbf{V})^{-1}\right]^n\,(\mathbf{I}+\lambda_p\mathbf{V})^{-1}\boldsymbol{\varepsilon}'\,.$$

Finally, the PT Range of the distribution of the number of cells during the $\langle\mathbf{p},\mathbf{B}\rangle$ distributed time-interval follows to be

$$\begin{aligned}\operatorname{Rng}(\xi(n)) &= \frac{1}{-\ln\min\{Re\,[\lambda\,(\lambda_p(\mathbf{B}+\lambda_p\mathbf{I})^{-1})]\}} = \frac{1}{-\ln\left[\lambda_p/(1/\operatorname{Rng}(\mathbf{B})+\lambda_p\right]}\\ &= \frac{1}{\ln\left[1+1/(\lambda_p\operatorname{Rng}(\mathbf{B})\right]}\,.\end{aligned}$$

Asymptotically, when $\lambda_p\operatorname{Rng}(\mathbf{B})$ is large, the Taylor approximation $\ln(1+x)\approx x$ can be used:

$$\operatorname{Rng}(\xi(n)) \approx \lambda_p\operatorname{Rng}(\mathbf{B})\;,$$

which justifies the definition of the MBS in Sect. 3.5.

Appendix D

Matrix-Analytic Queueing Models

D.1 Definition & Notation

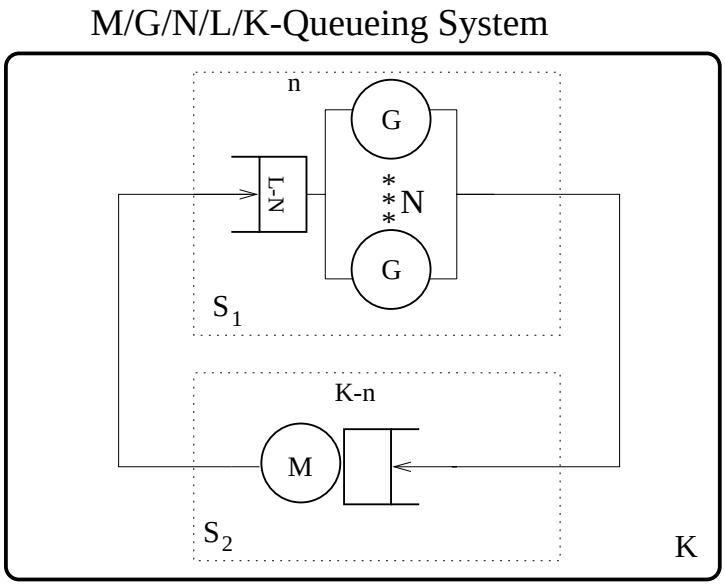

Figure D.1: **Example of a Closed M/G/N/L/K Queueing System:** Subsystem S_1 is the actual considered queue containing N identical servers with an unspecified (general) service-time distribution. S_2 generates the (Poisson-) arrivals to S_1. S_1 by itself is a loss-system; whenever there are L customers at S_1, additionally arriving customers are rejected. Here they go straight back to the queue of S_2 since the number of customers in the whole system is fixed at K.

The N-Burst model uses closed queuing systems of type M/G/N//K to describe the arrivals of bursts and their duration. Figure D.1 pictures an M/G/N/L/K-system. It consists of two subsystems: S_2 contains an exponential server and generates the arrivals to the subsystem S_1 which is the actual queue under consideration.

The so-called *Kendall notation*, $x/y/N/L/K$, describes the system: The first letter, x, describes the type of the server in S_2 that generates the arrivals to S_1. The second component is the type of the server in S_1. N is the number of servers in S_1. L is the maximum number of customer in S_1, i.e. the queue in front of the servers of S_1 can hold at most $L - N$ customers. K is the overall number of customers in the closed system: if n customers are at S_1, $K - n$ are left at S_2. That also implies that customers that are rejected at S_1 due to a full queue will go directly back to the queue of S_2 to keep the overall number of K customers. If any of the numbers K or L is left out, they are assumed to be infinite: If L is infinite, there will never be any rejections at the queue of S_1. If K is infinite, it implies that there is an infinite supply of customer in the queue of the arrival generating stage S_2. The server at S_2 is never idle, thus the arrivals to subsystem

S_1 are then distributed according to the distribution of the arrival-generating server at S_2.

The systems used herein have either $L = \infty$ or $K = \infty$. The case $K = \infty$ is called an open system, because S_2 has a never ending supply of customers in its own queue.

It turns out that the open loss-system M/G/N/K is mathematically equivalent to the closed system M/G/N//K (the former has a finite space for queueing in S_1 while the latter has infinite queueing-space in S_1 but a finite number of customers K). The critical situation is when there are K customers at S_1: In the open loss-system every additional customer gets dropped, while in the finite closed system no additional arrivals can happen, since all customers are at S_1. As soon as one customer at S_1 gets served, new arrivals to S_1 can happen. The difference of the two systems is, that in the M/G/N//K system, the arrival generating server at S_2 starts at the moment when a service at S_1 finishes, while in the M/G/N/K-system, the server at S_2 has already been working for some time. The memorylessness (1.5) of the exponential distribution causes these two scenarios to be equivalent. This also holds for load dependent arrivals in $M_{ld}/G_{ld}/N/[/]K$-systems, when the load-dependence is declared in terms of the load n at subsystem S_1.

A very important parameter of a queueing system is the *utilization*, ρ, which we define as the ratio of the average arrival rate at S_1 by the mean service-rate within S_1:

$$\rho = \frac{1/\,\mathbb{E}\{X_2\}}{N/\,\mathbb{E}\{X_1\}} = \frac{\mathbb{E}\{X_1\}}{N\;\mathbb{E}\{X_2\}},$$

where $\mathbb{E}\{X_1\}$ is the mean service time of one server in S_1 and $\mathbb{E}\{X_2\}$ is the mean time of the arrival-generating server at S_2. For load dependent arrivals or load dependent service times, this definition would have to be modified. For the open system ($K = \infty$), steady-state probabilities for the number of customers at S_1 exist for $\rho < 1$.

The mean queue-length, $\bar{q}$, is the average number of customers in subsystem S_1, which includes the customers in service.

Notation for Arrival Processes (i.e. type of server at S_2)

Abbreviation	Name	Explanation
M	Markovian	exponential
M_{ld}		load dependent exponential
IS	Independent Sources	special case of M_{ld}
GI	General Independent	general renewal process
SM	Semi Markov	correlated arrivals (here mostly MMPP)

Notation for Service Processes (i.e. type of servers at S_1)

Abbreviation	Name	Explanation
M	Markovian	exponential
G	General	general distribution
G_{ld}		load dependent general
G_{ldS}		server specific load dependence
G_{ph}		phase specific load dependence

The last two server-types are described in more detail in Sects. 2.5.2 and 2.5.3. The load dependence is always in terms of the load n in subsystem S_1.

LAQT-techniques are applicable when G is Matrix-Exponential (ME). In that case, the queueing system can be represented as a Markov Process with generator $\boldsymbol{\mathcal{Q}}$ (name is chosen on purpose, since the generator of M/G/N//K-queueing systems will be used as modulating process for MMPPs).

If K or L is finite, $\boldsymbol{\mathcal{Q}}$ is finite and the steady state solution of the system can then be calculated by $\boldsymbol{\pi}\boldsymbol{\mathcal{Q}} = \mathbf{0}$. However, the structure of the $\boldsymbol{\mathcal{Q}}$-matrix allows for more efficient solutions, as summarized in the next few sections.

D.2 $M_{ld}/G_{ld}/1//K$-Systems

The mean queue-length, $\bar{q}$, of the open M/G/1-System is determined solely by the first two moments of the service-time distribution (Pollaczek-Khinchin formula, see e.g. [Kleinrock 75]):

$$\mathbb{E}\{Q\} = \frac{\rho}{1-\rho}\left[1 + \rho \cdot \frac{C^2 - 1}{2}\right] \tag{D.1}$$

More advanced techniques have to be applied to obtain the distribution of the queue-length in a closed $M_{ld}/G_{ld}/1//K$-system The following is a generalization of the derivation of the formulas for the M/ME/1//K-system in [Lipsky 92] to load dependent stages S_1 and S_2. Let $\mu(n)$ be the load-dependent arrival rate. Also, let the service-time distribution of the single server at S_1 be represented by $< \mathbf{p}, \mathbf{B} >$. That distribution also depends on the load n in S_1 through the factors $\beta(n)$:

$$\mathbf{B}_n := \beta(n)\,\mathbf{B}, \quad n = 1, ..., K$$

The stationary distribution vector $\boldsymbol{\pi}$ of the system can be partitioned into components $\boldsymbol{\pi}(i)$, $i = 0, ..., K$ for the load i at S_1.

The balance equations of the system contain the following matrices:

$$\mathbf{A_i} := \mathbf{I} + \frac{1}{\mu(i)}\mathbf{B}_i - \boldsymbol{\varepsilon}'\mathbf{p}, \quad \mathbf{U_i} := \mathbf{A_i}^{-1}, \quad i = 1, \ldots, K-1,$$

which lead to the solution

$$\begin{aligned}
\boldsymbol{\pi}(0) &= c_0\mathbf{p} \\
\boldsymbol{\pi}(n) &= \boldsymbol{\pi}(0)\frac{\mu(0)}{\mu(n)}\,\mathbf{U_1}\cdot\ldots\mathbf{U_n}, \quad \text{for} \quad n = 1, \ldots, K-1. \\
\boldsymbol{\pi}(K) &= \boldsymbol{\pi}(0)\lambda_0\,\mathbf{U_1}\cdot\ldots\mathbf{U_{K-1}}\left[\mathbf{B}_K\right]^{-1}.
\end{aligned}$$

c_0 is derived from $\sum_{i=1}^{K}\boldsymbol{\pi}(i)\boldsymbol{\varepsilon}' = 1$.

D.3 $M_{ld}/G_{ld}/N//N$-Systems

The special case of an $M_{ld}/G_{ld}/N//N$-Systems deserves extra treatment, since there exist more efficient solutions. As mentioned in [Lipsky 92], due to the fact that there is no queueing at the N-server unit S_1, the ME-representation of the servers fulfill the conditions of a Jackson-network. Therefore, the mean queue-length at S_1 does not depend on the service-time distribution any more, but only on its mean (and on the load-dependence factors).

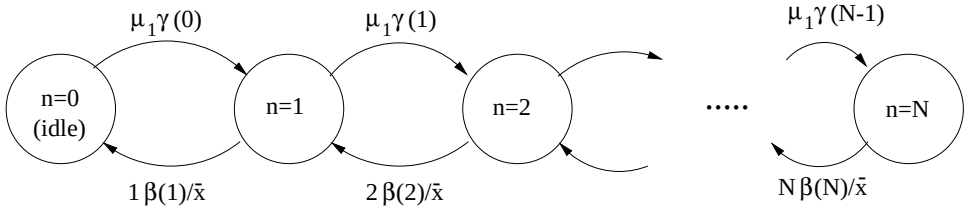

Figure D.2: **Simplification for $M_{ld}/G_{ld}/N//N$ Queue:** The stationary queue-length distribution of an $M_{ld}/G_{ld}/N//N$ system only depends on the mean $\bar{x}$ of the service time-distribution and on the load-dependence factors, not on the service-time distribution itself (see [Lipsky 92]). Therefore, the analytical treatment reduces to a standard birth-death process with the transition-rate diagram shown in this figure.

Thus, the general server units can be replaced by exponential servers with the same mean and the underlying Markov process reduces to a standard birth-death process as shown in Figure D.2.

The stationary queue-length distribution then is:

$$\pi(n) = \pi(0) \cdot \frac{\mu_1^n \prod_{i=0}^{n-1} \gamma(i)}{n! \prod_{i=1}^{n} \beta(i)} \cdot \bar{x}^n.$$

Again, $\pi(0)$ is derived by normalization: $\sum_i \pi(i) = 1$.

Remember that this formula does not hold if there is queueing outside of the N-server unit ($K > N$), or if the load-dependence factors depend on the individual phase of the ME distribution (Sect. 2.5.2), or if different load-dependence factors exist for the $N > 1$ servers in S_1 (see Sect. 2.5.3).

D.4 $GI/M/1$-Systems

M/G/$N//K$ queueing systems are used in the N-Burst model to describe the active bursts on the line, i.e. they represent the modulating process that determines the arrival rate at the switch. The actual performance analysis of network components (e.g. ATM switches) is done by a second queue of type MMPP/M/1 which uses the MMPP as arrival process[1]. The distinctive feature of an SM process is that it can generate correlated arrivals: Not only is the distribution of the inter-cell times important, but their order has to be taken into account as well.

However, to be able to compare the correlated arrivals of the SM/M/1-Queue (described by $< \boldsymbol{\mathcal{B}}, \boldsymbol{\mathcal{L}} >$) with independent arrivals with the same marginal distribution (described by $< \mathbf{p}, \mathbf{B} > = < \boldsymbol{\wp}, \boldsymbol{\mathcal{B}} >$) at an GI/M/1-queue, a short review of the GI/M/1-queue following [Lipsky 92] is given:

Let the exponential server have rate ν and let the arrival process be a renewal process defined by $< \mathbf{p}, \mathbf{B} >$ with mean $\bar{x} := \mathbf{pV}\boldsymbol{\varepsilon}'$ where $\mathbf{V} = \mathbf{B}^{-1}$.

The utilization of the GI/M/1 queue is

$$\rho = \frac{1}{\bar{x}\nu},$$

[1] Do NOT mix up the SM/M/1-queue that models e.g. a switch (see also Figure C.1) with the embedded M/G/$N//K$ queue that describes the burst-arrivals!

and the scalar[2] queue-length probabilities, $\hat{\pi}(k)$, for length k turn out to be:

$$\begin{aligned} \hat{\pi}(0) &= 1-\rho, \\ \hat{\pi}(k) &= (1-s)\rho s^{k-1} \quad \text{for} \quad k = 1,2,3,\dots, \end{aligned} \tag{D.2}$$

(D.3)

where s is the smallest positive eigenvalue of the matrix $\mathbf{A}$:

$$s = \min_i \left\{ \lambda_i(\mathbf{A}) \mid \lambda_i(\mathbf{A}) > 0 \right\}, \quad \mathbf{A} := \mathbf{I} + \frac{1}{\nu}\mathbf{B} - \boldsymbol{\varepsilon}'\mathbf{p}.$$

The mean queue-length of the GI/M/1-queue follows as:

$$\mathbb{E}\{Q\} = \sum_{k=1}^{\infty} k \cdot \hat{\pi}(k) = \frac{\rho}{1-s} \cdot \tag{D.4}$$

The mean system time (waiting time + service time) follows from Little's Law:

$$\mathbb{E}\{T\} = \bar{x}\,\mathbb{E}\{Q\} = \frac{1}{\nu(1-s)} \cdot$$

For an exponential arrival process (M/M/1-queue), $s = \rho$.

Since the arrival process is non-exponential, it makes a difference whether we look at the queue-length probabilities at randomly chosen points in time or at arrival times. The latter is important for determining buffer overflow probabilities for non-loss systems. If there is either a second backup buffer (assumed to be infinitely large, e.g. a harddisk) or assuming a feedback mechanism that signals the full primary buffer to the sender and thus delays the new packet (i.e. the backup buffer is spread out among the transmission sources), the probability that a newly arriving packet goes to the backup-buffer is:

$$\mathrm{BOP}(B) = \sum_{k=B}^{\infty} a(k),$$

where B is the size of the primary buffer and $a(k)$ is the probability that the queue-length at arrival times is k. It can be shown that (see [Lipsky 92], p. 219)

$$a(k) = (1-s)s^k. \tag{D.5}$$

$$\Longrightarrow \quad \mathrm{BOP}(B) = (1-s)\sum_{k=B}^{\infty} s^k = s^B.$$

D.5 $SM/M/1$-Systems

The SM/M/1-queue is treated in [Neuts 81] as a so-called *Quasi-Birth-Death Process*. That is a Markov process where the transition rate matrix is block-tridiagonal. The matrix notation of [Neuts 81] is slightly different but is changed here to make it consistent with the previous

[2] The state-space of the GI/M/1-queue is the product space of the possible queue-lengths $k = 0,1,2,\dots$ and the possible states of the arrival process. The scalar queue-length probabilities, $\hat{\pi}(k) := \hat{\boldsymbol{\pi}}(k)\boldsymbol{\varepsilon}'$, add up the probabilities of all the states with queue-length k.

formulas. The state space of the SM/M/1-queue is the product of the state space of the arrival process and the set of possible queue-lengths. The block-tridiagonal transition rate matrix (the generator) $\hat{\mathbf{Q}}$ for the SM/M/1-queue is the infinite matrix:

$$\hat{\mathbf{Q}} = \begin{bmatrix} \overline{\boldsymbol{\mathcal{A}}_1} & \boldsymbol{\mathcal{A}}_0 & \mathbf{O} & \dots \\ \boldsymbol{\mathcal{A}}_2 & \boldsymbol{\mathcal{A}}_1 & \boldsymbol{\mathcal{A}}_0 & \mathbf{O} \\ \mathbf{O} & \boldsymbol{\mathcal{A}}_2 & \boldsymbol{\mathcal{A}}_1 & \ddots \\ \mathbf{O} & \mathbf{O} & \ddots & \ddots \end{bmatrix}.$$

Let the SM arrival process be defined by $< \boldsymbol{\mathcal{B}}, \boldsymbol{\mathcal{L}} >$, then $\boldsymbol{\mathcal{Q}}$ has the following form:

$$\hat{\mathbf{Q}} = \begin{bmatrix} -\boldsymbol{\mathcal{B}} & \boldsymbol{\mathcal{L}} & \mathbf{0} & \mathbf{0} & \cdots \\ \nu\boldsymbol{\mathcal{I}} & -(\boldsymbol{\mathcal{B}} + \nu\boldsymbol{\mathcal{I}}) & \boldsymbol{\mathcal{L}} & \mathbf{0} & \cdots \\ \mathbf{0} & \nu\boldsymbol{\mathcal{I}} & -(\boldsymbol{\mathcal{B}} + \nu\boldsymbol{\mathcal{I}}) & \boldsymbol{\mathcal{L}} & \cdots \\ \mathbf{0} & \mathbf{0} & \nu\boldsymbol{\mathcal{I}} & -(\boldsymbol{\mathcal{B}} + \nu\boldsymbol{\mathcal{I}}) & \cdots \\ \vdots & \vdots & \ddots & \ddots & \ddots \end{bmatrix}.$$

The leftmost column of matrices represents the transitions to states with queue-length 0, the next column for transition to states with queue-length 1, etc. . The blocks on the main diagonal describe state transitions within the arrival process without a change of the queue-length. The upper diagonal ($\boldsymbol{\mathcal{L}}$) describes the arrivals while the lower diagonal ($\nu\boldsymbol{\mathcal{I}}$) represents the service of a cell (without changing the state of the arrival process).

The stationary queue-length distribution, $\hat{\boldsymbol{\pi}}$, can be partitioned into clusters of states with same queue-length:

$$\hat{\boldsymbol{\pi}} = [\hat{\boldsymbol{\pi}}(0), \hat{\boldsymbol{\pi}}(1), \hat{\boldsymbol{\pi}}(2), ...] .$$

The equation $\hat{\boldsymbol{\pi}}\hat{\mathbf{Q}} = \mathbf{0}$ then has a matrix-geometric solution:

$$\hat{\boldsymbol{\pi}}(k) = \hat{\boldsymbol{\pi}}(0)\boldsymbol{\mathcal{R}}^k,$$

where the rate-matrix $\boldsymbol{\mathcal{R}}$ has to fulfill the following quadratic matrix equation,

$$\boldsymbol{\mathcal{A}}_0 + \boldsymbol{\mathcal{R}}\boldsymbol{\mathcal{A}}_1 + \boldsymbol{\mathcal{R}}^2\boldsymbol{\mathcal{A}}_2 = \mathbf{O},$$

for $\boldsymbol{\mathcal{A}}_0 = \boldsymbol{\mathcal{L}}$, $\boldsymbol{\mathcal{A}}_1 = -(\boldsymbol{\mathcal{B}} + \nu\boldsymbol{\mathcal{I}})$, $\boldsymbol{\mathcal{A}}_2 = \nu\boldsymbol{\mathcal{I}}$.

Interpretation of the elements of $\boldsymbol{\mathcal{R}}$: $\boldsymbol{\mathcal{R}}_{jk}$ is the expected time spent in state (queue-length, state)$= (i+1, k)$ before first return to queue-length i, given that the process started at (i, j), in units of $(\boldsymbol{\mathcal{B}}_{jj} + \nu)^{-1}$, see [Latouche & Ramaswami 99]. That value is independent of i.

The boundary-condition is:

$$\hat{\boldsymbol{\pi}}(0)\left(\overline{\boldsymbol{\mathcal{A}}_1} + \boldsymbol{\mathcal{R}}\boldsymbol{\mathcal{A}}_2\right) = \mathbf{0}, \quad \text{here}: \quad \hat{\boldsymbol{\pi}}(0)\left(-\boldsymbol{\mathcal{B}} + \nu\boldsymbol{\mathcal{R}}\right) = \mathbf{0}.$$

Also, $\hat{\boldsymbol{\pi}}$ is a probability vector, thus:

$$1 = \sum_{i=0}^{\infty} \hat{\boldsymbol{\pi}}(i)\boldsymbol{\varepsilon}' = \sum_{i=0}^{\infty} \hat{\boldsymbol{\pi}}(0)\boldsymbol{\mathcal{R}}^i\boldsymbol{\varepsilon}' = \hat{\boldsymbol{\pi}}(0)\left(\boldsymbol{\mathcal{I}} - \boldsymbol{\mathcal{R}}\right)^{-1}\boldsymbol{\varepsilon}'.$$

These two conditions on $\hat{\boldsymbol{\pi}}(0)$ are hold for:

$$\hat{\boldsymbol{\pi}}(0) = \boldsymbol{\pi}\left(\boldsymbol{\mathcal{I}} - \boldsymbol{\mathcal{R}}\right),$$

where $\boldsymbol{\pi} = \boldsymbol{\wp}\boldsymbol{\mathcal{V}} / (\boldsymbol{\wp}\boldsymbol{\mathcal{V}}\boldsymbol{\varepsilon}')$. Note that in case of an MMPP arrival process, $\boldsymbol{\pi}$ is the steady state vector of the modulating process of the MMPP (see Section C.2).

The scalar steady-state probabilities for the queue-length are then obtained by

$$\hat{\boldsymbol{\pi}}(k)\,\boldsymbol{\varepsilon}' = \boldsymbol{\pi}(\boldsymbol{\mathcal{I}} - \boldsymbol{\mathcal{R}})\boldsymbol{\mathcal{R}}^k\boldsymbol{\varepsilon}', \quad k = 0, 1, \ldots \,. \tag{D.6}$$

The matrix-geometric series in the mean queue-length formula can be simplified when $\rho < 1$:

$$\mathbb{E}\{Q\} = \sum_{k=0}^{\infty} k\,\hat{\boldsymbol{\pi}}(k)\,\boldsymbol{\varepsilon}' = \boldsymbol{\pi}\boldsymbol{\mathcal{R}}(\boldsymbol{\mathcal{I}} - \boldsymbol{\mathcal{R}})^{-1}\boldsymbol{\varepsilon}'. \tag{D.7}$$

The vector part of (D.6), $\boldsymbol{\pi}(\boldsymbol{\mathcal{I}} - \boldsymbol{\mathcal{R}})\boldsymbol{\mathcal{R}}^k$, has as its components the probabilities that the queue-length is k and the arrival process is in the corresponding state. If we look at the queue at arrival times, the states of the SM-process with a high departure-rate[3] contribute more observation points. Therefore, the probability vector has to be scaled by the $\boldsymbol{\mathcal{L}}$-matrix and then normalized to give the probabilities at arrival times, see [Fiorini 97]. Adding up all the components results in:

$$a(k) = \frac{\boldsymbol{\pi}(\boldsymbol{\mathcal{I}} - \boldsymbol{\mathcal{R}})\boldsymbol{\mathcal{R}}^k\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}'}{\boldsymbol{\pi}\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}'} \,. \tag{D.8}$$

The queue-length at cell-arrival times, $Q^{(a)}$, is closely related to the cell-delay. The mean cell-delay and the cell-delay variation are very important performance parameters. Therefore, the first two moments of the queue-length distribution at cell-arrival times need to be studied:

$$\mathbb{E}\left\{Q^{(a)}\right\} = \frac{\boldsymbol{\pi}\boldsymbol{\mathcal{R}}(\boldsymbol{\mathcal{I}} - \boldsymbol{\mathcal{R}})^{-1}\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}'}{\boldsymbol{\pi}\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}'} \,. \tag{D.9}$$

$$\mathbb{E}\left\{\left[Q^{(a)}\right]^2\right\} = \frac{\boldsymbol{\pi}\boldsymbol{\mathcal{R}}(\boldsymbol{\mathcal{R}} + \boldsymbol{\mathcal{I}})(\boldsymbol{\mathcal{I}} - \boldsymbol{\mathcal{R}})^{-2}\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}'}{\boldsymbol{\pi}\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}'} \,. \tag{D.10}$$

Finally, the overflow probability for a buffer of size B comes out as:

$$\mathrm{BOP}(B) = \sum_{k=B}^{\infty} a(k) = \boldsymbol{\pi}\boldsymbol{\mathcal{R}}^B\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}' \tag{D.11}$$

The advantage of the matrix-geometric solution is that it is not necessary to do any calculations in the state-space described by $\hat{\mathbf{Q}}$. Independently of the buffer-size, the calculations can always be done with matrices of the dimension of the arrival process. Therefore, calculation of buffer-overflow probabilities can be done for arbitrary large buffer-sizes. That also holds for the SM/M/1/B loss-system (see Appendix D.6 and [Krieger et al. 98]).

Finally, it is fairly easy to generalize the service time from exponential to any other matrix-exponential distribution, e.g. Erlangian-M for ATM-switches. The only changes occur in the block-matrices $\boldsymbol{\mathcal{A}}_i$. However, their dimension increases by a factor of M.

[3] The process itself is an *arrival process* when fed into a queue. However, we talk of *departures* from the SM-process. Thereby we look at the phases of the sub-system that build up the arrival process. A departure from one of the phases (described by the elements of $\boldsymbol{\mathcal{L}}$) then corresponds to an arrival at the exponential server of the SM/M/1-queue.

D.6 SM/M/1/B Loss Models

For the SM/M/1/B loss model, the generator matrix $\hat{\mathbf{Q}}$ is finite with dimension $(B+1)\cdot\dim(\boldsymbol{\mathcal{B}})$:

$$\hat{\mathbf{Q}} = \begin{bmatrix} \overline{\boldsymbol{\mathcal{A}}_1} & \boldsymbol{\mathcal{A}}_0 & \mathbf{O} & \dots & \\ \boldsymbol{\mathcal{A}}_2 & \boldsymbol{\mathcal{A}}_1 & \boldsymbol{\mathcal{A}}_0 & \mathbf{O} & \dots \\ \mathbf{O} & \ddots & \ddots & \ddots & \mathbf{O} \\ & \mathbf{O} & \boldsymbol{\mathcal{A}}_2 & \boldsymbol{\mathcal{A}}_1 & \boldsymbol{\mathcal{A}}_0 \\ & & & \boldsymbol{\mathcal{A}}_2 & \underline{\boldsymbol{\mathcal{A}}_1} \end{bmatrix}.$$

$$\begin{aligned} \text{with} \quad & \boldsymbol{\mathcal{A}}_0 = \boldsymbol{\mathcal{L}}, \quad \boldsymbol{\mathcal{A}}_1 = -\boldsymbol{\mathcal{B}} - \nu\boldsymbol{\mathcal{I}}, \quad \boldsymbol{\mathcal{A}}_2 = \nu\boldsymbol{\mathcal{I}}, \\ \text{and} \quad & \overline{\boldsymbol{\mathcal{A}}_1} = \boldsymbol{\mathcal{A}}_1 + \boldsymbol{\mathcal{A}}_2 = -\boldsymbol{\mathcal{B}}, \quad \underline{\boldsymbol{\mathcal{A}}_1} = \boldsymbol{\mathcal{A}}_0 + \boldsymbol{\mathcal{A}}_1 = \boldsymbol{\mathcal{L}} - \boldsymbol{\mathcal{B}} - \nu\boldsymbol{\mathcal{I}} \end{aligned}$$

According to [Krieger et al. 98], the following mixed Matrix-Geometric solution for the steady-state queue-length probabilities exist:

$$\hat{\boldsymbol{\pi}}(i) = \mathbf{a}\boldsymbol{\mathcal{R}}^i + \mathbf{b}\boldsymbol{\mathcal{S}}^{B-i}, \quad i = 0, ..., B, \tag{D.12}$$

where the two Matrix-Geometric factors are the minimal solutions of the following quadratic matrix equations:

$$\begin{aligned} \boldsymbol{\mathcal{A}}_0 + \boldsymbol{\mathcal{R}}\boldsymbol{\mathcal{A}}_1 + \boldsymbol{\mathcal{R}}^2\boldsymbol{\mathcal{A}}_2 &= \mathbf{O}, \\ \boldsymbol{\mathcal{S}}^2\boldsymbol{\mathcal{A}}_0 + \boldsymbol{\mathcal{S}}\boldsymbol{\mathcal{A}}_1 + \boldsymbol{\mathcal{A}}_2 &= \mathbf{O}. \end{aligned}$$

Both matrix equations can be solved simultaneously with the spectral decomposition method, see Sect. F.1. For the case $\rho < 1$ (i.e. $\boldsymbol{\pi}\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}' < \nu$), the absolute value of all eigenvalues of $\boldsymbol{\mathcal{R}}$ turn out to be smaller than one, thus the inverse $(\boldsymbol{\mathcal{I}} - \boldsymbol{\mathcal{R}})^{-1}$ exists. $\boldsymbol{\mathcal{S}}$ has exactly one eigenvalue $\lambda_{\boldsymbol{\mathcal{S}},1} = 1$. Therefore, $(\boldsymbol{\mathcal{I}} - \boldsymbol{\mathcal{S}})$ is singular, but we will show below that sums of the form $\mathbf{b}\sum_k \boldsymbol{\mathcal{S}}^k$ can still be computed efficiently.

The coefficient vectors $\mathbf{a}$ and $\mathbf{b}$ in (D.12) can be derived from the boundary equations in $\hat{\boldsymbol{\pi}}\hat{\mathbf{Q}} = \mathbf{0}$,

$$[\mathbf{a} \,|\, \mathbf{b}] \left[\begin{array}{c|c} \boldsymbol{\mathcal{A}}_1 + \boldsymbol{\mathcal{A}}_2 + \boldsymbol{\mathcal{R}}\boldsymbol{\mathcal{A}}_2 & \boldsymbol{\mathcal{R}}^B\,(\boldsymbol{\mathcal{A}}_0 - \boldsymbol{\mathcal{R}}\boldsymbol{\mathcal{A}}_2) \\ \hline \boldsymbol{\mathcal{S}}^B\,(\boldsymbol{\mathcal{A}}_2 - \boldsymbol{\mathcal{S}}\boldsymbol{\mathcal{A}}_0) & \boldsymbol{\mathcal{S}}\boldsymbol{\mathcal{A}}_0 + \boldsymbol{\mathcal{A}}_1 + \boldsymbol{\mathcal{A}}_0 \end{array}\right] = \mathbf{0}, \tag{D.13}$$

and from the normalization of $\hat{\boldsymbol{\pi}}$:

$$\sum_{i=0}^{B} \hat{\boldsymbol{\pi}}(i)\,\boldsymbol{\varepsilon}' = 1 \quad \Longleftrightarrow \quad \mathbf{a}(\boldsymbol{\mathcal{I}} - \boldsymbol{\mathcal{R}})^{-1}\left(\boldsymbol{\mathcal{I}} - \boldsymbol{\mathcal{R}}^{B+1}\right) + \mathbf{b}(\boldsymbol{\mathcal{I}} - \tilde{\boldsymbol{\mathcal{S}}})^{-1}\left(\boldsymbol{\mathcal{I}} - \boldsymbol{\mathcal{S}}^{B+1}\right) = 1.$$

The definition of $\tilde{\boldsymbol{\mathcal{S}}}$ requires some elaboration:

Singularity of $(\boldsymbol{\mathcal{I}} - \boldsymbol{\mathcal{S}})$: Exactly one eigenvalue of $\boldsymbol{\mathcal{S}}$ is equal to one, $\lambda_{\boldsymbol{\mathcal{S}},1} = 1$. The corresponding left eigenvector is $\mathbf{u}_{\boldsymbol{\mathcal{S}},1} = \boldsymbol{\pi}$ which follows from

$$\begin{aligned} & \boldsymbol{\pi}\,(\boldsymbol{\mathcal{A}}_0 + \boldsymbol{\mathcal{A}}_1 + \boldsymbol{\mathcal{A}}_2) = \boldsymbol{\pi}(\boldsymbol{\mathcal{L}} - \boldsymbol{\mathcal{B}}) = \mathbf{0}, \\ \text{and} \quad & (\boldsymbol{\mathcal{I}} - \boldsymbol{\mathcal{S}})\,(\boldsymbol{\mathcal{A}}_0 + \boldsymbol{\mathcal{A}}_1 + \boldsymbol{\mathcal{S}}\boldsymbol{\mathcal{A}}_0) = \boldsymbol{\mathcal{A}}_0 + \boldsymbol{\mathcal{A}}_1 + \boldsymbol{\mathcal{A}}_2\,. \end{aligned}$$

The last equation together with $(\boldsymbol{\mathcal{A}}_0 + \boldsymbol{\mathcal{A}}_1 + \boldsymbol{\mathcal{A}}_2)\,\boldsymbol{\varepsilon}' = \mathbf{0}$, combines into

$$(\boldsymbol{\mathcal{I}} - \boldsymbol{\mathcal{S}})\,\left(\boldsymbol{\mathcal{S}}\boldsymbol{\mathcal{A}}_0\boldsymbol{\varepsilon}' - \boldsymbol{\mathcal{A}}_2\boldsymbol{\varepsilon}'\right) = \mathbf{0},$$

which in turn proves that the normalized right eigenvector for eigenvalue $\lambda_{\boldsymbol{\mathcal{S}},1} = 1$ is

$$\mathbf{v}'_{\boldsymbol{\mathcal{S}},1} = \frac{\boldsymbol{\mathcal{A}}_2\boldsymbol{\varepsilon}' - \boldsymbol{\mathcal{S}}\boldsymbol{\mathcal{A}}_0\boldsymbol{\varepsilon}'}{\nu(1-\rho)}\,.$$

The second set of equations in (D.13) together with $\boldsymbol{\mathcal{A}}_0\boldsymbol{\varepsilon}' - \boldsymbol{\mathcal{R}}\boldsymbol{\mathcal{A}}_2\boldsymbol{\varepsilon}' = \mathbf{0}$ proves that

$$\mathbf{b}\mathbf{v}'_{\boldsymbol{\mathcal{S}},1} = 0,$$

i.e. the right eigenvector for eigenvalue $\lambda_{\boldsymbol{\mathcal{S}},1} = 1$ of $\boldsymbol{\mathcal{S}}$ is orthogonal to the coefficient vector $\mathbf{b}$ of the mixed Matrix-Geometric solution (D.12).

Consequently, the matrix-geometric sum with the coefficient vector $\mathbf{b}$ and the factor matrix $\boldsymbol{\mathcal{S}}$ can be simplified:

$$\mathbf{b}\sum_{k=0}^{B}\boldsymbol{\mathcal{S}}^k = \mathbf{b}(\boldsymbol{\mathcal{I}} - \tilde{\boldsymbol{\mathcal{S}}})^{-1}\,\left(\boldsymbol{\mathcal{I}} - \boldsymbol{\mathcal{S}}^{B+1}\right),$$

where $\tilde{\boldsymbol{\mathcal{S}}} = \boldsymbol{\mathcal{S}} - \mathbf{v}'_{\boldsymbol{\mathcal{S}},1}\boldsymbol{\pi}$ is in principle the matrix $\boldsymbol{\mathcal{S}}$, but with the contribution of its eigenvalue $\lambda_{\boldsymbol{\mathcal{S}},1} = 1$ removed. This is possible in the summation above, since the right eigenvector $\mathbf{v}'_{\boldsymbol{\mathcal{S}},1}$ is orthogonal to the vector $\mathbf{b}$.

Having obtained the steady-state queue-length probabilities $\hat{\boldsymbol{\pi}}(i)$, the mean queue length of the SM/M/1/B model can be computed by:

$$\begin{aligned}\mathbb{E}\{Q\} = \sum_{i=1}^{B} i\,\hat{\boldsymbol{\pi}}(i)\boldsymbol{\varepsilon}' \;\; = \;\; & \mathbf{a}(\boldsymbol{\mathcal{I}} - \boldsymbol{\mathcal{R}})^{-1}\left[(\boldsymbol{\mathcal{I}} - \boldsymbol{\mathcal{R}})^{-1}\,(\boldsymbol{\mathcal{I}} - \boldsymbol{\mathcal{R}}^{B+2}) - (B+1)\,\boldsymbol{\mathcal{R}}^{B+1} - \boldsymbol{\mathcal{I}}\right]\boldsymbol{\varepsilon}' \\ & -\mathbf{b}(\boldsymbol{\mathcal{I}} - \tilde{\boldsymbol{\mathcal{S}}})^{-1}\,\left[(\boldsymbol{\mathcal{I}} - \tilde{\boldsymbol{\mathcal{S}}})^{-1}\,(\boldsymbol{\mathcal{I}} - \boldsymbol{\mathcal{S}}^{B+1}) - (B+1)\,\boldsymbol{\mathcal{I}}\right]\boldsymbol{\varepsilon}'.\end{aligned}$$

D.7 PT Range of Matrix-Geometric Queue-Length Distributions

Since the queue-length distribution of an N-Burst/M/1 queue with a Matrix-Exponential burst-length distribution has a matrix-geometric representation (see [Neuts 81]),

$$\boldsymbol{\pi}_k = \boldsymbol{\pi}_0\,\mathbf{R}^k\,,$$

the scalar queue-length probabilities (both at random observation points and at arrival instances) eventually decay geometrically as

$$\pi_k \sim \eta^k = \exp\left[-\ln\left(1/\eta\right)k\right],$$

where η is the eigenvalue with maximal absolute value of the matrix $\mathbf{R}$, see also [Neuts 86]. Consequently, the decay-rate of the queue-length distribution is $\ln(1/\eta)$, thus the PT Range of the queue-length probabilities can be computed as

$$\overline{q}[N] = \frac{1}{\ln(1/\eta)}\,.$$

Appendix E

Efficient Matrix Representations of N-Burst Models

E.1 Matrix Representation of N-Burst Models

The N-Burst traffic models are described in detail in Sects. 2.3 to 2.5. Those models have an MMPP representation, $< \boldsymbol{\mathcal{B}}, \boldsymbol{\mathcal{L}} >$ or alternatively $< \boldsymbol{\mathcal{Q}}, \boldsymbol{\mathcal{L}} >$ as it is introduced in Appendix C.2. This section provides the detailed matrix representations together with derived closed-form solutions for the steady-state vector $\boldsymbol{\pi}$ of the modulator $\boldsymbol{\mathcal{Q}}$, if available.

E.1.1 1-Burst Models

A single source ($N = 1$) ON/OFF process with exponential OFF periods with mean Z and general Matrix-Exponential ON periods with representation $\langle \mathbf{p}, \mathbf{B} \rangle$ (see Appendix B.1) has the MMPP representation:

$$\boldsymbol{\mathcal{Q}}_1 = \left[\begin{array}{c|c} -1/Z & 1/Z\,\mathbf{p} \\ \hline \mathbf{B}\boldsymbol{\varepsilon}' & -\mathbf{B} \end{array}\right], \quad \boldsymbol{\mathcal{L}}_1 = \left[\begin{array}{c|c} 0 & \\ \hline & \lambda_p \mathbf{I} \end{array}\right] + \lambda_0 \boldsymbol{\mathcal{I}} \cdot$$

The stationary distribution vector of the modulating process

$$\boldsymbol{\pi}_1 = \frac{1}{1 + \bar{x}/Z} \left[1,\ \frac{1}{Z}\,\mathbf{pV}\right], \quad \text{where} \quad \bar{x} := \mathbb{E}\left\{\langle \mathbf{p}, \mathbf{B} \rangle\right\} = \mathbf{pV}\boldsymbol{\varepsilon}',$$

satisfies the equation $\boldsymbol{\pi}_1 \boldsymbol{\mathcal{Q}}_1 = 0$ for the [M/G/1/1] model.

If we allow for waiting bursts in the N-Burst [M/G/1//K] model, the state-spaces has to be

extended to keep track of the number of waiting bursts:

$$\boldsymbol{\mathcal{Q}} = \begin{bmatrix} -\mu & \mu\mathbf{p} & \mathbf{0} & & & \\ \mathbf{B}\varepsilon' & -\mathbf{B}-\mu\mathbf{I} & \mu\mathbf{I} & \mathbf{O} & & \\ \mathbf{O} & \mathbf{B}\varepsilon'\mathbf{p} & -\mathbf{B}-\mu\mathbf{I} & \mu\mathbf{I} & \mathbf{O} & \\ & & \ddots & \ddots & \ddots & \\ & & \mathbf{O} & \mathbf{B}\varepsilon'\mathbf{p} & -\mathbf{B}-\mu\mathbf{I} & \mu\mathbf{I} \\ & & & \mathbf{O} & \mathbf{B}\varepsilon'\mathbf{p} & -\mathbf{B} \end{bmatrix}$$

Here, μ is the arrival-rate of bursts. The first row of $\boldsymbol{\mathcal{Q}}$ contains the transition rates from the state with no active bursts. The next K rows of blocks contain the rates for transitions from the sets of states with $i = 1$,..., $i = K$ active respectively waiting bursts.

The cell-rates of the 1-Burst model is zero for the OFF-State and λ_p for the ON-States. In addition, the model allows for smooth (i.e. Poisson) background traffic with rate λ_0. Hence, $\boldsymbol{\mathcal{L}}$ is the following matrix:

$$\boldsymbol{\mathcal{L}} = \begin{bmatrix} 0 & & & \\ & \lambda_p\mathbf{I} & & \\ & & \ddots & \\ & & & \lambda_p\mathbf{I} \end{bmatrix} + \lambda_0\,\boldsymbol{\mathcal{I}}.$$

There is no closed-form expression for $\boldsymbol{\pi}$ when $K > 1$. However, the Matrix-Geometric formulas in Section D.2 (with $\gamma \equiv 1$ and $\beta \equiv 1$ for the above scenario) provide an efficient method to compute $\boldsymbol{\pi}$.

E.1.2 N-Burst Models

For multiple, simultaneously active bursts, the current state of each of the up to N active servers in the modulating M/G/N//K queue has to be reflected in the state-space. An elegant formulation therefore is the use of *Kronecker-Products* and *Kronecker-Sums* of matrices , that

are defined as (see [GRAHAM 81]):

$$\mathbf{A} \otimes \mathbf{B} := \begin{bmatrix} a_{11}\mathbf{B} & \dots & a_{1n}\mathbf{B} \\ \vdots & & \vdots \\ a_{m1}\mathbf{B} & \dots & a_{mn}\mathbf{B} \end{bmatrix}, \quad \mathbf{A}^{\otimes n} := \mathbf{A} \otimes \cdots \otimes \mathbf{A}.$$

$$\mathbf{A} \oplus \mathbf{B} := \mathbf{A} \otimes \mathbf{I}_{\dim(B)} + \mathbf{I}_{\dim(A)} \otimes \mathbf{B}, \quad \mathbf{A}^{\oplus n} := \mathbf{A} \oplus \cdots \oplus \mathbf{A}.$$

The notation $\mathbf{A}^{\otimes i}$ expresses the Kronecker product with i factors $\mathbf{A}$. Analogously, $\mathbf{A}^{\oplus i}$ stands for the Kronecker sum with i summands.

The aggregation of the traffic from N identical ON/OFF sources in the IS model could be represented using N Kronecker sums of $\boldsymbol{\mathcal{Q}}_1$ respectively $\boldsymbol{\mathcal{L}}_1$. However, this is no longer possible for more general N-Burst models, such as the [M] model, see Sect. 2.4. Instead, a representation of $\boldsymbol{\mathcal{Q}}$ with a Quasi-Birth-Death structure is advantageous: the levels are defined by the number of active sources. Within in the block-matrices that represent states with more than one active source, Kronecker products have to be used. First, lets illustrate the simplest case, the 2-Burst [M/G/2//2] model in which μ is the rate of burst-starts.

[M/G/2//2] Model:

$$\boldsymbol{\mathcal{Q}} = \left[\begin{array}{c|c|c} -\mu & \mu\mathbf{p} & \mathbf{0} \\ \hline \mathbf{B}\boldsymbol{\varepsilon}' & -\mathbf{B} - \mu\mathbf{I} & \mu\,\mathbf{I} \otimes \mathbf{p} \\ \hline \mathbf{0} & \begin{array}{c}(\mathbf{B}\boldsymbol{\varepsilon}') \otimes \mathbf{I} \\ \mathbf{I} \otimes (\mathbf{B}\boldsymbol{\varepsilon}')\end{array} & -\mathbf{I} \otimes \mathbf{B} - \mathbf{B} \otimes \mathbf{I} \end{array}\right],$$

$$\boldsymbol{\mathcal{L}} = \left[\begin{array}{c|c|c} 0 & & \\ \hline & \mathbf{L_0} & \\ \hline & & \mathbf{L_0} \oplus \mathbf{L_0} \end{array}\right] + \lambda_0\,\boldsymbol{\mathcal{I}}\,.$$

Thereby, $\mathbf{L_0} := \lambda_p \mathbf{I}$.

The stationary distribution vector $\boldsymbol{\pi}$ with $\boldsymbol{\pi}\boldsymbol{\mathcal{Q}} = \mathbf{0}$ and $\boldsymbol{\pi}\boldsymbol{\varepsilon}' = 1$ can be derived as

$$\boldsymbol{\pi} = \frac{1}{1 + \mu\bar{x} + \mu^2\bar{x}^2/2!} \left[1,\ \mu\,\mathbf{p}\mathbf{V},\ \mu^2\,\mathbf{p}^{\otimes 2}\,(\mathbf{V} \otimes \mathbf{I})\,(\mathbf{B} \oplus \mathbf{B})^{-1}\right]. \tag{E.1}$$

General M_{ld}/G/N//N models: In the general case, the burst-start rate is not a constant μ, but it depends on the number i of active bursts, e.g. in the IS model $\mu(i) = (N - i)/Z$. The

generator matrix $\boldsymbol{\mathcal{Q}}$ for such general models and arbitrary large N has the form:

$$\boldsymbol{\mathcal{Q}}_N = \left[\begin{array}{c|c|c|c|c} \mathbf{Z_0} & \mathbf{X_0} & & & \\ \hline \mathbf{Y_1} & \mathbf{Z_1} & \mathbf{X_1} & & \\ \hline & \ddots & \ddots & \ddots & \\ \hline & & \mathbf{Y_{N-1}} & \mathbf{Z_{N-1}} & \mathbf{X_{N-1}} \\ \hline & & & \mathbf{Y_N} & \mathbf{Z_N} \end{array}\right],$$

$$\begin{aligned} \text{where:} \quad \mathbf{X_i} &= \mu(i)\,\mathbf{I}^{\otimes i} \otimes \mathbf{p}, \quad i = 0, ..., N-1, && \text{(E.2)} \\ \mathbf{Y_i} &= (\mathbf{B}\boldsymbol{\varepsilon}')^{\oplus i}, \quad i = 1, ..., N, && \text{(E.3)} \\ \mathbf{Z_i} &= -\mathbf{B}^{\oplus i} - \mu(i), \mathbf{I}^{\otimes i}, \quad i = 0, ..., N. && \text{(E.4)} \end{aligned}$$

Thereby, $\mu(N) := 0$ in order to obtain the right formula for $\mathbf{Z_N}$. Note that the main diagonal blocks $\mathbf{Z_i}$ of $\boldsymbol{\mathcal{Q}}_N$ are quadratic but with growing dimension T^i for i active sources. Consequently, the other two diagonals contain non-quadratic matrices $\mathbf{X_i}$ and $\mathbf{Y_i}$.

The matrix $\boldsymbol{\mathcal{L}}$ contains the corresponding cell-rates:

$$\boldsymbol{\mathcal{L}}_N = \left[\begin{array}{c|c|c|c|c} 0 & & & & \\ \hline & \lambda_p \mathbf{I} & & & \\ \hline & & 2\lambda_p \mathbf{I}^{\otimes 2} & & \\ \hline & & & \ddots & \\ \hline & & & & N\lambda_p \mathbf{I}^{\otimes N} \end{array}\right] + \lambda_0 \boldsymbol{\mathcal{I}}\,.$$

The vector $\boldsymbol{\pi}$ can be partitioned in $N+1$ sub-vectors according to the blocks in $\boldsymbol{\mathcal{Q}}_N$:

$$\boldsymbol{\pi} = [\pi_0,\ \boldsymbol{\pi_1},\ ...,\ \boldsymbol{\pi_N}]\,.$$

that can be computed by the following iterative approach:

$$\pi_0 = \left[\sum_{i=0}^{N} \frac{(\mu(i)\bar{x})^i}{i!}\right]^{-1}, \quad \boldsymbol{\pi_i} = \mu(i-1)\,(\boldsymbol{\pi_{i-1}} \otimes \mathbf{p})\,\left(\mathbf{B}^{\oplus i}\right)^{-1}, \quad i = 1, ..., N. \qquad \text{(E.5)}$$

Waiting Bursts: Finally, if bursts can queue up before they become active, when $K > N$, the $\boldsymbol{\mathcal{Q}}$-matrix has to be extended by the following $(K-N)$ blocks:

$$\boldsymbol{\mathcal{Q}} = \left[\begin{array}{c|c|c|c|c|c|c} \ddots & \ddots & \ddots & & & & \\ \hline & \mathbf{Y_{N-1}} & \mathbf{Z_{N-1}} & \mathbf{X_{N-1}} & & & \\ \hline & & \mathbf{Y_N} & \mathbf{Z_N}^{(0)} & \mathbf{X_N}^{(0)} & & \\ \hline & & & \mathbf{Y_N}\mathbf{I}^{\otimes N} \otimes \mathbf{p} & \mathbf{Z_N}^{(1)} & \mathbf{X_N}^{(1)} & \\ \hline & & & & \ddots & \ddots & \ddots \\ \hline & & & & & \mathbf{Y_N}\mathbf{I}^{\otimes N} \otimes \mathbf{p} & \mathbf{Z_N}^{(N-K)} \end{array}\right], \qquad \text{(E.6)}$$

$$\begin{aligned} \text{where:} \quad \mathbf{Y_N} &= (\mathbf{B}\boldsymbol{\varepsilon}')^{\oplus N}, \quad \text{as before,} \\ \mathbf{X_N}^{(j)} &= \mu(N+j)\,\mathbf{I}^{\otimes N} \otimes \mathbf{p}, \quad j = 0, ..., K-N-1, \\ \mathbf{Z_N}^{(j)} &= -\mathbf{B}^{\oplus N} - \mu(N+j), \mathbf{I}^{\otimes N}, \quad j = 0, ..., K-N. \end{aligned}$$

Again, $\mu(K) := 0$ has to be assumed so that the matrix $\mathbf{Z_N}^{(N-K)}$ is described correctly.

There is no comparable explicit formula for $\boldsymbol{\pi}$ any more, since the Jackson-network approach of Section D.3 does not apply when there is queueing in front of the N-server unit. Matrix-Analytic methods for M/G/N//K queues are discussed in [Lipsky 92], alternatively $\boldsymbol{\pi}$ can be derived directly from $\boldsymbol{\pi}\boldsymbol{\mathcal{Q}} = 0$ and $\boldsymbol{\pi}\boldsymbol{\varepsilon}' = 1$.

Traffic Mixes: If the traffic consists of several classes with different burst-length distributions and different cell-rates, as described in Sect. 2.5.2 and Appendix B.2, a more complicated burst-length distribution $\langle \mathbf{p}, \mathbf{B} \rangle$ results, and the $\boldsymbol{\mathcal{L}}$ matrix is slightly more general,

$$\boldsymbol{\mathcal{L}}_N = \begin{bmatrix} 0 & & & & \\ & \mathbf{L_0} & & & \\ & & \mathbf{L_0} \oplus \mathbf{L_0} & & \\ & & & \ddots & \\ & & & & \mathbf{L_0}^{\oplus N} \end{bmatrix} + \lambda_0 \boldsymbol{\mathcal{I}}\,,$$

since $\mathbf{L_0}$ can contain different intra-burst cell-rates $\lambda_p^{(j)}$ on its diagonal.

E.1.3 SHARED Model

In the model for elastic traffic in Chapter 9, the cell-rate during an individual burst depends on the number i of active bursts: each active source generates cells at rate $\beta(i)\lambda_p$. If all $\beta(i) \equiv 1$, the N-Burst model without throttling is obtained. For the SHARED model which is introduced in Sect. 9.2.1, the factors $\beta(i)$ are:

$$\beta(i) = \begin{cases} 1 & \text{when } i\lambda_p \le \nu \\ \frac{0.8\,\nu}{i\lambda_p} & \text{otherwise} \end{cases}\,,$$

i.e. no throttling is performed if $i\lambda_p < \nu$. In order to keep the number of cells per burst the same, the ON time duration has to be extended accordingly, which results in load dependent servers of the modulating IS/G$_{ld}$/N//N queue. Consequently, the load-dependence factors β affect the cell-rates in $\boldsymbol{\mathcal{L}}_N$ but also the blocks $\mathbf{Y_i}$ and $\mathbf{Z_i}$ in $\boldsymbol{\mathcal{Q}}_N$ that describe the burst-ends respectively the internal phase transitions of the ME burst-length distribution:

$$\tilde{\boldsymbol{\mathcal{Q}}}_N = \begin{bmatrix} \tilde{\mathbf{Z}}_\mathbf{0} & \mathbf{X_0} & & & \\ \tilde{\mathbf{Y}}_\mathbf{1} & \tilde{\mathbf{Z}}_\mathbf{1} & \mathbf{X_1} & & \\ & \ddots & \ddots & \ddots & \\ & & \tilde{\mathbf{Y}}_{\mathbf{N-1}} & \tilde{\mathbf{Z}}_{\mathbf{N-1}} & \mathbf{X_{N-1}} \\ & & & \tilde{\mathbf{Y}}_\mathbf{N} & \mathbf{Z_N} \end{bmatrix},$$

where:
$$\begin{aligned} \mathbf{X_i} &= \mu(i)\,\mathbf{I}^{\otimes i} \otimes \mathbf{p}, && i = 0, ..., N-1, \\ \tilde{\mathbf{Y}}_\mathbf{i} &= \beta(i)\,(\mathbf{B}\boldsymbol{\varepsilon}')^{\oplus i}, && i = 1, ..., N, \\ \tilde{\mathbf{Z}}_\mathbf{i} &= -\beta(i)\mathbf{B}^{\oplus i} - \mu(i), && i = 0, ..., N. \end{aligned}$$

Since the SHARED model is an extension of the Independent Source model, the load-dependent burst-arrival rate is $\mu(i) = (N-i)/Z$ in that case. The computation of $\boldsymbol{\pi}$ in (E.5) is also affected by the load-dependence factors $\beta(i)$:

$$\tilde{\boldsymbol{\pi}}_{\mathbf{i}} = \mu(i-1)\,(\tilde{\boldsymbol{\pi}}_{\mathbf{i-1}} \otimes \mathbf{p})\,\left(\beta(i)\mathbf{B}^{\oplus i}\right)^{-1}, \quad i = 1, ..., N,$$

where π_0 follows from normalization, $\boldsymbol{\pi}\boldsymbol{\varepsilon}' = 1$.

Finally, the matrix $\boldsymbol{\mathcal{L}}$ that describes the Poisson cell-arrival rates has to be modified:

$$\tilde{\boldsymbol{\mathcal{L}}}_N = \begin{bmatrix} 0 & & & & \\ & \beta_1\,\lambda_p\mathbf{I} & & & \\ & & \beta_2\,2\lambda_p\mathbf{I}^{\otimes 2} & & \\ & & & \ddots & \\ & & & & \beta_N\,N\lambda_p\mathbf{I}^{\otimes N} \end{bmatrix}$$

Variations: More general variations of the load-dependent burst-durations are mentioned in Sects. 2.5.2 and 2.5.3:

- If different traffic classes show different back-off behavior as proposed at the end of Sect. 2.5.2, different load-dependent factors $\beta_j(i)$, $i = 1, ..., N$, for every phase $j = 1, ..., T$ of the burst-length distribution $\langle \mathbf{p}, \mathbf{B} \rangle$ can be necessary. Thus, for i active bursts, the burst-length distribution is described by the matrix

  $$\tilde{\mathbf{B}}(i) = \mathbf{diag}\,(\beta_1(i),\ ...,\ \beta_T(i))\,\mathbf{B}.$$

  The matrix $\tilde{\mathbf{B}}(i)$ has to replace $\mathbf{B}$ in the blocks $\mathbf{Z_i}$ and $\mathbf{Y_i}$ of $\boldsymbol{\mathcal{Q}}_N$, $i = 1, ..., N$.

  Also, the i-th block of $\boldsymbol{\mathcal{L}}$ is changed from $\mathbf{L}_0{}^{\oplus i}$ to $(\,\mathbf{diag}\,(\beta_1(i), ..., \beta_T(i)) \cdot \mathbf{L}_0)^{\oplus i}$.

- If all the i active bursts do not get throttled by the same factor $\beta(i)$, but the j-th active burst is assigned its 'own' back-off factor $\beta^{(j)}(i)$, $j = 1, ..., i$, the Kronecker-Sums of the form

  $$\mathbf{B}^{\oplus i} \quad \text{or} \quad \left(\mathbf{B}\boldsymbol{\varepsilon}'\right)^{\oplus i}$$

  in the blocks $\mathbf{Z_i}$ and $\mathbf{Y_i}$ of $\boldsymbol{\mathcal{Q}}_N$ need to be replaced by

  $$\left(\beta^{(1)}(i)\mathbf{B}\right) \oplus ... \oplus \left(\beta^{(i)}(i)\mathbf{B}\right),$$

  respectively the equivalent Kronecker-Sum of $\mathbf{B}\boldsymbol{\varepsilon}'$. This back-off behavior is discussed in Sect. 2.5.3.

  The same procedure is applied to the Kronecker-Sums of the matrix $\mathbf{L_0}$ that build up the blocks of $\boldsymbol{\mathcal{L}}$.

E.2 Efficient Algorithms

In the previous section, we described the details of the MMPP representation of all kinds of N-Burst models. Once, those matrices are defined, properties of the N-Burst traffic stream such as moments and correlation can be computed numerically from the analytic model, see Appendix C.2 and C.3.

However, if the ME representation of the burst-length distribution requires more than only a few states, the state-space of N-Burst models with larger N becomes large very quickly, see the next section. Hence, efficient algorithms and possibly approaches to reduce the state-space are discussed in this chapter. We concentrate on the evaluation of the N-Burst arrival process in this chapter and move the efficient analysis of MMPP/M/1 queueing models to Appendix F.

An experimental evaluation of the efficiency of the various algorithms is part of this chapter. The experiments were conducted on a SUN Ultra 1, with 167 MHZ and 128 MB of main memory running Solaris 2.5. The algorithms in this chapter are implemented in the software package that is described in detail in [Schwefel 00a].

E.2.1 Setup of LAQT-Matrix $\mathcal{Q}$

An important criterion for the performance of the algorithms is the size of the participating matrices $\mathcal{Q}$, $\mathcal{B}$, and $\mathcal{L}$. The matrix $\mathcal{Q}$ of the [M/G/N//K] model (E.6) has dimension:

$$\dim \mathcal{Q} = \sum_{n=i}^{N} T^i + (K-N)\,T^N = \frac{T^{N+1}-1}{T-1} + (K-N)\,T^N, \quad \text{for} \quad T := \dim \mathbf{B} > 1.$$

Every set of states with $i \leq N$ active bursts has to keep track of all T possible states for each of the i active servers, which leads to T^i states at each i-active level.

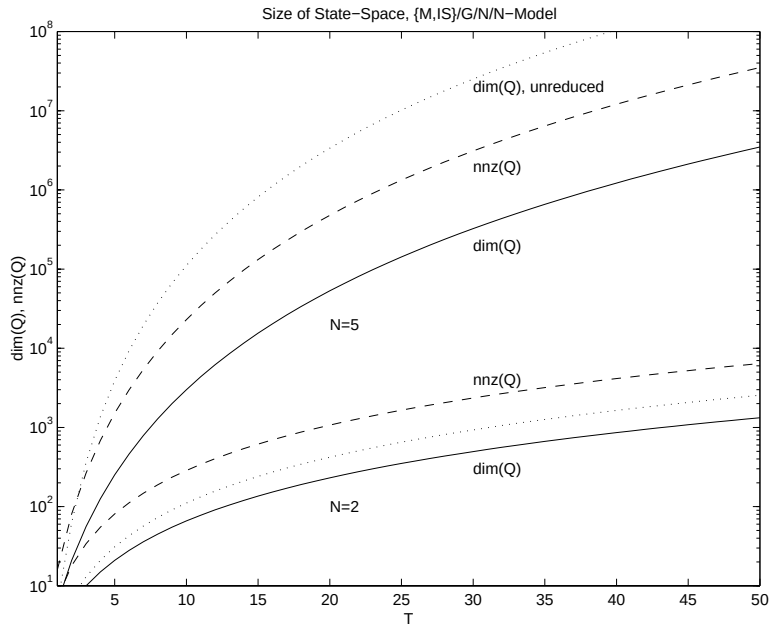

Figure E.1: **State-Space Sizes of N-burst Models:** The dimension of the $\mathcal{Q}$-matrix in the [M/G/N//K] model can be strongly reduced: The dotted lines show dim $\mathcal{Q}$ for $N = K = 2$ respectively 5 and varying number of states in the burst-length distribution, $T = \dim \mathbf{B}$. The solid line shows the number of states in the reduced representation and the dashed line the number of non-zero elements in $\mathcal{Q}$ for diagonal $\mathbf{B}$: $\mathcal{Q}$ is very sparse, i.e. the number of non-zero elements is very small compared to the overall number of elements (which is $(\dim \mathcal{Q})^2$).

However, an improvement is possible: In case of N identical servers[1] in the embedded queueing model it does not matter whether server 1 is in state k and server 2 in state j or

vice versa. Thus all states of this kind can be identified (for details see [Schwefel 97] and [Lipsky & Liefvoort 95]), and the number of states reduces to

$$\dim \boldsymbol{\mathcal{Q}} = \binom{N+T}{N} + (K-N)\binom{N+T-1}{N}.$$

Figure E.1 illustrates the gain of this reduction. Furthermore, the matrices are very sparsely populated, most of their $(\dim \boldsymbol{\mathcal{Q}})^2$ elements are 0. For a diagonal representation $\mathbf{B}$ of the burst-length distribution (as it is the case for the LRD model that uses TPT distributions, see Sect. 3.4) the number of non-zero elements in $\boldsymbol{\mathcal{Q}}$ is for the reduced representation:

$$\mathrm{nnz}\,(\boldsymbol{\mathcal{Q}}) = \dim \boldsymbol{\mathcal{Q}} + 2T\binom{N+T-1}{T} \quad \Longrightarrow \quad \frac{\mathrm{nnz}\,(\boldsymbol{\mathcal{Q}})}{\dim \boldsymbol{\mathcal{Q}}} = \frac{2N}{N/T+1} + 1 < 2N+1.$$

Thus, the number of non-zero elements grows at most linearly with the dimension of $\boldsymbol{\mathcal{Q}}$ and not quadratic as the overall number of elements does. Since $\boldsymbol{\mathcal{L}}$ is diagonal, $\boldsymbol{\mathcal{B}} = \boldsymbol{\mathcal{Q}} + \boldsymbol{\mathcal{L}}$ has the same number of non-zero elements as $\boldsymbol{\mathcal{Q}}$.

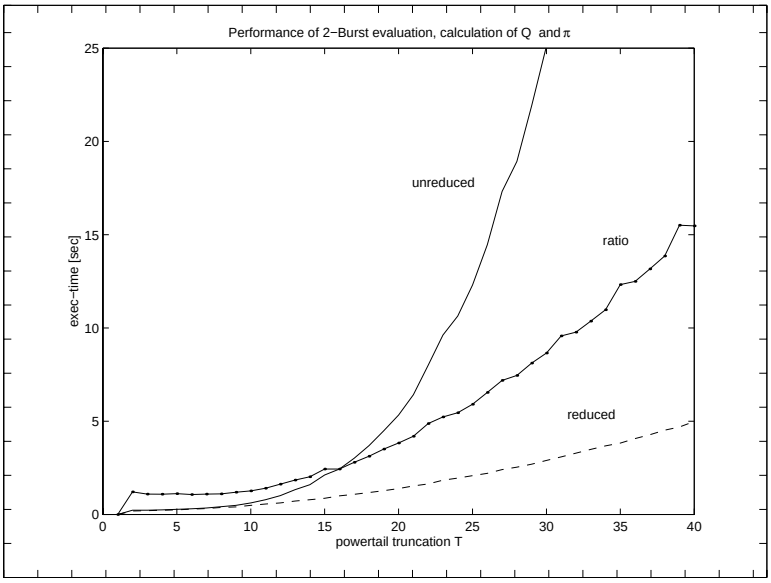

Figure E.2: **Execution Time for Matrix Setup and π Computation:** The state-space reduction has an enormous influence on the execution times, here shown for the 2-Burst [M/G/2//2] model with TPT-T distributed burst-lengths. The graph shows the execution times for the setup of $\boldsymbol{\mathcal{Q}}$ and the calculation of $\boldsymbol{\pi}$ for the unreduced state space (upper curve) versus the reduced state space (lower curve) and their ratio. For higher values of N the performance effects are even more dramatic.

Making use of the sparse structure of the matrices results in an essential improvement for both, computing time (see Figure E.2) and memory usage.

E.2.2 Calculation of Steady State Vector $\boldsymbol{\pi}$

The steady state distribution vector $\boldsymbol{\pi}$ of the modulating process with generator $\boldsymbol{\mathcal{Q}}$ is the solution of

$$\boldsymbol{\pi}\boldsymbol{\mathcal{Q}} = \mathbf{0}, \quad \text{and} \quad \boldsymbol{\pi}\boldsymbol{\varepsilon}' = 1. \tag{E.7}$$

[1] i.e. for all the variations of Sect. 2.5 except for the server-dependent scaling of Section 2.5.3.

$\boldsymbol{\mathcal{Q}}$ is singular since, $\boldsymbol{\mathcal{Q}}\boldsymbol{\varepsilon}' = 0$. So the system (E.7) can be transformed to a modified $\boldsymbol{\pi}\tilde{\boldsymbol{\mathcal{Q}}} = [\mathbf{0}, 1]$, where the last column of $\boldsymbol{\mathcal{Q}}$ is substituted by $\boldsymbol{\varepsilon}'$.

However, there exist explicit formulas to compute $\boldsymbol{\pi}$ for N-Burst $[\mathrm{M}_{ld}/\mathrm{G}_{ld}/N//N]$ models, see Sect. E.1.2, which can be computed much faster than solving Equations (E.7). Unfortunately, for models with waiting bursts $K > N$ or with more complicated types of load-dependent burst-lengths (as in Sect. 2.5.2 and 2.5.3) no simple explicit formulas are available.

If (E.5) cannot be applied, the sparse system (E.7) should be solved by exact methods (LU-decomposition, Gauß-algorithm), only if the state-space is small, say $\dim \boldsymbol{\mathcal{Q}} < 400$. For larger state-spaces, iterative methods such as *GMRES* or *Biconjugate Gradients methods* (see [Stewart et al. 92]) are much more efficient without substantial loss of numerical accuracy. Such methods iterate on the approximate solution of the linear equations until the error-vector becomes small. To speed-up convergence, *preconditioning* by incomplete LU-decomposition can be used, i.e. a partial and thus sparse LU-decomposition of the coefficient matrix is used as an approximation for its inverse.

E.2.3 Evaluation of Cell-Arrival Process

$\boldsymbol{\mathcal{Q}}$, $\boldsymbol{\mathcal{B}}$, and $\boldsymbol{\mathcal{L}}$ are sparse. Unfortunately an inverse of a sparse matrix is not sparse any more. Thus, calculating $\boldsymbol{\mathcal{V}}$ or $\boldsymbol{\mathcal{Y}}$ is inefficient in terms of computing time and memory usage. If possible, it should be avoided. According to (C.1) the moments of an SM-process are:

$$\mathbb{E}\left\{X^k\right\} = \boldsymbol{\wp}\boldsymbol{\mathcal{V}}^k\boldsymbol{\varepsilon}'.$$

For MMPPs, the first moment (mean) of the process can be derived directly from the $\boldsymbol{\pi}$-vector:

$$\mathbb{E}\{X\} = \frac{1}{\boldsymbol{\pi}\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}'}.$$

Also, the product $\boldsymbol{\wp}\boldsymbol{\mathcal{V}}$ which is needed for higher moments has a simpler form for MMPPs:

$$\boldsymbol{\wp}\boldsymbol{\mathcal{V}} = \frac{\boldsymbol{\pi}\boldsymbol{\mathcal{B}}\boldsymbol{\mathcal{V}}}{\boldsymbol{\pi}\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}'} = \mathbb{E}\{X\} \cdot \boldsymbol{\pi}.$$

since $\boldsymbol{\wp} = \boldsymbol{\pi}\boldsymbol{\mathcal{L}}/(\boldsymbol{\pi}\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}') = \boldsymbol{\pi}\boldsymbol{\mathcal{B}}/(\boldsymbol{\pi}\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}')$. Still, for higher moments the product $(\boldsymbol{\pi}\boldsymbol{\mathcal{V}})\boldsymbol{\mathcal{V}}^{k-1}$ needs to be calculated. Due to the reasoning above, it is much more efficient to solve the system $\mathbf{x_{k+1}}\boldsymbol{\mathcal{B}} = \mathbf{x_k}$, $\mathbf{x_1} = \boldsymbol{\wp}\boldsymbol{\mathcal{V}}$ for successive k to obtain $\mathbf{x_k} = \boldsymbol{\wp}\boldsymbol{\mathcal{V}}^k$ without calculating $\boldsymbol{\mathcal{V}} = \boldsymbol{\mathcal{B}}^{-1}$ explicitly.

As in the calculation of $\boldsymbol{\pi}$ from (E.7), iterative solvers such as GMRES for the linear equations are superior when the state-space is large. Also, the incomplete LU-decomposition for preconditioning only has to be done once, since the coefficient matrix never changes.

An equivalent procedure has to be applied to calculate the covariance:

$$\mathrm{Cov}(X, X_{+k}) = \boldsymbol{\wp}\boldsymbol{\mathcal{V}}\left[(\boldsymbol{\mathcal{V}}\boldsymbol{\mathcal{L}})^k - \boldsymbol{\varepsilon}'\boldsymbol{\wp}\right]\boldsymbol{\mathcal{V}}\boldsymbol{\varepsilon}' = \boldsymbol{\wp}\boldsymbol{\mathcal{V}}\left(\boldsymbol{\mathcal{V}}\boldsymbol{\mathcal{L}}\right)^k\boldsymbol{\mathcal{V}}\boldsymbol{\varepsilon}' - \mathbb{E}\{X\}^2.$$

However, if the autocorrelation coefficient needs to be calculated for large[2] lags k, the calculation of $\boldsymbol{\mathcal{Y}} = \boldsymbol{\mathcal{V}}\boldsymbol{\mathcal{L}}$ and its spectral decomposition[3], $\boldsymbol{\mathcal{Y}} = \boldsymbol{\mathcal{W}}\boldsymbol{\Lambda}\boldsymbol{\mathcal{U}}$, $\boldsymbol{\mathcal{W}} = \boldsymbol{\mathcal{U}}^{-1}$, will at some point get less costly than the iterated solution of the linear system.

[2] As it is the case to verify the existence of Long-Range Dependence.

[3] i.e. $\boldsymbol{\mathcal{Y}} = \sum_{i=1}^{S} \Lambda_i \mathbf{w_i}'\mathbf{u_i} =: \boldsymbol{\mathcal{W}}\boldsymbol{\Lambda}\boldsymbol{\mathcal{U}}$, where $\mathbf{u_i}$ are the left eigenvectors, $\mathbf{u_i}\boldsymbol{\mathcal{Y}} = \Lambda_i\mathbf{u_i}$, and Λ_i the eigenvalues of $\boldsymbol{\mathcal{Y}}$.

Even for a sparse matrix, which $\boldsymbol{\mathcal{Y}}$ is not, there do not exist any algorithms for a spectral decomposition that take advantage of the sparsity. Also, the matrix of eigenvectors is in general fully populated. Thus the spectral decomposition is advisable only for the case of large lags k since the term $\boldsymbol{\mathcal{Y}}^k$ can be calculated efficiently, once the decomposition is done.

Since $\boldsymbol{\wp}\boldsymbol{\mathcal{Y}} = \boldsymbol{\wp}$ and $\boldsymbol{\mathcal{Y}}\boldsymbol{\varepsilon}' = \boldsymbol{\mathcal{V}}\boldsymbol{\mathcal{B}}\boldsymbol{\varepsilon}' = \boldsymbol{\varepsilon}'$, $\boldsymbol{\wp}$ and $\boldsymbol{\varepsilon}'$ are the left and right eigenvectors for eigenvalue 1. The subtraction of $\boldsymbol{\varepsilon}'\boldsymbol{\wp}$ in the formula for the covariance is equivalent to neglecting eigenvalue $\Lambda_i = 1$ of $\boldsymbol{\mathcal{Y}}$. Then, $r(k)$ can be calculated as:

$$r(k) = \frac{\boldsymbol{\wp}\boldsymbol{\mathcal{V}}\boldsymbol{\mathcal{W}}\tilde{\boldsymbol{\Lambda}}^k\boldsymbol{\mathcal{U}}\boldsymbol{\mathcal{V}}\boldsymbol{\varepsilon}'}{\mathbb{E}\{X^2\} - (\mathbb{E}\{X\})^2},$$

where $\boldsymbol{\mathcal{Y}} = \boldsymbol{\mathcal{W}}\boldsymbol{\Lambda}\boldsymbol{\mathcal{U}}$, $\tilde{\boldsymbol{\Lambda}}$ is the $\boldsymbol{\Lambda}$-matrix with eigenvalue 1 reset to 0. The eigenvalue analysis delivers the left eigenvectors as columns of $\boldsymbol{\mathcal{W}}$. As usual, $\boldsymbol{\mathcal{U}}$ is not calculated explicitly but the vector $\mathbf{h} := \boldsymbol{\mathcal{U}}\boldsymbol{\mathcal{V}}\boldsymbol{\varepsilon}'$ is derived as solution of the linear system $\boldsymbol{\mathcal{W}}\mathbf{h} = \boldsymbol{\mathcal{V}}\boldsymbol{\varepsilon}'$.

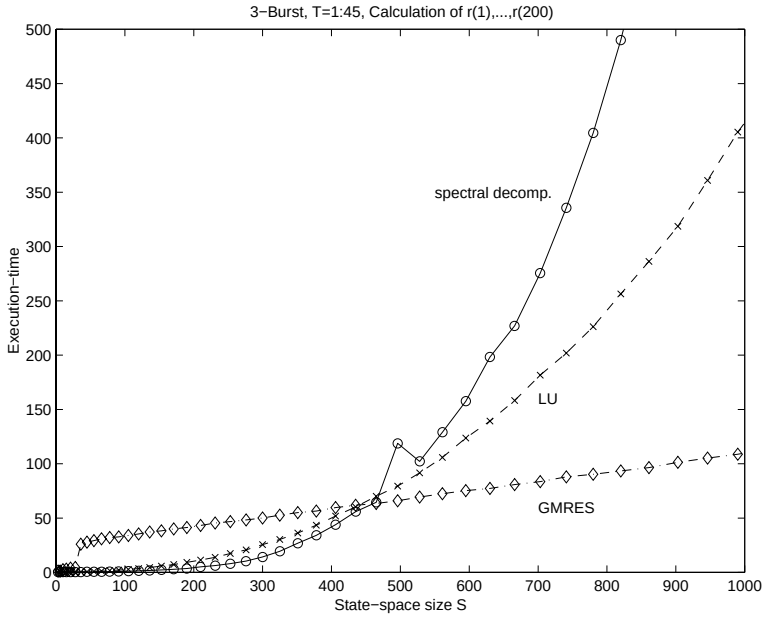

Figure E.3: **Comparison of Three Different Methods for Calculating the Coefficient of Autocorrelation:** $r(k)$ is computed for lags $k = 1, ..., 200$. The spectral decomposition is fastest for small state-spaces ($S < 465$). Applying the iterative solver GMRES $k = 200$ times is slow for small state-spaces but its execution time only grows linearly with $S = \dim \boldsymbol{\mathcal{Q}}$. The execution time of the exact solver via LU-decomposition is in between these two methods.

Figure E.3 demonstrates the advantage of the spectral decomposition for smaller state spaces in terms of execution time. Also, for large lags, the execution times increase for the iterative application of GMRES and the exact LU decomposition, while it remains nearly constant for spectral decomposition, compare also to Fig. F.2.

Appendix F

Matrix-Algebraic Algorithms for MMPP/M/1 Queue Performance

This chapter discusses algorithms that are necessary for the computation of performance parameters of MMPP/M/1 queues, in particular for the steady-state queue-length distribution and for the transient mean First Passage Time. Most algorithms can be applied to the more general SM/M/1 queues with only minor modifications.

Let $< \boldsymbol{\mathcal{B}}, \boldsymbol{\mathcal{L}} >$ be the matrix representation of the MMPP arrival process with stationary state-probability vector $\boldsymbol{\pi}$, see Appendix C.2. Furthermore, let ν be the service-rate of the exponential server in the MMPP/M/1 queue.

F.1 Steady-State Queue-Length Distribution

The principle of the matrix-geometric solution of the SM/M/1-queue is reviewed in Section D.5. The hard part is solving the quadratic matrix equation for the rate-matrix $\boldsymbol{\mathcal{R}}$:

$$\boldsymbol{\mathcal{A}}_0 + \boldsymbol{\mathcal{R}}\boldsymbol{\mathcal{A}}_1 + \boldsymbol{\mathcal{R}}^2\boldsymbol{\mathcal{A}}_2 = \mathbf{O}, \tag{F.1}$$

where $\boldsymbol{\mathcal{A}}_0 = \boldsymbol{\mathcal{L}}$, $\boldsymbol{\mathcal{A}}_1 = -(\boldsymbol{\mathcal{B}} + \nu\boldsymbol{\mathcal{I}})$, $\boldsymbol{\mathcal{A}}_2 = \nu\boldsymbol{\mathcal{I}}$. Before we discuss and compare different algorithms to solve (F.1), a few properties of $\boldsymbol{\mathcal{R}}$ can already be determined beforehand.

One prerequisite for the queue to be stable is that the utilization

$$\rho = \frac{1}{\nu\,\mathbb{E}\{X\}} = \frac{\boldsymbol{\pi}\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}'}{\nu},$$

is smaller than 1. In the stable case, the rate-matrix $\boldsymbol{\mathcal{R}}$ has eigenvalues with absolute value smaller than 1, thus $\boldsymbol{\mathcal{R}}^k \to \mathbf{O}$ for $k \to \infty$.

$\boldsymbol{\mathcal{R}}$ has only positive elements and from (F.1) it follows for the row-sums of $\boldsymbol{\mathcal{R}}$:

$$\boldsymbol{\mathcal{R}}\boldsymbol{\mathcal{A}}_2\boldsymbol{\varepsilon}' = \boldsymbol{\mathcal{A}}_0\boldsymbol{\varepsilon}'. \tag{F.2}$$

$$\text{in our scenario:} \quad \nu\boldsymbol{\mathcal{R}}\boldsymbol{\varepsilon}' = \boldsymbol{\mathcal{B}}\boldsymbol{\varepsilon}' = \boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}' \quad \Longrightarrow \quad \boldsymbol{\pi}\boldsymbol{\mathcal{R}}\boldsymbol{\varepsilon}' = \rho.$$

Once the $\boldsymbol{\mathcal{R}}$-matrix is obtained, the queue-length probabilities and overflow probabilities can be calculated as (see Section D.5):

$$\hat{\boldsymbol{\pi}}(k) = \hat{\boldsymbol{\pi}}(0)\boldsymbol{\mathcal{R}}^k = \boldsymbol{\pi}(\boldsymbol{\mathcal{I}} - \boldsymbol{\mathcal{R}})\boldsymbol{\mathcal{R}}^k, \quad k = 0, 1, \ldots .$$

It follows for the mean queue-length

$$\mathbb{E}\{Q\} = \sum_{k=0}^{\infty} k\,\hat{\boldsymbol{\pi}}(k)\boldsymbol{\varepsilon}' = \boldsymbol{\pi}\boldsymbol{\mathcal{R}}(\boldsymbol{\mathcal{I}} - \boldsymbol{\mathcal{R}})^{-1}\boldsymbol{\varepsilon}' = \frac{1}{\nu}\boldsymbol{\pi}(\boldsymbol{\mathcal{I}} - \boldsymbol{\mathcal{R}})^{-1}\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}',$$

and its second moment

$$\mathbb{E}\left\{Q^2\right\} = \sum_{k=0}^{\infty} k^2\, \hat{\boldsymbol{\pi}}(k)\boldsymbol{\varepsilon}' = \boldsymbol{\pi}\boldsymbol{\mathcal{R}}(\boldsymbol{\mathcal{R}}+\boldsymbol{\mathcal{I}})(\boldsymbol{\mathcal{I}}-\boldsymbol{\mathcal{R}})^{-2}\boldsymbol{\varepsilon}' = \frac{1}{\nu}\boldsymbol{\pi}(\boldsymbol{\mathcal{R}}+\boldsymbol{\mathcal{I}})(\boldsymbol{\mathcal{I}}-\boldsymbol{\mathcal{R}})^{-2}\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}'.$$

The queue-length-probabilities at cell-arrival times are

$$a(k) = \frac{\boldsymbol{\pi}(\boldsymbol{\mathcal{I}}-\boldsymbol{\mathcal{R}})\boldsymbol{\mathcal{R}}^k\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}'}{\boldsymbol{\pi}\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}'} = \frac{1}{\rho}\boldsymbol{\pi}(\boldsymbol{\mathcal{I}}-\boldsymbol{\mathcal{R}})\boldsymbol{\mathcal{R}}^{k+1}\boldsymbol{\varepsilon}'\,. \tag{F.3}$$

The last equality is due to $\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}' = \nu\boldsymbol{\mathcal{R}}\boldsymbol{\varepsilon}'$.

Finally, the buffer overflow probability, i.e. the probability that the primary buffer of size B is full at cell-arrival times, is:

$$\mathbb{P}\left\{Q^{(a)} \geq B\right\} = \sum_{k=B}^{\infty} a(k) = \mathbb{E}\{X\}\ \boldsymbol{\pi}\boldsymbol{\mathcal{R}}^B\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}' = \frac{1}{\rho}\boldsymbol{\pi}\boldsymbol{\mathcal{R}}^{B+1}\boldsymbol{\varepsilon}' \tag{F.4}$$

In the following Sections F.1.1 to F.1.3, different algorithms for the computation of the $\boldsymbol{\mathcal{R}}$-matrix are described. Their efficiency and their numerical accuracy is compared in Section F.1.4.

F.1.1 Simple Functional Iteration

The first approach to solve (F.1) for the rate matrix $\boldsymbol{\mathcal{R}}$ is a simple iterative procedure. Eq. (F.1) is solved for the linear term of $\boldsymbol{\mathcal{R}}$, and the resulting equation is used as the definition of one iteration-step:

$$\boldsymbol{\mathcal{R}}_{n+1} = -(\boldsymbol{\mathcal{A}}_0 + \boldsymbol{\mathcal{R}}_n^2\boldsymbol{\mathcal{A}}_2)\boldsymbol{\mathcal{A}}_1^{-1}. \tag{F.5}$$

Starting with $\boldsymbol{\mathcal{R}}_0 = \mathbf{O}$ (or another suitable initialization), this formula is repeatedly applied until convergence is reached. [Neuts 81] showed that this procedure converges and delivers a correct solution.

The individual steps of this algorithm are not very costly: The matrix inversion can be done beforehand, yet the sparsity of $\boldsymbol{\mathcal{A}}_1$ is lost thereby. Each iteration only needs two matrix-multiplications and one addition of fully populated matrices. The problem is the speed of convergence: Especially for heavy traffic situations with long queues[1] the number of necessary iterations grows very quickly and the algorithm gets unusable.

One property of the $\boldsymbol{\mathcal{R}}$ matrix can be deduced from this algorithm by induction: If $\boldsymbol{\mathcal{A}}_0$ has a row-vector with all its components being zero, then the corresponding row-vector in $\boldsymbol{\mathcal{R}}$ has zero elements as well. Thus the OFF-state of the N-Burst process with no background Poisson traffic ($\lambda_0 = 0$) causes the first row of $\boldsymbol{\mathcal{R}}$ to contain only zeros.

There are variants of the functional iteration method that speed up convergence slightly (see [Meini 97]). However, they all suffer from the same problem of an exploding number of iterations ($\gg 1000$) for scenarios with long queue lengths.

[1] In classical model long queues are due to high utilization ρ. For LRD arrival processes, long queues can happen at low utilization as well, see Chapter 5. In particular for such LRD N-Burst arrival processes, this algorithm is useless due to its inefficiency.

F.1.2 Quadratically Convergent Methods

Logarithmic Reduction [Latouche & Ramaswami 93] describe a more efficient algorithm for a time-discrete model. Since the *Cyclic Reduction* algorithm below is equally powerful, the detailed description of the Logarithmic Reduction is spared here. The interested reader be referred to [Latouche & Ramaswami 93]. Only the fundamental formulas and the outline of the algorithm are mentioned here:

The continuous process for the SM/M/1-queue is transformed to a discrete process by using the embedded Markov chain at epochs of level-transitions[2]. The algorithm then works with the following matrices for the discrete process:

- $G_{ij} = \mathbb{P}\{\text{return to level } \nu-1 \text{ in state } (\nu-1,\mathrm{j}) \mid \text{starting at } (\nu,\mathrm{i})\}$.
- $R_{ij} = \mathbb{E}\{\text{number of visits to } (\nu+1,\mathrm{j}) \text{ until return to level } \nu \mid \text{started in } (\nu,\mathrm{i})\}$.
- $U_{ij} = \mathbb{P}\{\text{return to level } \nu \text{ in state } (\nu,\mathrm{j}) \text{ under taboo of level } \nu-1 | \text{ starting at } (\nu,\mathrm{i})\}$.

The matrix $\mathbf{G}$ can be calculated efficiently from

$$\mathbf{G} = \sum_{k=0}^{\infty} \left(\prod_{i=0}^{k-1} \mathbf{B_0}^{(i)} \right) \mathbf{B_2}^{(k)}, \tag{F.6}$$

where

$$\mathbf{B_i}^{(0)} = (-\boldsymbol{\mathcal{A}}_1)^{-1}\boldsymbol{\mathcal{A}}_i, \quad \mathbf{B_i}^{(k)} = (\boldsymbol{\mathcal{I}} - \mathbf{A_1}^{(k)})^{-1}\mathbf{A_i}^{(k)}, \quad i = 0, 2,$$

with

$$\mathbf{A_0}^{(k+1)} = \left(\mathbf{B_0}^{(k)}\right)^2, \quad \mathbf{A_1}^{(k+1)} = \mathbf{B_0}^{(k)}\mathbf{B_2}^{(k)} + \mathbf{B_2}^{(k)}\mathbf{B_0}^{(k)}, \quad \mathbf{A_2}^{(k+1)} = \left(\mathbf{B_2}^{(k)}\right)^2.$$

Though all the used matrices have the dimension of $\boldsymbol{\mathcal{A}}_i$, they are written non-caligraphic to emphasize that they refer to the time-discrete model.

The $\boldsymbol{\mathcal{R}}$-matrix of the continuous SM/M/1-process is then derived by

$$\boldsymbol{\mathcal{R}} = \boldsymbol{\mathcal{A}}_0(-\mathbf{U})^{-1} = -\boldsymbol{\mathcal{A}}_0(\boldsymbol{\mathcal{A}}_1 + \boldsymbol{\mathcal{A}}_0\mathbf{G})^{-1}.$$

$\mathbf{G}$ is approximated by truncating the infinite sum in (F.6) at some index. The truncation condition is that $\mathbf{G}$ is stochastic within some error-bound δ_0:

$$||\mathbf{G}\boldsymbol{\varepsilon}' - \boldsymbol{\varepsilon}'|| < \delta_0.$$

Leaving out optimizations, the algorithm has the following form in MATLAB syntax, see [MathWorks 96]:

[2] in this context, level i means all the states with queue-length i.

```
i=0;
B0=-inv(A1)*A0;
B2=-inv(A1)*A2;

I=speye(size(A0,1));
eps=ones(size(A0,1),1);

S=B2; P=B0;
delta=1;

while delta>1e-12,
  i=i+1;                      % number of iterations (summands for S)
  B0 = inv(I-B0*B2-B2)*B0^2;
  B2 = inv(I-B0*B2-B2)*B2^2;
  S=S+P*B2;                   % S is finite sum approximating G
  P=P*B0;
  delta  = norm(eps'-S*eps');  % until S stochastic
  end;
U=A1+A0*S;
R=-A0/U;
```

Cyclic Reduction An equally powerful algorithm is based on algebraic transformations of Eq. (F.1). The algorithm for time-discrete processes and results about its behavior are given in [Bini & Meini 96]. As in Section F.1.2, it is possible to use the discrete embedded process at epochs of level-transitions. However, a variant of the algorithm for continuous processes such as our SM/M/1-model is developed here:

The algorithm transforms the Eq. (F.1) for $\boldsymbol{\mathcal{R}}$ to equations for

$$\boldsymbol{\mathcal{R}}^{(k)} := \left(\boldsymbol{\mathcal{R}}^{(k-1)}\right)^2 = \boldsymbol{\mathcal{R}}^{2^k},$$

which have the form

$$\boldsymbol{\mathcal{A}}_0^{(k)} + \boldsymbol{\mathcal{R}}^{(k)}\boldsymbol{\mathcal{A}}_1^{(k)} + \left(\boldsymbol{\mathcal{R}}^{(k)}\right)^2 \boldsymbol{\mathcal{A}}_2^{(k)} = 0. \tag{F.7}$$

Equation (F.7) is solved for $\boldsymbol{\mathcal{R}}^{(k)}$:

$$\boldsymbol{\mathcal{R}}^{(k)} = \left(-\boldsymbol{\mathcal{A}}_0^{(k)} - \left(\boldsymbol{\mathcal{R}}^{(k)}\right)^2 \boldsymbol{\mathcal{A}}_2^{(k)}\right)\left(\boldsymbol{\mathcal{A}}_1^{(k)}\right)^{-1} = \left(-\boldsymbol{\mathcal{A}}_0^{(k)} - \boldsymbol{\mathcal{R}}^{(k+1)}\boldsymbol{\mathcal{A}}_2^{(k)}\right)\left(\boldsymbol{\mathcal{A}}_1^{(k)}\right)^{-1}. \tag{F.8}$$

(F.8) is then substituted in $\left(\boldsymbol{\mathcal{R}}^{(k)} \cdot (F.7)\right)$:

$$\left[-\boldsymbol{\mathcal{A}}_0^{(k)}\left(\boldsymbol{\mathcal{A}}_1^{(k)}\right)^{-1}\boldsymbol{\mathcal{A}}_0^{(k)}\right] + \boldsymbol{\mathcal{R}}^{(k+1)}\left[\boldsymbol{\mathcal{A}}_1^{(k)} - \boldsymbol{\mathcal{A}}_2^{(k)}\left(\boldsymbol{\mathcal{A}}_1^{(k)}\right)^{-1}\boldsymbol{\mathcal{A}}_0^{(k)} - \boldsymbol{\mathcal{A}}_0^{(k)}\left(\boldsymbol{\mathcal{A}}_1^{(k)}\right)^{-1}\boldsymbol{\mathcal{A}}_2^{(k)}\right]$$

$$+(\boldsymbol{\mathcal{R}}^{(k+1)})^2\left[-\boldsymbol{\mathcal{A}}_2^{(k)}\left(\boldsymbol{\mathcal{A}}_1^{(k)}\right)^{-1}\boldsymbol{\mathcal{A}}_2^{(k)}\right] = 0.$$

Now, the coefficient matrices $\boldsymbol{\mathcal{A}}_i^{(k)}$ of Eq. (F.7) can be read off:

$$\begin{aligned}
\boldsymbol{\mathcal{A}}_0^{(k+1)} &= -\boldsymbol{\mathcal{A}}_0^{(k)} \left(\boldsymbol{\mathcal{A}}_1^{(k)}\right)^{-1} \boldsymbol{\mathcal{A}}_0^{(k)}, \\
\boldsymbol{\mathcal{A}}_1^{(k+1)} &= \boldsymbol{\mathcal{A}}_1^{(k)} - \boldsymbol{\mathcal{A}}_2^{(k)} \left(\boldsymbol{\mathcal{A}}_1^{(k)}\right)^{-1} \boldsymbol{\mathcal{A}}_0^{(k)} - \boldsymbol{\mathcal{A}}_0^{(k)} \left(\boldsymbol{\mathcal{A}}_1^{(k)}\right)^{-1} \boldsymbol{\mathcal{A}}_2^{(k)}, \\
\boldsymbol{\mathcal{A}}_2^{(k+1)} &= -\boldsymbol{\mathcal{A}}_2^{(k)} \left(\boldsymbol{\mathcal{A}}_1^{(k)}\right)^{-1} \boldsymbol{\mathcal{A}}_2^{(k)}. \\
\boldsymbol{\mathcal{A}}_i^{(0)} &= \boldsymbol{\mathcal{A}}_i, \quad i = 0, 1, 2.
\end{aligned}$$

In the stable case ($\rho < 1$), all eigenvalues of $\boldsymbol{\mathcal{R}}$ are smaller than 1, thus $\boldsymbol{\mathcal{R}}^k$ converges to $\mathbf{O}$. Then, $\boldsymbol{\mathcal{R}}^{(k)} = \boldsymbol{\mathcal{R}}^{2^k}$ converges very fast (quadratically) to $\mathbf{O}$. At some k, it can be assumed that within numerical accuracy:

$$\boldsymbol{\mathcal{R}}^{(k+1)} \approx 0.$$

Consequently, (F.7) with $\boldsymbol{\mathcal{R}}^{(k+1)} = 0$ provides an approximation for $\boldsymbol{\mathcal{R}}^{(k)}$:

$$\boldsymbol{\mathcal{R}}^{(k)} \approx -\boldsymbol{\mathcal{A}}_0^{(k)} \left[\boldsymbol{\mathcal{A}}_1^{(k)}\right]^{-1}.$$

$\boldsymbol{\mathcal{R}} = \boldsymbol{\mathcal{R}}^{(0)}$ can then be obtained by stepwise backwards substitution using (F.8):

$$\boldsymbol{\mathcal{R}}^{(k)} = -\left[\boldsymbol{\mathcal{A}}_0^{(k)} + \boldsymbol{\mathcal{R}}^{(k+1)} \boldsymbol{\mathcal{A}}_2^{(k)}\right] \left(\boldsymbol{\mathcal{A}}_1^{(k)}\right)^{-1}.$$

This backwards substitution (and thus the necessary storage of all the matrices $\left(\boldsymbol{\mathcal{A}}_i^{(j)}\right)$ can be spared with the following approach: Starting off from a rewritten form of (F.1),

$$\boldsymbol{\mathcal{R}} = -\boldsymbol{\mathcal{A}}_0 \left[\boldsymbol{\mathcal{A}}_1 + \boldsymbol{\mathcal{R}}\boldsymbol{\mathcal{A}}_2\right]^{-1} =: -\boldsymbol{\mathcal{A}}_0 \left[\hat{\boldsymbol{\mathcal{A}}}_1^{(0)} + \boldsymbol{\mathcal{R}}^{(0)} \boldsymbol{\mathcal{A}}_2^{(0)}\right]^{-1},$$

repeatedly apply (F.8) to get:

$$\boldsymbol{\mathcal{R}} = -\boldsymbol{\mathcal{A}}_0 \left[\hat{\boldsymbol{\mathcal{A}}}_1{}^{(k)} - \boldsymbol{\mathcal{R}}^{(k)} \boldsymbol{\mathcal{A}}_2^{(k)}\right]^{-1},$$

$$\text{where} \quad \hat{\boldsymbol{\mathcal{A}}}_1{}^{(\mathrm{k}+1)} = \hat{\boldsymbol{\mathcal{A}}}_1{}^{(\mathrm{k})} - \boldsymbol{\mathcal{A}}_0^{(\mathrm{k})} (\boldsymbol{\mathcal{A}}_1^{(\mathrm{k})})^{-1} \boldsymbol{\mathcal{A}}_2^{(\mathrm{k})}, \quad \hat{\boldsymbol{\mathcal{A}}}_1{}^{(0)} = \boldsymbol{\mathcal{A}}_1.$$

The elements of $\boldsymbol{\mathcal{A}}_0^{(k)}$, $\boldsymbol{\mathcal{A}}_2^{(k)}$, and $\boldsymbol{\mathcal{R}}^{(k)}$ are all non-negative and the limits for $k \to \infty$ are all bounded. For $\boldsymbol{\mathcal{A}}_0^{(k)}$ and $\boldsymbol{\mathcal{R}}^{(k)}$ the limits are actually $\mathbf{O}$ which can be used as stopping criterion for the iterations. Other possible criteria would be either the row-sums of $\boldsymbol{\mathcal{R}}$,

$$||\boldsymbol{\mathcal{R}}\boldsymbol{\mathcal{A}}_2\boldsymbol{\varepsilon}' - \boldsymbol{\mathcal{A}}_0\boldsymbol{\varepsilon}'|| < \delta,$$

or the convergence of $\boldsymbol{\mathcal{R}}^{(k)}$ to $\mathbf{O}$,

$$||\boldsymbol{\mathcal{R}}^{(k+1)}|| < \delta.$$

Most efficient is the use of $\boldsymbol{\mathcal{A}}_0^{(k)}$, since the other criteria require at least the computation of $\boldsymbol{\mathcal{R}}^{(k)}$ in each iteration. When using $\boldsymbol{\mathcal{A}}_0^{(k)} \to \mathbf{O}$ as stopping criterion, the algorithm has the following form:

```
A0k=A0; A1k=A1; A2k=A2;
A1d=A1;
k=0;
delta=1; dold=2;

while (delta>1e-20) & (abs(delta-dold)>1e-20),
        dold=delta;
        H0=A0k*inv(-A1k);
        H2=A2k*inv(-A1k);

        A1d=A1d+H0*A2k;

        A1k=A1k+H2*A0k+H0*A2k;
        A2k=H2*A2k;
        A0k=H0*A0k;
        k=k+1;
        delta=max(max(A0k));
end;

R=-A0/(A1d-A0k/A1k*A2k);
```

The number of iterations is determined by the speed of the convergence of $\boldsymbol{\mathcal{R}}^{(k)} \to \mathbf{O}$. If the eigenvalues of $\boldsymbol{\mathcal{R}}$ are close to 1, which is the case for LRD N-Burst traffic in the blow-up regions, see 5.1, the number of iterations increases but due to the quadratic convergence of $\boldsymbol{\mathcal{R}}^{(k)} = \boldsymbol{\mathcal{R}}^{2^k}$ not nearly as much as for the simple functional iteration of Sect. F.1.1.

F.1.3 Spectral Expansion

Spectral Expansion tries to overcome the problem that the number of iterations depends on the resulting queue-lengths (i.e. the closeness of the eigenvalues of $\boldsymbol{\mathcal{R}}$ to 1).

The method calculates the eigenvalues, λ_i, and the left eigenvectors, $\mathbf{u}_i$, of the rate-matrix $\boldsymbol{\mathcal{R}}$:

$$\mathbf{u}_i \boldsymbol{\mathcal{R}} = \lambda_i \mathbf{u}_i.$$

Let $\boldsymbol{\mathcal{U}}$ be the matrix with rows $\mathbf{u}_i$, and let $\boldsymbol{\mathcal{W}}$ be the matrix whose columns $\mathbf{w}_i'$ are the right eigenvalues of $\boldsymbol{\mathcal{R}}$ such that

$$\boldsymbol{\mathcal{W}} = \boldsymbol{\mathcal{U}}^{-1}.$$

Furthermore, let $\boldsymbol{\Lambda} = \mathbf{diag}\,(\lambda_1, ..., \lambda_S)$. Then, the spectral decomposition of $\boldsymbol{\mathcal{R}}$ is:

$$\boldsymbol{\mathcal{R}} = \sum_{i=1}^{S} \lambda_i \mathbf{w}_i' \mathbf{u}_i = \boldsymbol{\mathcal{W}} \boldsymbol{\Lambda} \boldsymbol{\mathcal{U}}.$$

The left eigenvectors of $\boldsymbol{\mathcal{R}}$ and its eigenvalues can be computed directly from (F.1) by the following: Multiplying (F.1) by $\mathbf{u}_i$ and substituting $\mathbf{u}_i \boldsymbol{\mathcal{R}} = \lambda_i \mathbf{u}_i$ yields a general, quadratic eigenvalue (EV) problem:

$$\mathbf{u}_i \boldsymbol{\mathcal{A}}_0 + \lambda_i \mathbf{u}_i \boldsymbol{\mathcal{A}}_1 + \lambda_i^2 \mathbf{u}_i \boldsymbol{\mathcal{A}}_2 = \mathbf{0}.$$

The quadratic EV-equation can be reduced to a standard EV-problem which has twice the dimension: First introduce the vector $\tilde{\mathbf{u}_i}$ and rewrite the quadratic EV-equation:

$$\tilde{\mathbf{u}_i} := \lambda_i \mathbf{u}_i \quad \Longrightarrow \quad \mathbf{u}_i \boldsymbol{\mathcal{A}}_0 \boldsymbol{\mathcal{A}}_2^{-1} + \tilde{\mathbf{u}_i} \boldsymbol{\mathcal{A}}_1 \boldsymbol{\mathcal{A}}_2^{-1} + \lambda_i \tilde{\mathbf{u}_i} = \mathbf{0}. \tag{F.9}$$

Then the standard EV-problem combines both equalities from (F.9):

$$[\mathbf{u_i},\, \tilde{\mathbf{u}}_\mathbf{i}] \left[\begin{array}{c|c} \mathbf{O} & -\boldsymbol{\mathcal{A}}_0\boldsymbol{\mathcal{A}}_2^{-1} \\ \hline \mathbf{I} & -\boldsymbol{\mathcal{A}}_1\boldsymbol{\mathcal{A}}_2^{-1} \end{array}\right] = \lambda_i\,[\mathbf{u}_i,\, \tilde{\mathbf{u}}_i] \tag{F.10}$$

Equation (F.10) is of twice the dimension of $S = \dim(\boldsymbol{\mathcal{R}})$. Thus, $(2S)$ possible solutions for λ_i exist. All combinations of S eigenvalues out of those form a solution of the quadratic Eq. (F.1). If the SM/M/1-queue is stable ($\rho < 1$), half of the eigenvalues from (F.10) are smaller than 1, one eigenvalue is exactly one, and the rest is greater than one (in absolute value since the eigenvalues can be complex). The S eigenvalues with smallest absolute value belong to the minimal solution $\boldsymbol{\mathcal{R}}$. Actually, the other half of the eigenvalues is not calculated in vain, but can be used to get the minimal solution $\boldsymbol{\mathcal{S}}$ of the equation

$$\boldsymbol{\mathcal{S}}^2\boldsymbol{\mathcal{A}}_0 + \boldsymbol{\mathcal{S}}\boldsymbol{\mathcal{A}}_1 + \boldsymbol{\mathcal{A}}_2 = \mathbf{O}. \tag{F.11}$$

$\boldsymbol{\mathcal{R}}^{-1}$ is a possible solution of (F.11), however it is not minimal. Yet, the minimal solution for $\boldsymbol{\mathcal{S}}$ can be derived from the larger half of the eigenvalues, which together with the corresponding eigenvectors form the spectral decomposition of $\boldsymbol{\mathcal{S}}^{-1}$. $\boldsymbol{\mathcal{S}}$ is important for the computation of the queue-length distribution for SM/M/1/B loss models in Appendix D.6, see also [Krieger et al. 98].

The standard EV-problem (F.10) has a sparse coefficient matrix for the SM/M/1-queue, since all the participating matrices are sparse and $\boldsymbol{\mathcal{A}}_2$ is diagonal (thus its inverse is also sparse). There are algorithms (see [Vorst & Golub 97]) that efficiently calculate some eigenvalues and eigenvectors of large sparse matrices. However, the complete decomposition is normally done by standard (non-sparse) methods, such as the QR-algorithm, which is used herein, see [MathWorks 96].

The equivalent formula to (F.2) for the row-sum of $\boldsymbol{\mathcal{R}}$ in its spectral decomposition follows:

$$\boldsymbol{\mathcal{U}}\boldsymbol{\mathcal{R}} = \boldsymbol{\Lambda}\boldsymbol{\mathcal{U}} \quad\Longrightarrow\quad \boldsymbol{\Lambda}\boldsymbol{\mathcal{U}}\boldsymbol{\mathcal{A}}_2\boldsymbol{\varepsilon}' = \boldsymbol{\mathcal{U}}\boldsymbol{\mathcal{R}}\boldsymbol{\mathcal{A}}_2\boldsymbol{\varepsilon}' = \boldsymbol{\mathcal{U}}\boldsymbol{\mathcal{A}}_0\boldsymbol{\varepsilon}'.$$

One advantage of the spectral expansion method is that it delivers the spectral decomposition of $\boldsymbol{\mathcal{R}}$ right away. Then, the calculation of $\boldsymbol{\mathcal{R}}^k = \boldsymbol{\mathcal{W}}\boldsymbol{\Lambda}^k\boldsymbol{\mathcal{U}}$ and of $(\boldsymbol{\mathcal{I}} - \boldsymbol{\mathcal{R}})^{-1}$ can be done very efficiently. The formulas for the performance parameters have the following form:

$$\text{Queue}-\text{length Probabilities}: \quad \hat{\pi}(k) = \boldsymbol{\pi}\boldsymbol{\mathcal{W}}(\boldsymbol{\mathcal{I}} - \boldsymbol{\Lambda})\boldsymbol{\Lambda}^k\boldsymbol{\mathcal{U}}\boldsymbol{\varepsilon}'. \tag{F.12}$$

$$\text{Mean Queue}-\text{length}: \quad \mathbb{E}\{Q\} = \boldsymbol{\pi}\boldsymbol{\mathcal{W}}(\boldsymbol{\mathcal{I}} - \boldsymbol{\Lambda})^{-1}\boldsymbol{\Lambda}\boldsymbol{\mathcal{U}}\boldsymbol{\varepsilon}'. \tag{F.13}$$

$$\text{Second Moment}: \quad \mathbb{E}\{Q^2\} = \boldsymbol{\pi}\boldsymbol{\mathcal{W}}(\boldsymbol{\mathcal{I}} + \boldsymbol{\Lambda})(\boldsymbol{\mathcal{I}} - \boldsymbol{\Lambda})^{-2}\boldsymbol{\Lambda}\boldsymbol{\mathcal{U}}\boldsymbol{\varepsilon}'. \tag{F.14}$$

$$\text{Buffer Overflow Probability}: \quad \mathbb{P}\left\{Q^{(a)} \geq B\right\} = \mathbb{E}\{X\}\,\boldsymbol{\pi}\boldsymbol{\mathcal{W}}\boldsymbol{\Lambda}^B\boldsymbol{\mathcal{U}}\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}'. \tag{F.15}$$

The right eigenvectors of $\boldsymbol{\mathcal{R}}$, which are the columns of $\boldsymbol{\mathcal{W}}$, always appear as $\boldsymbol{\pi}\boldsymbol{\mathcal{W}}$ in these formulas. Thus, $\boldsymbol{\mathcal{W}} = \boldsymbol{\mathcal{U}}^{-1}$ does not have to be calculated explicitly, but only $\mathbf{y}\boldsymbol{\mathcal{U}} = \boldsymbol{\pi}$ has to be solved for $\mathbf{y} = \boldsymbol{\pi}\boldsymbol{\mathcal{W}}$.

Also, $\boldsymbol{\mathcal{L}}\boldsymbol{\varepsilon}'$ in the buffer-overflow formula can be replaced by $\nu\boldsymbol{\mathcal{R}}\boldsymbol{\varepsilon}'$ (F.2) which yields:

$$\mathbb{P}\left\{Q^{(a)} \geq B\right\} = \frac{1}{\rho}\mathbf{y}\boldsymbol{\Lambda}^{B+1}\boldsymbol{\mathcal{U}}\boldsymbol{\varepsilon}'.$$

For large buffers, the eigenvalues of $\boldsymbol{\mathcal{R}}$ that are close to zero do not matter in $\boldsymbol{\mathcal{R}}^B$. Thus only the ones that are close to 1 would be sufficient and those can be calculated by special algorithms that take advantage of the sparsity of the coefficient matrix in (F.10). However, to be able to calculate the vector $\mathbf{y}$ with $\mathbf{y}\boldsymbol{\mathcal{U}} = \boldsymbol{\pi}$, the complete decomposition is necessary (which is done by non-sparse algorithms).

F.1.4 Execution Times and Numerical Error

Finally, the different methods of solving Eq. (F.1) for the rate-matrix $\boldsymbol{\mathcal{R}}$ of the SM/M/1-queue are compared in terms of their execution time and their numerical accuracy. The execution time is measured in seconds of CPU-time on a `SUN Ultra 1`, `167 MHz`, with `128 MB` of main memory running `Solaris 2.5`.

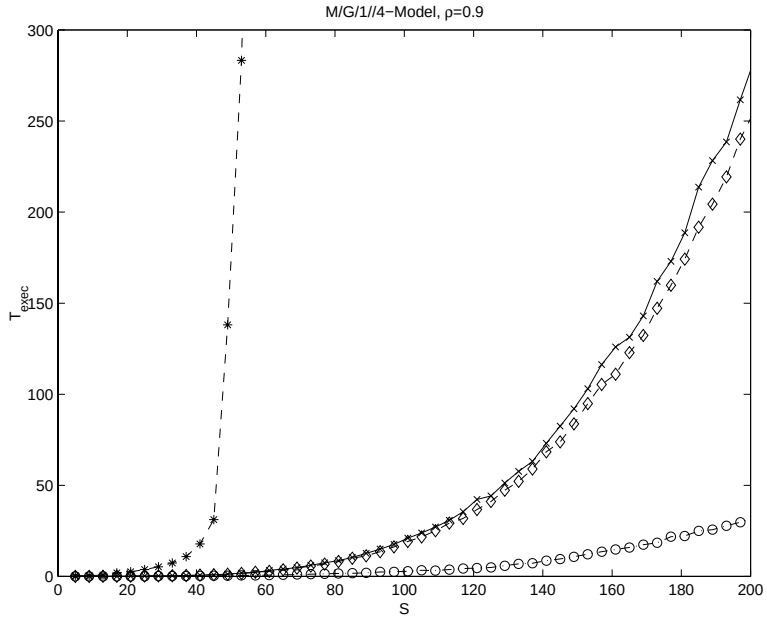

Figure F.1: **Comparison of the Execution Time of Four Different Methods for Calculating the $\boldsymbol{\mathcal{R}}$:** Simple functional iteration (stars), Logarithmic Reduction (crosses), Cyclic Reduction (diamonds), and Spectral Expansion (circles). The graph shows the execution times for an [M/G/1//4] model with $\rho = 0.9$ and TPT-distributed burst-lengths with $T = 1, ..., 52$ which results in state spaces of dimension $S = 5, ..., 209$. For the high utilization $\rho = 0.9$, Spectral Expansion outperforms all the other methods. Simple functional iteration requires such a high number of iterations that it is practically useless.

Figure F.1 compares the execution times of all four methods for the 1-Burst [M/G/1//4] arrival process with an increasing number of states for the TPT-distribution of the burst-lengths. While the simple functional iteration becomes useless very quickly due to the large number of iterations, Spectral Expansion outperforms all other methods.

However, the number of iterations of the iterative algorithms depends on the resulting queue-length distribution: The longer the queues, the more iterations are necessary. Thus, for small utilization ρ, the efficient iterative approaches are faster than Spectral Expansion. But as Figure F.2 shows, the higher ρ gets, the longer these iterative algorithms take, while spectral expansion has almost constant execution time. The particularly strong increase in execution time at about $\rho = 0.7$ in Figure F.2 is a consequence of the blow-up point that is located at exactly that value of ρ, see Sect. 5.1.

On the other hand, it also needs to be investigated, how accurate the numerical results of the individual algorithms are. There are two possible approaches to define the numerical error:

$$\delta = ||\boldsymbol{\mathcal{A}}_0 + \boldsymbol{\mathcal{R}}\boldsymbol{\mathcal{A}}_1 + \boldsymbol{\mathcal{R}}^2\boldsymbol{\mathcal{A}}_2||,$$

$$\text{or} \qquad \delta^* = ||\boldsymbol{\mathcal{R}}\boldsymbol{\mathcal{A}}_2\boldsymbol{\varepsilon}' - \boldsymbol{\mathcal{A}}_0\boldsymbol{\varepsilon}'||$$

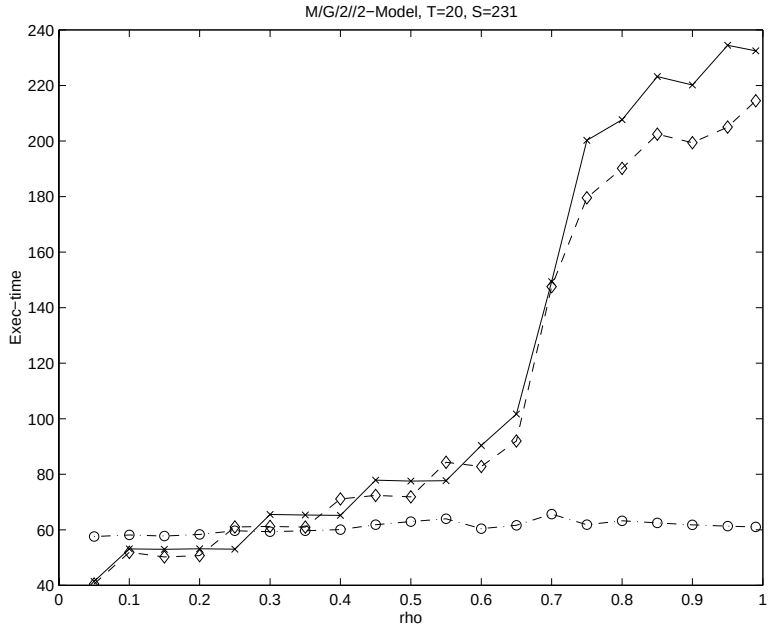

Figure F.2: **Execution Times for N-Burst/M/1 Queues with Varying Utilization:** The number of iterations in Logarithmic Reduction (crosses) and Cyclic Reduction (diamonds) depends on the resulting queue-length distribution. Spectral Expansion (circles) on the other hand requires almost constant computing time, i.e. it is slower than the iterative algorithms for smaller values of ρ, but much more efficient for higher utilization. The numerical error δ (see below) did not show any trends over that range of ρ, it was in the order of 10^{-16} for spectral expansion and 10^{-18} for the iterative algorithms.

The definition of δ follows straight from (F.1). δ^* uses property (F.2) of $\boldsymbol{\mathcal{R}}$.

Figure F.3 shows δ for different state-space sizes of the [M/G/1//4] model with $\rho = 0.5$. While the numerical error in the spectral decomposition method is determined by the eigenvalue and eigenvector analysis, the iterative methods can decrease the error in additional iterations. In Figure F.3 the numerical error, δ, seems to be a huge disadvantage of the spectral decomposition method: It is larger by a factor of about 10^{11} for state spaces with more than 100 states of the [M/G/1//4] model. δ^* shows a similar behavior.

The numerical error depends heavily on the model under consideration: E.g. for the [M/G/2//2] model and $200 < S < 800$, δ for the spectral decomposition is always below 10^{-15}, still about a factor of 20 worse than the iterative approaches, but not nearly as bad in for the [M/G/1//4] case.

Furthermore, even though $\delta = 8.3 \cdot 10^{-9}$ when using Spectral Expansion for the [M/G/1//4] arrival process with $T = 40$ states, the relative numerical errors for $\mathbb{E}\{Q\}$ and $\mathbb{E}\{Q^2\}$ are of the order of $6 \cdot 10^{-15}$, which is in the order of the numerical accuracy of the used floating point numbers.

For buffer-overflow probabilities with a buffer of $B = 10^3$ cells, the relative numerical error increases to about $2 \cdot 10^{-12}$, but it is still acceptable. Thus the numerical error δ can be rather large without leading to numerically unacceptable results for the performance parameters of the SM/M/1 queue.

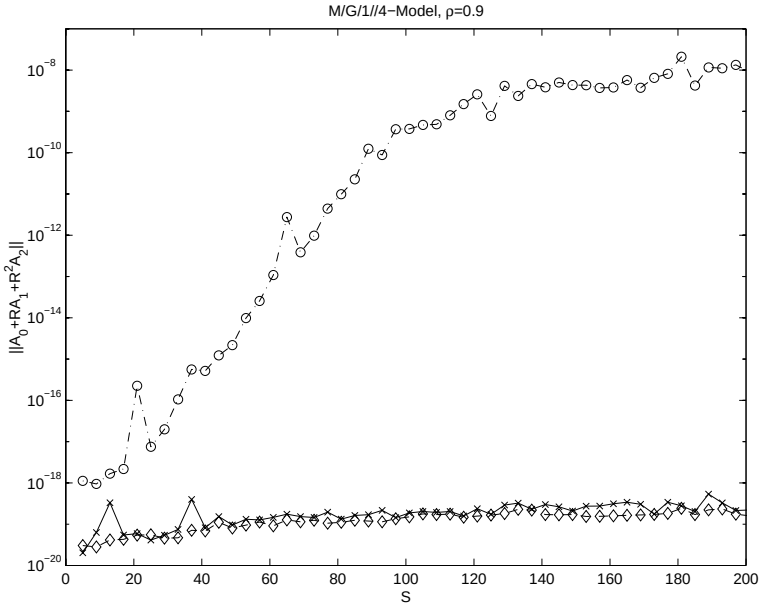

Figure F.3: **Numerical error $\delta = ||\mathcal{A}_0 + \mathcal{R}\mathcal{A}_1 + \mathcal{R}^2\mathcal{A}_2||$ for the Three Algorithms:** While the iterative methods hold the error below 10^{-18}, it grows up to 10^{-8} for the spectral expansion (circles) in the considered case of the [M/G/1//4] model.

Conclusion: Spectral expansion needs almost constant execution times for a fixed state-space size, but the numerical errors δ can get rather high. The iterative approaches are quicker for some parameter sets of the model, but their execution time strongly depends on the model-parameters, especially for N-Burst arrival processes with LRD properties as in Figure F.2. But their advantage is the better control of the numerical error.

Summary of Execution Times

The comparison of execution times is now extended to the whole N-Burst package (see [Schwefel 00a]), which includes the matrix-setup and the arrival process evaluation. The algorithms for the latter two tasks are discussed in Appendix E.2.

Since the SM/M/1-evaluation with the spectral expansion does not depend as much on the parameters of the arrival process as the other iterative algorithms, that method is used for queueing analysis in this part.

As in Section F.1.4, the execution time is measured in seconds of CPU-time on a `SUN Ultra 1, 167 MHz`.

Figure F.4 shows the execution times of all the program modules of the package for the [M/G/3//3] model, reduced state space, with TPT-distributions of varying T for the burst-lengths. The queueing analysis and the calculation of the moments with exact algorithms do not take advantage of the sparsity of the participating matrices and therefore are much worse in terms of execution time (and also memory usage). All the purely sparse algorithms can handle huge state-spaces of up to 30 000 states very efficiently with $T_{exec} < 20\,sec$ (Figure F.4, bottom).

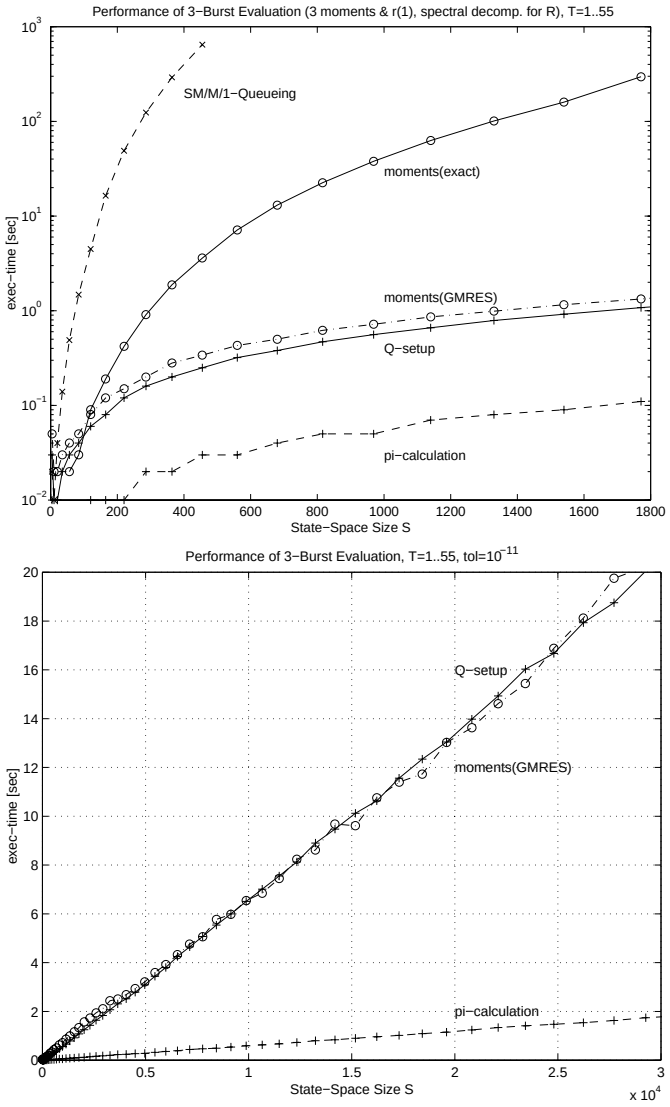

Figure F.4: **Execution Times of Various Program-Modules for Large State-Spaces:** 3-Burst models with up to 30 000 states are considered in the lower graph. The execution times still remain sustainable due to the sparsity of the involved matrices. The graphs show the execution times for the setup of the $\boldsymbol{Q}$-matrix of the [M/G/3//3] model, the calculation of $\boldsymbol{\pi}$, the calculation of the first three moments and $r(1)$ (using GMRES versus exact solver), and finally SM/M/1-Queueing using spectral expansion.

F.2 Mixed Matrix-Geometric Solution of TCP Queueing Model

The infinite generator matrix of the TCP_{B_1} queueing model in Sect. 9.2.2 has the following Quasi-Birth-Death Structure:

$$\hat{\mathbf{Q}} = \left[\begin{array}{ccccccc} \overline{\mathbf{A}_1} & \mathbf{A}_0 & & & & & \\ \mathbf{A}_2 & \mathbf{A}_1 & \mathbf{A}_0 & & & & \\ & \ddots & \ddots & \ddots & & & \\ & & \mathbf{A}_2 & \mathbf{A}_1 & \mathbf{A}_0 & & \\ \hline & & & \mathbf{C}_2 & \mathbf{C}_1 & \mathbf{C}_0 & \\ & & & & \ddots & \ddots & \ddots \end{array}\right],$$

where $\overline{\mathbf{A}_1} = \boldsymbol{\mathcal{Q}}_N - \boldsymbol{\mathcal{L}}_N$, $\mathbf{A}_0 = \boldsymbol{\mathcal{L}}_N$, $\mathbf{A}_1 = \boldsymbol{\mathcal{Q}}_N - \boldsymbol{\mathcal{L}}_N - \nu\boldsymbol{\mathcal{I}}$ refer to the unthrottled N-Burst base model, see Appendix E.1.2. From level B_1 on those matrices are replaced by the shared-bandwidth model SHARED, see Appendix E.1.3: $\mathbf{C}_0 = \tilde{\boldsymbol{\mathcal{L}}}_N$, $\mathbf{C}_1 = \tilde{\boldsymbol{\mathcal{Q}}}_N - \tilde{\boldsymbol{\mathcal{L}}}_N - \nu\boldsymbol{\mathcal{I}}$. The service process is always exponential with rate ν: $\mathbf{A}_2 = \mathbf{C}_2 = \nu\boldsymbol{\mathcal{I}}$.

It can be shown (see [Krieger & Naumov 99] for the necessary proofs in a comparable scenario) that the block-partitioned steady-state probability distribution of such a process, $\boldsymbol{\pi} = [\boldsymbol{\pi}_0, \boldsymbol{\pi}_1, ...]$ with $\boldsymbol{\pi}\hat{\mathbf{Q}} = \mathbf{0}$, can be expressed in the following mixed matrix-geometric form

$$\begin{aligned} \boldsymbol{\pi}_i &= \mathbf{a}\,\mathbf{R}^i + \mathbf{b}\,\mathbf{S}^{B_1-1-i}, \quad && i = 0, ..., B_1 - 1 \\ \boldsymbol{\pi}_k &= \boldsymbol{\pi}_{B_1}\,\mathbf{T}^{k-B_1}, && k = B_1, B_1+1, ... \end{aligned}$$

where the matrix factors $\mathbf{R}, \mathbf{S}, \mathbf{T}$ are the minimal solutions of the following quadratic matrix equations:

$$\mathbf{A}_0 + \mathbf{R}\mathbf{A}_1 + \mathbf{R}^2\mathbf{A}_2 = 0, \quad \mathbf{A}_2 + \mathbf{S}\mathbf{A}_1 + \mathbf{S}^2\mathbf{A}_0 = 0,$$
$$\mathbf{C}_0 + \mathbf{T}\mathbf{C}_1 + \mathbf{T}^2\mathbf{C}_2 = 0.$$

The vectors $\mathbf{a}$, $\mathbf{b}$, and $\boldsymbol{\pi}_{B_1}$ follow from the boundary equations at level 0, $B_1 - 1$, B_1 and from normalization $\boldsymbol{\pi}\boldsymbol{\varepsilon}' = 1$ as the solution of the following system of linear equations:

$$[\mathbf{a},\ \mathbf{b}] \left[\begin{array}{c|c|c} \mathbf{L_a} & \mathbf{R_a} & \mathbf{d_1} \\ \hline \mathbf{L_b} & \mathbf{R_b} & \mathbf{d_2} \end{array}\right] = [\mathbf{0},\ \mathbf{0},\ 1]\ ,$$

where

$$\begin{aligned} \mathbf{L_a} &= \overline{\mathbf{A}_1} + \mathbf{R}\,\mathbf{A}_2, \\ \mathbf{R_a} &= \mathbf{R}^{B_1-2}\ \left(\mathbf{R}\mathbf{A}_0 - (\mathbf{A}_0 + \mathbf{R}\mathbf{A}_1)\,\mathbf{C}_2^{-1}\cdot\right. \\ &\qquad \left.\cdot(\mathbf{C}_1 + \mathbf{T}\mathbf{C}_2)\right), \\ \mathbf{L_b} &= \mathbf{S}^{B_1-2}\ \left(\mathbf{S}\overline{\mathbf{A}_1} + \mathbf{A}_2\right), \\ \mathbf{R_b} &= \mathbf{A}_0 - (\mathbf{S}\mathbf{A}_0 + \mathbf{A}_1)\ \mathbf{C}_2^{-1}\ (\mathbf{C}_1 + \mathbf{T}\mathbf{C}_2), \\ \mathbf{d_1} &= \sum_{k=0}^{B_1-1} \mathbf{R}^k\boldsymbol{\varepsilon}' - \\ &\qquad -\mathbf{R}^{B_1-2}\,(\mathbf{A}_0 + \mathbf{R}\mathbf{A}_1)\,\mathbf{C}_2^{-1}\ \left(\sum_{k=0}^{\infty}\mathbf{T}^k\boldsymbol{\varepsilon}'\right), \\ \mathbf{d_2} &= \sum_{k=0}^{B_1-1} \mathbf{S}^k\boldsymbol{\varepsilon}' - (\mathbf{S}\mathbf{A}_0 + \mathbf{A}_1)\ \mathbf{C}_2^{-1}\ \left(\sum_{k=0}^{\infty}\mathbf{T}^k\boldsymbol{\varepsilon}'\right). \end{aligned}$$

Finally,

$$\boldsymbol{\pi}_{B_1} = -\mathbf{a}\mathbf{R}^{B_1-2}\,(\mathbf{A}_0 + \mathbf{R}\mathbf{A}_1) + \mathbf{b}\,(\mathbf{S}\mathbf{A}_0 + \mathbf{A}_1)\,\mathbf{C}_2^{-1}\,.$$

The equations for $B_1 = 0$ (SHARED model), $B_1 = 1$, and $B_1 = \infty$ (N-Burst model) are somewhat simpler, but they have to be treated separately.

F.3 Computation of mean First Passage Times

An algorithms for the computation of mFPTs of M/ME/1 queues is derived in [Lipsky 92], pp. 174-186. That algorithm can be extended to MMPP/M/1 queues where the arrival process is an MMPP, described by $< \boldsymbol{\mathcal{B}}, \boldsymbol{\mathcal{L}} >$. The mFPT to buffer level n can be expressed as the sum of the mean times, it takes to reach level $\ell + 1$ from level ℓ, for $\ell = 0, ..., n-1$:

$$\text{mFPT}(n) = \sum_{\ell=0}^{n-1} \mathbf{p_u}(\ell)\boldsymbol{\tau}'_{\mathbf{u}}(\ell)\,.$$

Thereby, the vector

$$\mathbf{p_u}(\ell) = \boldsymbol{\wp}_0 \cdot \mathbf{H_u}(0) \cdot ... \cdot \mathbf{H_u}(\ell - 1)\,,$$

contains the state probabilities of the arrival process when reaching level ℓ for the first time. $\boldsymbol{\wp}_0$ is the initial state-probability vector.

The vector, $\boldsymbol{\tau}'_{\mathbf{u}}(\ell)$, has as its components the mean time that is necessary to reach level $\ell+1$ for the first time, given that the process started at level ℓ with the arrival process in the corresponding state. $\boldsymbol{\tau}'_{\mathbf{u}}(\ell)$ can be computed by the following recursive formulas:

$$\boldsymbol{\tau}'_{\mathbf{u}}(0) = \boldsymbol{\mathcal{B}}^{-1}\boldsymbol{\varepsilon}', \quad \boldsymbol{\tau}'_{\mathbf{u}}(n) = \left[\boldsymbol{\mathcal{B}} + \nu\boldsymbol{\mathcal{I}} - \nu\mathbf{H_u}(n-1)\right]^{-1}\,\left[\boldsymbol{\varepsilon}' + \nu\boldsymbol{\tau}'_{\mathbf{u}}(n-1)\right]\,.$$

Finally, a recursive formula to compute the matrix $\mathbf{H_u}(\ell)$ can be derived. Its elements $(H_u)_{ij}$ reflect the probability that the arrival process is in state j when reaching level $\ell + 1$ for the first time, given that the queue started at level ℓ with the arrival process in state i:

$$\mathbf{H_u}(\ell) = \left[\boldsymbol{\mathcal{B}} + \nu\boldsymbol{\mathcal{I}} - \nu\mathbf{H_u}(\ell-1)\right]^{-1}\boldsymbol{\mathcal{L}}, \quad \mathbf{H_u}(0) = \boldsymbol{\mathcal{B}}^{-1}\boldsymbol{\mathcal{L}}\,.$$

The described algorithm requires n iterations to compute mFPT(n). For larger n, a more efficient algorithm is presented in [Latouche & Ramaswami 99], which only require of the order of $\log n$ iterations.

Appendix G

Derivation of the Asymptotic Behavior of the Conditional Overflow Ratio

The conditional Cell-Loss Ratio ($\mathrm{CLR}_c(t_0, B)$) in the finite-buffer loss model and its equivalent in the infinite buffer back-up model, the conditional Buffer-Overflow Ratio ($\mathrm{BOR}_c(t_0, B)$), are defined and discussed in Sect. 6.4. In this chapter, we derive the asymptotic behavior for large buffers for both transient performance measures.

G.1 Conditional Cell Loss Ratio for Infinite Power-Tails

Overflows for N-Burst/M/1 queueing models for large B and ρ not too close to 1 ($1-\rho$ is the fraction of the time that the server is idle) are caused by long over-saturation periods, during which the mean arrival rate is temporarily raised beyond the service-rate. According to Sect. 5.2, such over-saturation periods for the N-Burst model in blow-up region i_0 have a duration X that is also Power-Tailed, but with exponent $\beta = i_0(\alpha - 1) + 1$ (see Eq. (5.4)).

$$R(x) \sim c/x^{\beta}, \quad f(x) \sim \beta c/x^{\beta+1}\,, \tag{G.1}$$

when the individual burst-lengths are Power-Tail distributed with exponent α.

During those over-saturation periods, the average cell-arrival rate is

$$\Lambda_{i_0} = \lambda + i_0(\lambda_p - \kappa),$$

i.e. cells accumulate in the queue with average rate $(\Lambda_{i_0} - \nu) > 0$. Assuming that the queue is empty at the start of the over-saturation period, the length of the over-saturation period has to be at least

$$X > \frac{B}{\Lambda_{i_0} - \nu} =: x_0 \tag{G.2}$$

to be able to fill up the buffer to level B with the average rate $(\Lambda_{i_0} - \nu) > 0$. In the following, we look at over-saturation periods, X', which are long enough to cause cell losses, i.e. X' is the conditional random variable $X' = (X \mid X > x_0)$. The density function and reliability function of X' are obtained by normalization:

$$f_{X'}(x) = \frac{f(x)}{R(x_0)}, \quad R_{X'}(x) = \frac{R(x)}{R(x_0)}, \quad \text{for} \quad x > x_0.$$

Hence, assuming a PT distribution (G.1) with exponent β for X, the expected value of the over-saturation periods X' that are conditioned on having at least length x_0 is

$$\mathbb{E}\left\{X' - x_0\right\} = \int_{x_0}^{\infty} (x - x_0)\, f_{X'}(x)\, dx \approx \frac{1}{\beta - 1}\, x_0, \quad \text{for large } x_0. \tag{G.3}$$

Note that the expected value of X' grows linearly with the threshold x_0. That behavior is a property that is peculiar for Power-Tail distributions, see Sect. 3.3.

In the finite-buffer N-Burst/M/1/$[B]$ model, a long over-saturation period $X' \gg x_0$ causes on average $N_{(cl)}$ cell losses with

$$N_{(cl)} = \left(X' - x_0\right) \left(\Lambda_{i_0} - \nu\right) .$$

When looking at the asymptotic behavior for large B, the probability that such long over-saturation periods X' occur during time t_0 becomes very small. Thus, if overflows occur within time t_0, then – with high probability – they will be caused by only a single long over-saturation period. Hence, asymptotically for large B,

$$\mathrm{CLR}_c(t_0, B) \sim \frac{\mathbb{E}\left\{N_{(cl)}\right\}}{\lambda \cdot t_0} = \frac{\mathbb{E}\left\{X' - x_0\right\}\left(\Lambda_{i_0} - \nu\right)}{\lambda \cdot t_0} \ . \tag{G.4}$$

Together with Eq. (G.3) and with the definition (G.2) of x_0, we obtain for the asymptotic behavior of $\mathrm{CLR}_c(t_0, B)$ for large buffers

$$\mathrm{CLR}_c(B, t_0) \sim \frac{1}{\beta - 1} \cdot \frac{B}{\lambda\, t_0} = \frac{1}{\alpha - 1} \cdot \frac{1}{i_0} \cdot \frac{B}{\lambda\, t_0} \ . \tag{G.5}$$

Note that the derivation of the asymptotic behavior in (G.5) does not take into account that the over-saturation period X' could last longer than the observation interval t_0. In particular for very large buffers B such a truncation of the over-saturation period by the end of the observation interval is likely to happen.

For a mathematically rigorous derivation of the asymptotic behavior, the limit $B \to \infty$ would not be sufficient, but a simultaneous limit $t_0 \to \infty$ with some restrictions on the speed of growth of t_0 in relation to the speed of growth of B has to be considered. However, for finite t_0 and large B (G.5) provides a good approximation for the CLR_c as it is shown in Sect. 6.4.

G.2 Conditional Buffer-Overflow Ratio for Infinite Power-Tails

The situation for the infinite-buffer N-Burst/M/1 model becomes somewhat more complicated. Up to Equation (G.3) the argumentation is identical: buffer-overflows are caused by long over-saturation periods $X' > x_0$. However, in the infinite-buffer model, the queue-length can grow to level $Q_1 > B$ during the over-saturation period, where

$$Q_1 = X' \left(\Lambda_{i_0} - \nu\right) .$$

As soon as buffer-occupancy B_s is reached, all arriving cells cause overflow events. Thus, the number of overflow events during the over-saturation period is on average

$$N_{(ov1)} = \left(X' - x_0\right) \Lambda_{i_0} .$$

The critical difference to the finite-buffer model is that even after the over-saturation period X' ends, additional buffer-overflow events occur until the buffer has drained below the occupancy of B cells. In worst case, the over-saturation period with i_0 long-term active sources ends and all $i_0 - 1$ remaining sources stay active during the whole drain period, i.e. the queue-length reduces with average rate $(\nu - \Lambda_{i_0 - 1})$ and the duration of the drain period is

$$T_{(dr)} = \frac{Q_1 - B}{\nu - \Lambda_{i_0 - 1}} \cdot$$

The probability of that worst-case scenario that all $i_0 - 1$ sources remain active can be computed from the duration of the time-period during which the same $(i_0 - 1)$ sources are active that contribute their peak-rate during $T_{(dr)}$. That time-period, called Z, started an unknown time $Z_0 >= 0$ before the over-saturation period X'. Thus, the probability that all those $i_0 - 1$ sources remain active during the drain-period of duration $T_{(dr)}$ is

$$\mathbb{P}\left\{Z > Z_0 + X' + T_{(dr)} | Z > Z_0 + X'\right\} = \frac{(X' + Z_0)^{\beta[i_0-1]}}{(Z_0 + X' + T_{(dr)})^{\beta[i_0-1]}} = \left(\frac{1}{1 + \frac{T_{(dr)}}{X'+Z_0}}\right)^{\beta[i_0-1]} \quad \text{(G.6)}$$

$$\geq \left(\frac{1}{1 + \frac{T_{(dr)}}{X'}}\right)^{\beta[i_0-1]} \quad \text{(G.7)}$$

Equation (G.6) follows from the fact that time-periods with $(i_0 - 1)$ long-term active sources show a Power-Tailed duration with exponent $\beta[i_0 - 1] = (i_0 - 1) \cdot (\alpha - 1) + 1$, see Sect. 3.3. Frequently, the probability (G.6) of the worst case scenario is rather high ($> 50\%$), so worst-case can be considered to be a common scenario here.

If the assumption that the other $i_0 - 1$ sources remain active during $T_{(dr)}$ does not hold, the queue will drain more quickly, at best with rate $(\nu - \lambda)$, i.e. all of the $i_0 - 1$ long-term active sources finish very quickly:

$$T'_{(dr)} = \frac{Q_1 - B}{\nu - \lambda} \cdot$$

During the drain period of duration $T_{(dr)}$ (respectively $T'_{(dr)}$), on average another $N_{(dr)}$ ($N'_{(dr)}$) buffer-overflow events occur:

$$N_{(dr)} = T_{(dr)} \cdot \Lambda_{i_0-1} = (Q_1 - B)\frac{\Lambda_{i_0-1}}{\nu - \Lambda_{i_0-1}},$$

$$N'_{(dr)} = T'_{(dr)}\lambda = (Q_1 - B) \cdot \frac{\lambda}{\nu - \lambda} = (Q_1 - B) \cdot \frac{\rho}{1 - \rho} \cdot$$

Hence, during a single large over-saturation period X', the expected number of overflow events is in worst case

$$\mathbb{E}\left\{N_{(ov1)} + N_{(dr)}\right\} = \mathbb{E}\left\{X' - x_0\right\}\left[\Lambda_{i_0} + (\Lambda_{i_0} - \nu)\frac{\Lambda_{i_0-1}}{\nu - \Lambda_{i_0-1}}\right] =$$

$$= \mathbb{E}\left\{X' - x_0\right\}\frac{\lambda}{\rho}\frac{1}{1 - i_\Delta}. \quad \text{(G.8)}$$

The last equation follows from (5.1) and (5.3) by

$$\Lambda_{i_0+k} - \nu = \kappa \frac{b}{1 - b}\,(i_\Delta + k)\ .$$

In the case of faster draining with rate $(\nu - \lambda)$, we get instead of Eq. (G.8):

$$\mathbb{E}\left\{N_{(ov1)} + N'_{(dr)}\right\} = \mathbb{E}\left\{X' - x_0\right\}\left[\Lambda_{i_0} + (\Lambda_{i_0} - \nu)\,\frac{\rho}{1 - \rho}\right] = \mathbb{E}\left\{X' - x_0\right\}\,\frac{\Lambda_{i_0} - \lambda}{1 - \rho}\ \cdot \quad \text{(G.9)}$$

For a large buffer-size B, it is very likely that overflows during the observation period t_0 are caused by a single over-saturation period. Therefore, asymptotically for large B:

$$\frac{\mathbb{E}\left\{N_{(ov1)} + N'_{(dr)}\right\}}{\lambda \cdot t_0} \leq \mathrm{BOR}_c(t_0, B) \leq \frac{\mathbb{E}\left\{N_{(ov1)} + N_{(dr)}\right\}}{\lambda \cdot t_0} \,. \tag{G.10}$$

Putting the formulas (G.8), (G.9), (G.3), and (G.2) together and using the burstiness $b = 1-\kappa/\lambda_p$ of the sources and Eqs. (5.1) and (5.4), we get:

The conditional overflow ratio, $\mathrm{BOR}_c(B)$ in an N-Burst/M/1 model grows asymptotically linearly with increasing buffer-size B,

$$\frac{B}{\alpha-1} \cdot \frac{1}{i_\Delta} \cdot \frac{1}{\lambda\, t_0} \cdot \frac{1}{1-\rho} \;\leq\; \mathrm{BOR}_c(t_0, B) \;\leq\; \frac{B}{i_0\,(\alpha-1)} \cdot \frac{1}{\lambda\, t_0} \cdot \frac{N^{\frac{1-b}{b}\frac{1}{\rho}}}{i_\Delta\,(1-i_\Delta)}\,, \tag{G.11}$$

when the model operates in blow-up region i_0 and the burst-length distribution shows infinite Power-Tails.

The two asymptotic bounds in (G.11) are identical for $i_0 = 1$:

$$\mathrm{BOR}_c(t_0, B) \sim \frac{B}{\alpha-1} \cdot \frac{1}{i_\Delta} \cdot \frac{1}{\lambda\, t_0} \cdot \frac{1}{1-\rho} \quad \text{for} \quad i_0 = 1 \tag{G.12}$$

Note that for a mathematically rigorous derivation of the asymptotic behavior again a simultaneous limit $B, t_0 \to \infty$ has to be considered, see the end of the previous section.

The following list summarizes the assumptions or approximations that were made in the derivation of Eq. (G.11):

1. The overflow events are caused by a *single* long over-saturation period with i_0 long-term active sources. This is very likely if $\gamma(t_0, B) \ll 1$.
2. The queue is empty at the start of the long over-saturation period, which is a reasonable assumption if ρ is not too close to 1.
3. t_0 is large enough such that it contains the whole over-saturation period X' and the drain period $T_{(dr)}$.
4. Additional overflows that occur after draining to level B are neglected.
5. The distribution of the duration of the over-saturation periods shows PT behavior at least from x_0 on. Since a large buffer-size, B, causes the necessary time, x_0, for queue-growth to level B to grow, the assumption holds for larger buffer sizes.
6. The use of average rates $(\Lambda_{i_0} - \nu)$ and $(\nu - \Lambda_{i_0-1})$ (respectively $(\nu - \lambda)$) for the queue-growth and the subsequent draining neglects the variations due to the exponential distributions in the intra-burst times and service-times, as well as variations due to shorter bursts. If the time-periods X' and $T_{(dr)}$ are large, this is justified, see also Sect. 6.5.
7. Over-saturation periods caused by periods with $i_0 + 1$ or more long-term active sources are neglected. Since the duration of such over-saturation periods is Power-tailed as well, but with a larger exponent (quicker decay), they are less likely, the larger B is.

If for a fixed t_0, B is increased too far, such that $\mathrm{BOR}_c^{(PT)}(t_0, B) \geq 1$, then Assumption (3) is obviously violated.

G.3 Truncated Tails

The main reason for the asymptotically linear increase of $\mathrm{CLR}_c(B)$ and $\mathrm{BOR}_c(B)$ in Eqs. (G.5) and (G.11) is a particular property of Power-Tail distributions: the expected residual time after some known time x_0 grows linearly with the threshold x_0, see Sect. 3.3. For truncated tails as described in Sect. 3.4, a different behavior is expected.

Truncated tails of the burst-length distributions with Power-Tail Range $\overline{x_T}$ cause a truncation of the PT distribution of the over-saturation periods with i_0 long-term active sources with PT Range $\overline{x_T}/i_0$, see Eq. (5.12). Therefore, if B is large enough, such that x_0 is larger than the Power-Tail Range $\overline{x_T}/i_0$, the duration of the long over-saturation periods that cause overflows/losses is not described by the Power-Tailed region of X, but instead by the exponential drop-off with mean $\overline{x_T}/i_0$. Due to the memoryless property of the exponential distribution, the mean duration of the long over-saturation periods X' is then given by

$$\mathbb{E}\left\{X' - x_0\right\} = \frac{\overline{x_T}}{i_0}. \tag{G.13}$$

In the derivation of the asymptotic behavior of $\mathrm{CLR}_c(t_0, B)$, combining (G.13) and (G.4) provides the asymptotic behavior for truncated tails:

$$\mathrm{CLR}_c^{(TPT)}(t_0, B) \sim \frac{\overline{x_T}}{i_0} \cdot \frac{\Lambda_{i_0} - \nu}{\lambda \cdot t_0} = \frac{b\, i_\Delta}{i_0} \cdot \frac{\mathrm{MBS}}{\lambda\, t_0}\,, \tag{G.14}$$

where $\mathrm{MBS} := \lambda_p \cdot \overline{x_T}$ is the PT Range of the distribution of the number of cells in a single burst, see Sect. 3.5.

In the infinite-buffer model, the combination of (G.13), (G.10), (G.8), and (G.9) lead to the asymptotic bounds

$$\frac{b}{1-\rho} \cdot \frac{\mathrm{MBS}}{\lambda\, t_0} \;\le\; \mathrm{BOR}_c^{(TPT)}(t_0, B) \;\le\; \frac{N}{i_0} \cdot \frac{1-b}{\rho\,(1-i_\Delta)} \cdot \frac{\mathrm{MBS}}{\lambda\, t_0}\;. \tag{G.15}$$

Appendix H

Martingale Computation for Transient Overflow Probabilities

This chapter summarizes the equations that allow to compute the expected value $\mathbb{E}\tau$ and the Laplace transform $\mathbb{E}e^{-s\tau}$ of the density of τ. The detailed background and the derivations of the theorems are given in [Asmussen et al. 00a]. First, the M/M/1 queue is considered as a basic model before the algorithm is generalized to the Markov modulated setting.

H.1 The M/M/1 Queue

We assume here that $Q(t)$ is the queue length (number in system) at time t in the M/M/1 queue with arrival intensity λ and service intensity $\mu \neq \lambda$. The First Passage Time is defined as $\tau_n := \inf\{t > 0 : Q(t) = n\}$. Let $\{N_\lambda(t)\}$, $\{N_\mu(t)\}$ be independent Poisson processes with intensities λ, resp. μ, and $X(t) = N_\lambda(t) - N_\mu(t)$. Then $\{X(t)\}$ is a Lévy process with Lévy exponent

$$\kappa(\alpha) \;=\; \log \mathbb{E}\left\{e^{\alpha X(1)}\right\} \;=\; \lambda(e^{\alpha} - 1) + \mu(e^{-\alpha} - 1)$$

(see e.g. [Bertoin 90]). Let $\sigma = \inf\{t > 0 : Q(t) = 0\}$ and let

$$L(t) \;=\; \begin{cases} 0 & \text{for } t < \sigma \\ -\inf_{\sigma \leq s \leq t}(X(s) - X(\sigma)) & \text{for } t \geq \sigma \end{cases}$$

denote the local time. Then the queue length process can be generated as

$$Q(t) \;=\; Q(0) + X(t) + L(t).$$

Note that $L(t)$ is purely discontinuous, in fact equal to the number of times $\{N_\mu(t)\}$ has an event (a 'dummy service') while the queue is empty. In other words, $L(t)$ is the number of elements of the random set

$$\mathcal{M}(t) \;=\; \{s \leq t : Q(s) = 0,\, N_\mu(s) \neq N_\mu(s-)\}\,.$$

Let γ be the non–zero root of $\kappa(\gamma) = 0$ (it is easily seen that $\gamma = -\log\rho$). Furthermore, let m be the drift of the process $\{X(t)\}$:

$$m := \kappa'(0) = \lambda - \mu\,.$$

Theorem H.1.1 *It holds that*

$$\mathbb{E}\{L(\tau)\} \;=\; \frac{e^{\gamma n} - \mathbb{E}\left\{e^{\gamma Q(0)}\right\}}{1 - e^{-\gamma}} \;=\; \frac{\rho^{-n} - \mathbb{E}\left\{\rho^{-Q(0)}\right\}}{1 - \rho}, \tag{H.1}$$

$$\mathbb{E}\{\tau\} \;=\; \frac{\mathbb{E}\{X(\tau)\}}{m} \;=\; \frac{n - \mathbb{E}\{Q(0)\} - \mathbb{E}\{L(\tau)\}}{m} \tag{H.2}$$

Further, for each β, one can compute $x = x(\beta) = \mathbb{E}\left\{e^{\beta\tau}\right\}$ as the first component of the solution $(x, y) = (x(\beta), y(\beta))$ of the linear equations

$$\begin{aligned} e^{\alpha_1 n}x - (1 - e^{-\alpha_1})y &= \mathbb{E}\left\{e^{\alpha_1 Q(0)}\right\}, \\ e^{\alpha_2 n}x - (1 - e^{-\alpha_2})y &= \mathbb{E}\left\{e^{\alpha_2 Q(0)}\right\}, \end{aligned}$$

where α_1, α_2 are the roots of $\kappa(\alpha) = -\beta$:

$$\alpha_{1,2} = \log \frac{\lambda + \mu - \beta \pm \sqrt{(\lambda + \mu - \beta)^2 - 4\lambda\mu}}{2\lambda}.$$

Higher moments of the distribution of τ can also be derived from the equations in Theorem H.1.1. In particular, the second moment $\mathbb{E}\left\{\tau^2\right\}$ turns out to be useful for the computation of some approximations for $\mathbb{P}\{\tau \le T\}$, see Section 3 of [Asmussen et al. 00a]. The detailed formulas for the computation of the second moment are given in Appendix B of [Asmussen et al. 00a].

H.2 The MMPP/M/1 Queue

Let $\{J(t)\}_{t\ge 0}$ be the background Markov process, say with p states, and $\boldsymbol{Q} = (q_{ij})_{i,j=1,\dots,p}$ its intensity matrix, $\boldsymbol{\pi} = (\pi_i)_{i=1,\dots,p}$ the stationary row vector. Assume that the arrival rate is λ_i when $J(t) = i$ and the service rate μ_i. We use notation like $\boldsymbol{\Delta_\lambda}$, $\boldsymbol{\Delta_\pi}$ for the diagonal matrices with the λ_i, π_i on the diagonal.

We let $\{Q(t)\}_{t\ge 0}$ denote the queue length process and $\{X(t)\}_{t\ge 0}$ the unreflected version, that is, the difference between two MMPP's determined by the λ_i, resp. μ_i and the same driving Markov process $\{J(t)\}_{t\ge 0}$. The mean drift $m = \lim_{t\to\infty} X(t)/t$ is

$$m \;=\; \sum_{i=1}^{p} \pi_i(\lambda_i - \mu_i) \;=\; \boldsymbol{\pi}(\boldsymbol{\Delta_\lambda} - \boldsymbol{\Delta_\mu})\mathbf{1}$$

where $\mathbf{1} = (1 \dots 1)^T$ is the column vector with 1 at all entries.

Let the local time $L(t)$ and $\sigma = \inf\{t > 0 : Q(t) = 0\}$ be defined in terms of $\{Q(t)\}$, $\{X(t)\}$ precisely as in the M/M/1 case. Again, $L(t)$ is the number of elements of the random set

$$\mathcal{M}(t) \;=\; \{s \le t : Q(s) = 0,\, N_\mu(s) \ne N_\mu(s-)\}.$$

Write $\mathcal{M}_i(t) = \{s \in \mathcal{M} : J(s) = i\}$ and let $L_i(t)$ be the number of elements of $\mathcal{M}_i(t)$. Then $\{L_i(t)\}$ is the local time in state i, and $L(t) = L_1(t) + \cdots + L_p(t)$. By $\boldsymbol{L}(t)$ we denote the row vector with ith component $L_i(t)$. We have $Q(t) = X(t) + L(t)$.

Define

$$\boldsymbol{F}[\alpha] \;=\; \boldsymbol{Q} + \boldsymbol{\Delta_\lambda}(e^{\alpha} - 1) + \boldsymbol{\Delta_\mu}(e^{-\alpha} - 1).$$

The matrix $\boldsymbol{F}[\alpha]$ is the matrix c.g.f. of $\{(J(t), X(t))\}$ in the sense that the matrix with ijth element

$$\mathbb{E}\left\{e^{\alpha X(t)};\, J(t) = j\right\}_i$$

is given by $e^{t\boldsymbol{F}[\alpha]}$. For a real α, we denote by $\kappa(\alpha)$ the eigenvalue with maximal real part of $\boldsymbol{F}[\alpha]$ and by $\boldsymbol{h}^{(\alpha)}$ the corresponding right column eigenvector.

We can write

$$\det(\boldsymbol{F}[\alpha]) \;=\; \sum_{j=-p}^{p} f_j e^{j\alpha}$$

and so there exist [up to] $2p$ roots $\gamma_1, \ldots, \gamma_{2p}$. Strictly speaking, the α_j are only determined up to a multiple of $2\pi i$ but the precise choice of this multiple is immaterial for the following. The corresponding right column eigenvector are denoted by $\boldsymbol{h}^{(1)}, \ldots, \boldsymbol{h}^{(2p)}$ (do not confuse with $\boldsymbol{h}^{(\alpha)}$ defined above!).

In the following, we use the normalization $\boldsymbol{\pi}\boldsymbol{h}^{(\alpha)} = 1$ and define

$$\boldsymbol{k} \;:=\; \frac{d}{d\alpha}\boldsymbol{h}^{(\alpha)}\bigg|_{\alpha=0}.$$

Proposition H.2.1 *In the* MMPP/M/1 *model,* $\boldsymbol{k} \;=\; (\boldsymbol{Q} - \boldsymbol{1\pi})^{-1}(m\boldsymbol{I} + \boldsymbol{\Delta_\mu} - \boldsymbol{\Delta_\lambda})\boldsymbol{1}$.

Theorem H.2.2 *Define* $q_i \;=\; \mathbb{P}\{J(\tau) = i\}$, $\ell_i \;=\; \mathbb{E}\{L_i(\tau)\}$. *Then* $q_1, \ldots, q_p$, $\ell_1, \ldots, \ell_p$ *are determined as the solution of the linear equations*

$$e^{\gamma_j n}\sum_{i=1}^{p} h_i^{(j)} q_i \;-\; (1 - e^{-\gamma_j})\sum_{i=1}^{p} h_i^{(j)}\ell_i \;=\; \mathbb{E}\left\{e^{\gamma_j Q(0)} h_{J(0)}^{(j)}\right\}, \quad j = 1, \ldots, 2p. \tag{H.3}$$

Further

$$\mathbb{E}\{\tau\} \;=\; \frac{\mathbb{E}\{X(\tau)\}\, m}{} = \frac{n - \mathbb{E}\{Q(0) - \ell - c\}}{m} \tag{H.4}$$

where $\ell = \mathbb{E}\{L(\tau)\} = \sum_1^p \ell_i$ *and*

$$c \;=\; \mathbb{E}\left\{k_{J(0)}\right\} \;-\; \mathbb{E}\left\{k_{J(\tau)}\right\} \;=\; \sum_{i=1}^{p} k_i(\,\mathbb{P}\{J(0) = i\} - q_i.$$

Remark: If the MMPP contains one or more states $j \in J_0$ with $\lambda_j = 0$ (so-called OFF states), then the background process cannot possibly be in those states at time τ, so from the above definition of q_i it follows that $q_j = 0$ for $j \in J_0$.

Theorem H.2.3 *For each* β, *define* $\gamma_1(\beta), \ldots, \gamma_{2p}(\beta)$ *as the roots of* $0 = \det(\boldsymbol{F}[\gamma] + \beta\boldsymbol{I})$ *and let* $\boldsymbol{h}^{(j;\beta)}$ *be a non–zero column vector satisfying* $(\boldsymbol{F}[\gamma_j(\beta)] + \beta\boldsymbol{I})\boldsymbol{h}^{(j;\beta)} = \boldsymbol{0}$. *Then one can compute* $x = x(\beta) = \mathbb{E}\left\{e^{\beta\tau}\right\}$ *as* $x = x_1 + \cdots + x_p$ *by solving the* $2p$ *linear equations*

$$\sum_{i=1}^{p} e^{\gamma_j(\beta)n} h_i^{(j;\beta)} x_i - (1 - e^{-\gamma_j(\beta)})\sum_{i=1}^{p} h_i^{(j;\beta)} y_i \;=\; \mathbb{E}\left\{e^{\gamma_j(\beta)Q(0)} h_{J(0)}^{(j;\beta)}\right\}$$

for $x_1, \ldots, x_p$, $y_1, \ldots, y_p$.

Remark: For the OFF states $j \in J_0$ with $\lambda_j = 0$, the corresponding x_j are automatically $x_j = 0$.

H.3 Approximations

The algorithms of the last two sections allow the exact numerical computation of the probabilities $\mathbb{P}\{\tau_n \leq T\}$. However, approximation formulas can be very convenient due to their largely reduced complexity. Several approximations are listed here that follow from the asymptotic theory as discussed in [Asmussen et al. 00a].

In the stable case $\rho < 1$, [Asmussen et al. 00a] provides four different approximations:

- **(EXP):** The approximately exponential form of the density of f_τ suggest the approximation
$$\mathbb{P}\{\tau_n \leq t_0\} \approx 1 - \exp\left(\frac{t_0}{\mathbb{E}\{\tau_n\}}\right).$$
- **(CYC):** Another approximation follows from the analysis of so-called busy-cycles with duration C:
$$\mathbb{P}\{\tau_n \leq t_0\} \approx \frac{f_n t_0}{\mathbb{E}\{C\}} \cdot \tag{H.5}$$
Thereby, C is the random variable that expresses the duration of a busy cycle, and $f_n = \mathbb{P}\{\tau_n \leq C\}$ is the probability that buffer-level n is reached in a busy cycle.
In the setting of the M/M/1 queue, we have:
$$f_n = \mathbb{P}\{\tau_n \leq C\} = \frac{1-\rho}{1-\rho^n}\rho^{n-1} \quad \text{and} \quad \mathbb{E}\{C\} = \frac{1}{\lambda} + \frac{1}{\mu - \lambda}.$$
- **(GUMB):** A limit theorem for the maximum queue-length until time $T_{\tilde{m}}$ leads to a so-called Gumbel distribution, which can be used in the following way:
$$T_{\tilde{m}} = e^{\gamma\tilde{m}}\frac{\mathbb{E}\{C\}}{K} \quad \Longrightarrow \quad \mathbb{P}\{\tau_{\tilde{m}+n} > T_{\tilde{m}}\} \to e^{-e^{-\gamma n}}, \tag{H.6}$$
where K is the tail-constant in the exponential decay $f_n \sim K\exp(-\gamma n)$. Since $K = (1-\rho)/\rho$ in the M/M/1 setting, the limit theorem there simplifies to
$$T_{\tilde{m}} = \frac{1}{\mu(1-\rho)^2}\rho^{-\tilde{m}} \quad \Longrightarrow \quad \mathbb{P}\{\tau_{\tilde{m}+n} > T_{\tilde{m}}\} \to e^{-\rho^n}.$$
- **(LD):** A large deviation limit can be derived which has the following form for $\rho < 1$
$$\frac{n}{T_n} \to m^* > \kappa'(\gamma) \quad \Longrightarrow \quad \frac{1}{n}\log\,\mathbb{P}\{\tau_n \leq T_n\} \to \frac{\kappa(\theta_{m^*})}{m^*} - \theta_{m^*}, \tag{H.7}$$
with $\kappa'(\theta_{m^*}) = m^*$. In the stable M/M/1 setting, the conditions simplify:
$$m^* > \mu - \lambda, \quad \theta_{m^*} = \log\frac{m^* + \sqrt{m^{*2} + 4\lambda\mu}}{2\lambda},$$
$$\text{and} \quad \kappa(\theta_{m^*}) = \lambda\left(e^{\theta_{m^*}} - 1\right) + \mu\left(e^{-\theta_{m^*}} - 1\right).$$

The use of the limits (EXP) and (CYC) for an approximation of $\mathbb{P}\{\tau_n \leq T\}$ for given λ, ρ, n and T is straightforward. The (GUMB) limit can be used in the following way:
$$\tilde{m} = \frac{1}{\gamma}\log\frac{TK}{\mathbb{E}\{C\}}, \quad \mathbb{P}\{\tau_n \leq T\} \approx 1 - e^{-e^{-\gamma(n-\tilde{m})}}.$$
Similarly, the (LD) approximation is derived from (H.7)
$$\mathbb{P}\{\tau_n \leq T\} \approx \left[\exp\left(\frac{\kappa(\theta_{m^*})}{m^*} - \theta_{m^*}\right)\right]^n, \quad \text{with} \quad m^* = \frac{n}{T}.$$

H.4 Solving the Eigenvalue Problem

The computation of the transform of the distribution of τ in Theorem H.2.3 requires to solve a generalized eigenvalue problem for $\gamma_i(\beta)$ and $\boldsymbol{h}^{(i;\beta)}$, $i = 1, ..., 2p$:

$$(\boldsymbol{F}[\gamma_j(\beta)] + \beta \boldsymbol{I})\ \boldsymbol{h}^{(j;\beta)} = 0. \tag{H.8}$$

The same eigenvalue problem with $\beta = 0$ has to be solved for the computation of $\mathbb{E}\{\tau\}$ in Theorem H.2.2.

Using the definition of $\boldsymbol{F}[\alpha]$ and substituting $\xi_i := e^{\gamma_i(\beta)} - 1$, we get:

$$\left[\boldsymbol{Q} + \beta \boldsymbol{I} + \boldsymbol{\Delta_\lambda}\, \xi_i - \boldsymbol{\Delta_\mu} \frac{\xi_i}{\xi_i + 1}\right]\ \boldsymbol{h}^{(i;\beta)} = 0.$$

Introducing a new vector variable,

$$\widetilde{\boldsymbol{h}}_i\ :=\ \frac{1}{\xi_i + 1}\, \boldsymbol{h}^{(i;\beta)},$$

results in the new eigenvalue problem of dimension $2p$:

$$\begin{bmatrix} \boldsymbol{Q} + \beta \boldsymbol{I} & \boldsymbol{0} \\ \boldsymbol{I} & -\boldsymbol{I} \end{bmatrix} \cdot \begin{bmatrix} \boldsymbol{h}^{(i;\beta)} \\ \widetilde{\boldsymbol{h}}_i \end{bmatrix} \ =\ \xi_i \cdot \begin{bmatrix} -\boldsymbol{\Delta_\lambda} & \boldsymbol{\Delta_\mu} \\ \boldsymbol{0} & \boldsymbol{I} \end{bmatrix} \cdot \begin{bmatrix} \boldsymbol{h}^{(i;\beta)} \\ \widetilde{\boldsymbol{h}}_i \end{bmatrix}\ . \tag{H.9}$$

If all $\lambda_i > 0$, then (H.9) can easily be transformed to a standard Eigenvalue (EV) problem which can be solved by standard algorithms (QR or QZ). Even though the QZ algorithm also works for the generalized EV problem (H.9) if the matrix on the right-hand side is singular (at least one OFF state, $|J_0| \geq 1$, where $\lambda_j = 0$ for all $j \in J_0$), it is advisable to reduce the EV problem to dimension $2p - |J_0|$: In order to achieve this, (H.9) has to be multiplied from the left by a transformation matrix to erase all elements except for the diagonal element in column j of the matrix on the left-hand side for all $j \in J_0$. Thereafter, those columns and rows can be deleted and a standard eigenvalue problem of dimension $2p - |J_0|$ results. The remaining components $h_j^{(i,\beta)}$ for $j \in J_0$ are still necessary, but they can easily be determined from (H.8) since all the $\gamma_i(\beta) = \log(\xi_i + 1)$, $i = 1, ..., 2p - |J_0|$ are now known.

The case $\beta = 0$, which is needed in Theorem H.2.2, does not require any special treatment.

H.5 Laplace Transform Inversion

The Laplace Transform of a function $f(t)$,

$$\hat{f}(s) = \int_0^\infty e^{-st} f(t)\, dt,$$

can be numerically inverted using the EULER algorithm, see [Abate & Whitt 92]:

$$f(t) = \frac{2e^{at}}{\pi} \int_0^\infty \mathrm{Re}\left(\hat{f}(a + iu)\right) \cos ut\, du\,,$$

where $a > 0$ can be chosen freely, but its choice influences the numerical properties of the algorithm. Numerical integration via the trapezoidal rule (with $h = \pi/2\,t$) yields:

$$f_h(t) = \frac{e^{A/2}}{t} \left[\frac{1}{2}\,\mathrm{Re}\left(\hat{f}(A/2t)\right) + \sum_{k=1}^{\infty} (-1)^k \mathrm{Re}\left(\hat{f}\left(A/2t + \frac{k\pi}{t}\, i\right)\right)\right], \quad A := 2 \cdot t \cdot a.$$

The evaluation of the infinite sum should not be done by simple truncation, but the numerical accuracy is greatly improved by more sophisticated summation methods, such as the *Euler summation*. See [Abate & Whitt 92] for more details.

If $f(t)$ is assumed to be a density function, the Laplace transform of the cumulative distribution function $F(t)$ with $F'(t) = f(t)$ can be easily obtained from the Laplace transform of $f(t)$:

$$\hat{F}(s) = \frac{1}{s}\,\hat{f}(s)\,.$$

The computation of the cumulative distribution function via Laplace transform inversion is used in Sect. 6.3 for the transient overflow probabilities $\gamma(t_0, B)$.

Probabilistic Scaling: If small probabilities need to be obtained via the inversion of the Laplace Transform, the numerical accuracy of the result can be greatly improved by using *Probabilistic Scaling*. See [Choudhury & Whitt 97] for the details.

Appendix I

Tables

I.1 Tail Constants of PT Distributions

Power-Tail Distributions show the following asymptotic behavior (see Section 3.2):

$$R(x) \sim \frac{c_{PT}(\alpha, \bar{x})}{x^{\alpha}} =: \frac{c_{PT}^{(1)}(\alpha) \cdot \bar{x}^{\alpha}}{x^{\alpha}},$$

where $c_{PT}(\alpha, \bar{x})$ is called the *tail-constant*. $c_{PT}^{(1)}(\alpha)$ is the tail-constant of the distribution when scaled to have expected value $\bar{x} = 1$.

Thus $c_{PT}^{(1)}(\alpha)$ is the second important parameter of the Power-Tail of a distribution, in addition to the tail-exponent α . The following table provides the tail-constants for two different distribution types. First, the TPT distributions from Sect. 3.4. Secondly, for a so-called Pareto distribution, which has the reliability function:

$$R(x) = \frac{1}{\left(\frac{x}{(\alpha-1)\,\bar{x}} + 1\right)^{\alpha}}.$$

Consequently, its tail constant is

$$c_{Pareto}^{(1)}(\alpha) = (\alpha - 1)^{\alpha}.$$

For different values of α, the numerical values of the tail-constant are:

α	1.2	1.4	1.8	2.0	2.2	2.5	3.0
$c_{TPT}^{(1)}(\alpha)$	0.1290	0.2138	0.3860	0.4951	0.6320	0.9158	1.7544
$c_{Pareto}^{(1)}(\alpha)$	0.1450	0.2773	0.6692	1.0000	1.4935	2.7557	8.0000

I.2 Measurements

A set of measurements is described in Sect. 4.1 that is used to obtain parameter estimates for realistic settings of 1-Burst models. The statistical properties of some measurements together with the calibration results are presented in the following.

I.2.1 Statistical Properties

Name	Duration [s]	Mean $\overline{X}$ [μs]	$C^2(X)$	Max [ms]	95% quant. [μs]	$r(1)$
TX3	48.2	61.4	13.7	27.0	376	0.13
TX11	56.8	72.4	13.1	10.8	475	0.14
TX12	53.0	67.5	14.5	8.01	431	0.14
TX13	69.9	89.1	15.0	33.3	580	0.14
TX14	76.1	97.0	13.5	12.1	654	0.17
TX17	139	82.2	14.5	14.7	541	0.20
TX18	111	64.4	17.2	10.9	388	0.22
TX20	81.9	49.3	27.2	33.3	244	0.11
TX21	79.8	48.0	24.5	36.3	226	0.14
TX22	78.4	47.0	26.9	28.3	209	0.12
TX23	136	82.9	13.9	9.64	540	0.20
TX24	101	61.1	19.1	28.3	377	0.16
TX25	107	64.9	15.2	9.71	402	0.19
RX3	40.5	51.6	22.3	7.17	196	$6.8 \cdot 10^{-2}$
Ethernet	52.4 min	3.14 ms	3.22	342	8.15 ms	0.20

Number of Samples: The number of samples is $7.8 \cdot 10^5$ in measurements TX3 to TX14 and RX3. The remaining TX17 to TX25 contain about twice as many inter-cell times, between $1.6 \cdot 10^6$ and $1.7 \cdot 10^6$. The sample set of the Ethernet data that is discussed in [Leland et al. 94] contains 10^6 inter-packet times.

Minimum and Quantiles: The minimal inter-cell time is $2\,\mu$s in all ATM measurements. The 5%-quantile is 5μs for the TX direction and $3\,\mu$s for RX. Finally, the median is $6\,\mu$s in all TX measurements, and slightly higher for the incoming RX direction.

I.2.2 1-Burst Calibration Results

The application of the methods for 1-Burst calibration in Sect. 4.3 provides the following results for the 1-Burst parameters:

	T_0 [μs]	T_1 [ms]	κ $\left[\frac{\text{cells}}{\text{ms}}\right]$	λ_p $\left[\frac{\text{cells}}{\text{ms}}\right]$	$\overline{x_p}$ [μs]	Z [ms]	λ_0 $\left[\frac{\text{cells}}{\text{ms}}\right]$	b	$\overline{n_p}$
TX3	40.9	1.08	16.3	131	84.5	0.616	0.435	0.876	11.1
TX11	43.3	3.59	13.8	135	93.8	0.849	0.375	0.898	12.7
TX12	43.3	3.59	14.8	132	88.4	0.723	0.426	0.888	11.7
TX13	43.5	3.46	11.2	136	117	1.34	0.256	0.917	16
TX14	43.3	3.59	10.3	133	103	1.28	0.342	0.923	13.8
TX17	43.3	3.59	12.2	135	118	1.25	0.509	0.91	15.9
TX18	43.3	3.59	15.5	140	228	1.88	0.417	0.889	31.8
TX20	46.7	1.44	20.3	141	175	1.05	0.231	0.856	24.6
TX21	46.7	1.44	20.8	141	166	0.973	0.339	0.852	23.4
TX22	46.7	1.44	21.3	140	169	0.959	0.28	0.848	23.7
TX23	44.2	3.48	12.1	135	111	1.19	0.519	0.911	15
TX24	46.7	1.44	16.4	137	127	0.969	0.434	0.881	17.4
TX25	44.2	3.48	15.4	138	140	1.15	0.502	0.888	19.3
RX3	49.8	0.823	19.4	135	117	0.706	0.142	0.856	15.8

Note that the choice of the thresholds T_0 and T_1 influences the result. In particular the choice of T_1 is not clear in all measurements, some do not show as clear plateaus as the simulations of 1-Burst models in Sect. 4.3.3. This is an indication that the OFF-time distribution is not quite exponential in the data. However, as Sect. 5.7 indicates, the actual shape of the OFF-time distribution is not critical.

The LRD estimators of Sect. 4.4 indicate values between $\alpha = 1.4$ and $\alpha = 1.8$, see also [Rimkus 99] and [Gogl 00]. Note that the background Poisson rate λ_0 only accounts for less than 5% of the overall average rate κ. The peak cell-rate λ_p during ON periods is by a factor 7 to 13 larger than the average rate.

The Ethernet data that appears in the previous section does not appear to be well modeled by a 1-Burst process, see Sect. 4.5.

I.3 Solution of Non-Linear Equation

In Section 6.4.3, it is derived that the conditional Buffer-Overflow Ratio shows a minimum at some buffer-size B_{min} which can be computed by Eq. (6.20). However, as part of this solution, the value of $z_0(\beta)$ is necessary which is defined only implicitly as the root of Eq. (6.19):

$$\exp(-z_0) = \frac{1}{1 + \frac{\beta}{\beta-1} z_0} \cdot$$

z_0 depends only on the PT exponent β of the over-saturation periods of the switch. Numerical values are listed in the following table:

β	1.2	1.4	1.6	1.8	2.0	2.2	2.6
$z_0(\beta)$	2.918	2.138	1.721	1.450	1.256	1.111	0.9035

References

[Abate & Whitt 92] Joseph Abate and Ward Whitt: *The Fourier-series method for inverting transforms of probability distributions.* Queueing Systems **10**, pp. 5-88, 1992.

[Adler et al. 98] Robert Adler, Raisa Feldman, and Murad Taqqu (eds.): *A PRACTICAL GUIDE TO HEAVY TAILS: STATISTICAL TECHNIQUES AND APPLICATIONS.* Birkhäuser, Boston, 1998.

[Andersen & Nielsen 98] A. Andersen and B.F. Nielsen: *A Markovian approach for modeling packet traffic with long-range dependence.* IEEE J. Sel. Areas Comm. **16**, 719–732, 1998.

[Arvidson & Karlsson 99] Ake Arvidson and Pär Karlsson: *On Traffic Models for TCP/IP.* In P. Key and D. Smith (eds.), 'Teletraffic Engineering in a Competitive World, Vol 3A', pp. 457-466. Elsevier Science B.V., 1999.

[Asmussen et al. 00a] Soeren Asmussen, Manfred Jobmann, and Hans-Peter Schwefel: *Exact buffer overflow calculations for queues via martingales.* Submitted to Queueing Systems.

[Asmussen et al. 00b] Soeren Asmussen, Pascal Frantz, Manfred Jobmann, and Hans-Peter Schwefel: *Large deviations and fast simulation in the presence of boundaries.* To appear in Stochastic Processes and Applications.

[Beran 94] Jan Beran: *STATISTICS FOR LONG-MEMORY PROCESSES.* Chapman & Hall. 1994.

[Bertoin 90] J. Bertoin: *LÉVY PROCESSES.* Cambridge University Press. 1990.

[Bini & Meini 96] Dario Bini and Beatrice Meini: *On the solution of a nonlinear Matrix equation arising in queueing problems.* SIAM J. Matrix Anal. Appl., **17**, p. 906-926, 1996.

[Cahn 98] Robert Cahn: *Wide Area Network Design – Concepts and Tools for Optimization.* Morgan Kaufmann Publishers Inc., 1998.

[Chaudhry & Templeton 83] M. Chaudhry and J. Templeton: *A FIRST COURSE IN BULK QUEUES.* John Wiley & Sons. 1983.

[Choudhury & Whitt 97] Gagan Choudhury and Ward Whitt: *Probabilistic Scaling for the Numerical Inversion of Non-Probability Transforms.* INFORMS Journal on Computing **9**, No. 2, pp 175-184. 1997.

[Coevering 95] M. van de Coevering: *Computing transient performance measures for the M/M/1 queue.* OR Spektrum, Nr. 17, pp. 19-22. 1995.

[Cohen, 82] Jacob Cohen: *THE SINGLE SERVER QUEUE.* North-Holland Publishing Company. Amsterdam, 1982.

[Crovella & Bestavros 96] Mark Crovella and Azer Bestavros: *Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes.* Proceedings of the ACM Sigmetrics, 1996.

[Daley 00] Daryl Daley: Personal Correspondence. September 2000.

[Dumas & Simonian 00] V. Dumas and A. Simonian: *Asymptotic bounds for the fluid queue fed by sub-exponential ON/OFF sources.* Advances in Applied Probability **32**, No 1. March 2000.

[Fan & Georganas 96] Yanhe Fan and Nicolas Georganas: *Performance Analysis of ATM Switches with Self-Similar Input Traffic.* Technical Report, Dept. of Electrical Engineering, University of Ottawa. December 1996.

[Feldman & Whitt 98] Anja Feldman and Ward Whitt: *Fitting mixtures of exponentials to long-tail distributions to analyze network performance models.* Performance Evaluation **31**, pp. 245-279. 1998.

[Fiorini et al. 95] Pierre Fiorini, Lester Lipsky, Wen-Jung Hsin, and Appie van de Liefvoort: *Auto-Correlation of Lag-k For Customers Departing From Semi-Markov Processes.* Technical Report TUM-I9506, Institut für Informatik, Technische Universität München. July 1995.

[Fiorini 97] Pierre Fiorini: *Modeling Telecommunication Systems with Self-Similar Data Traffic.* PhD Thesis, University of Connecticut, September 1997.

[Fischer 00] Volker Fischer: *Evolutionary Design of Corporate Networks under Uncertainty.* Dissertation, Technische Universität München. March 2000.

[Garg et al. 92] Sharad Garg, Lester Lipsky, and Maryann Robbert: *The Effect of Power-Tail Distributions on the Behavior of Time Sharing Computer Systems.* 1992 ACM SIGAPP Symposium on Applied Computing, Kansas City, MO. March 1992.

[Gogl 98] Helmut Gogl: *Validation of the N-Burst Model.* Technical Report #10, Project 'Verkehrsstatistiken und Anwendungsprofile in ATM-Netzen', Institut für Informatik, Technische Universität München, Juni 1998.

[Gogl 00] Helmut Gogl: *Measurement and Characterization of Traffic Streams in High-Speed Wide Area Networks.* Dissertation, Technische Universität München. October 2000.

[Graham 81] Alexander Graham: *KRONECKER PRODUCTS AND MATRIX CALCULUS WITH APPLICATIONS.* Ellis Horwood. 1981.

[Greiner et al. 99] Michael Greiner, Manfred Jobmann, and Lester Lipsky: *The Importance of Power-tail Distributions for Telecommunication Traffic Models.* Operations Research **47**, No. 2, pp 313-326, March 1999.

[Greiner et al. 99b] Michael Greiner, Manfred Jobmann, and Claudia Klüppelberg: *Telecommunication traffic, queueing models and subexponential distributions.* Queueing Systems **33**, pp. 125-152. 1999

[Gribble et al. 98] Steven Gribble, Gurmeet Manku, Drew Roselli, and Eric Brewer: *Self-Similarity in File-Systems.* Proceedings of the ACM Sigmetrics, 1998.

[Guerin et al., 00] C.A. Guerin, H. Nyberg, O. Perrin, S. Resnick, H. Rootzen, and C. Starica: *Empirical Testing of the Infinite Source Poisson Data Traffic Model.* Preprint, 2000.

[Heffes 80] H. Heffes: *A class of data traffic processes - covariance function characterizationand related queueing results.* Bell Syst. Tech. J. 59, pp. 897-930. 1980.

[Heyman & Sobel 82] D. Heyman and M. Sobel: *STOCHASTIC MODELS IN OPERATIONS RESEARCH, VOL. I.* McGraw-Hill. 1982.

[Heyman et al. 97] D. Heyman, T. Lakshman, and A. Neidhardt: *A New Method for Analysing Feedback-Based Protocols with Applications to Engineering Web Traffic over the Internet.* Proceedings of the ACM Sigmetrics, pp. 24-38, 1997.

[Heyman 00] Daniel Heyman: *Performance Implications of Very Large Service-Time Variances.* Performance Evaluation **40**, pp. 47-70. 2000.

[Heyman 00b] Daniel Heyman: *Estimation of MMPP Models of IP Traffic.* Technical report, AT&T Labs. 2000

[Hill 75] B. Hill: *A simple general approach to inference about the tail of a distribution.* Annals of Statistics **3**, pp. 1163-1174. 1975.

[Jelenkovic & Lazar 97] Predrag Jelenkovic and Aurel Lazar: *Multiplexing On-Off Sources with Subexponential On Periods: Part II.* In V. Ramaswami and P. Wirth (eds.), 'Teletraffic Contributions for the Information Age, Vol 2A', pp. 965-974. Elsevier Science B.V., 1997.

[Jelenkovic & Lazar 99] Predrag Jelenkovic and Aurel Lazar: *Asymptotic results for multiplexing subexponential on-off sources.* Advances of Applied Probability **31**, No. 2. 1999.

[Kella & Whitt 92] O. Kella and W. Whitt: *Useful martingales for stochastic storage processes with Lévy input.* Journal of Applied Probability **29**, pp 396-403, 1992.

[Kesidis et al. 93] George Kesidis, Jean Walrand, and Cheng-Shang Chang: *Effective Bandwidths for Multiclass Markov Fluids and Other ATM Sources.* IEEE/ACM Transactions on Networking, Vol.1 , No.4, August 1993.

[Kleinrock 75] Leonard Kleinrock: *QUEUEING SYSTEMS, Volume I: Theory.* John Wiley & Sons,New York, 1975.

[Klinger 97] Wolfgang Klinger: *On the Convergence of Sums of Power-Tail Samples to Their α-Stable Distribution.* M.Sc. Thesis, Department of Computer Science, University of Connecticut. August 1997.

[Krieger et al. 98] Udo Krieger, Valeri Naoumov, and Dietmar Wagner: *Analysis of a Finite FIFO Buffer in an advanced Packet-Switched Network.* IEICE Trans. Commun., VOL. E00-B, No.5, May 1998.

[Krieger & Naumov 99] Udo Krieger and Valery Naumov: *Analysis of a delay-loss System with a superimposed Markovian Arrival Process and state-dependent Service Times.* Proceedings of the MMB conference, 1999.

[Krishnan 96] K. Krishnan: *A new class of performance results for a fractional Brownian traffic model.* Queueing Systems 22, pp. 277-285, 1996.

[Kulkarni & Li 98] Lalita A. Kulkarni and San-qi Li: *Transient Behaviour of Queuing Systems with Correlated Traffic.* Commun.Statist.-Stochastical Models, No. 14, Vol. 4, pp. 933-978, 1998.

[Latouche & Ramaswami 93] Guy Latouche and V. Ramaswami: *A Logarithmic Reduction Algorithm for Quasi-Birth-Death Processes.* Journal of Applied Probability **30**, pp. 650-674, 1993.

[Latouche & Ramaswami 99] G. Latouche and V. Ramaswami: *INTRODUCTION TO MATRIX ANALYTIC METHODS IN STOCHASTIC MODELING.* ASA-SIAM Series on Statistics and Applied Probability 5. 1999.

[Leguesdron et al. 93] P. Leguesdron, J. Pellaumail, G. Rubino, and B. Sericola: *Transient Analysis of the M/M/1 Queue.* Adv. Appl. Prob., Nr. 25, pp. 702-713, 1993.

[Leland & Ott 86] Will Leland and Teunis Ott: *Load-Balancing Heuristics and Process Behavior.* Proceedings of ACM SigMetrics 1986, in: Performance Evaluation Review, Vol. 14, No 1. May 1986.

[Leland et al. 94] Will Leland, Murad Taqqu, Walter Willinger, and Daniel V. Wilson: *On the Self-Similar Nature of Ethernet Traffic (Extended Version).* Proc. of IEEE/ACM Trans. on Networking, **2**, 1. February 1994.

[Li et al. 98] Jian-Min Li, Indra Widjaja, and Marcel Neuts: *Congestion detection in ATM networks.* Performance Evaluation **34**, pp. 147-168, 1998.

[Lipsky 92] Lester Lipsky: *QUEUEING THEORY: A Linear Algebraic Approach.* MacMillan Publishing Company, New York, 1992.

[Lipsky & Liefvoort 95] Lester Lipsky and Appie van de Liefvoort: *Transformation of the Kronecker Product of Identical Servers to a Reduced Product Space.* Technical Report TUM-I9505, Institut für Informatik, Technische Universität München. February 1995.

[Lipsky et al. 97] Lester Lipsky, Mark Crovella, Pierre Fiorini, Michael Greiner, John Hatem, Manfred Jobmann, and Hans-Peter Schwefel: *How To Model Telecommunications (and Other) Systems Where Self-Similar (Power-Tail) Behavior is Observed: (Background Review and Research Proposal).*, Technical Report, CSE/BRC01297, University of Connecticut. May, 1997.

[Lipsky et al. 99a] Lester Lipsky, Pierre Fiorini, and Hans-Peter Schwefel: *Analytical Models of Performance in Telecommunication Systems, Based on On-Off Traffic Sources with Self-Similar Behavior.* 7th International Conference on Telecommunication Systems Modeling and Analysis. Nashville, March 1999.

[Lipsky et al. 99b] Lester Lipsky, Hans-Peter Schwefel, and Tianyun Zhang: *Simulation Techniques and Moduls.* Technical Report #E1, Project 'Pufferdimensionierung in ATM-Weitverkehrsnetzen mit selbstähnlichen Verkehrsströmen (...)', Technische Universität München & University of Connecticut, July 1999.

[Lipsky & Schwefel 00] Lester Lipsky and Hans-Peter Schwefel: *Variations of Packet Size In The Analytic N-Burst Model and Comparisons with Other Approximations.* Technical Report #E2, Project 'Pufferdimensionierung in ATM-Weitverkehrsnetzen mit selbstähnlichen Verkehrsströmen (...)', Technische Universität München & University of Connecticut, August 2000.

[Lucantoni 93] David Lucantoni: *The BMAP/G/1 Queue: A Tutorial.* In Donatiello and Nelson, eds., 'Models and Techniques for Performance Evaluation of Computer and Communications Systems'. Springer Verlag, 1993.

[MathWorks 96] *MATLAB, Documentation, Version 5.2.* MathWorks, Inc. 1996.

[Meier-Hellstern & Fischer 92] Kathy Meier-Hellstern and Wolfgang Fischer: *MMPP Cookbook.* Performance Evaluation **18**, p. 149-171. 1992.

[Meini 97] Beatrice Meini: *New convergence results on functional iteration techniques for the numerical solution of M/G/1 type Markov chains.* Numerische Mathematik **78**, pp. 39-58. 1997.

[Melamed 91] B. Melamed: *TES: A class of methods for generating autocorrelated uniform variates.* ORSA Journal on Computing **3**, no. 4, pp. 317-329. 1991.

[Neuts 81] Marcel Neuts: *MATRIX-GEOMETRIC SOLUTIONS IN STOCHASTIC MODELS.* John Hopkins University Press, London, 1981.

[Neuts 86] Marcel Neuts: *The caudal characteristic curve of queues.* Adv. Appl. Prob., **18**, pp. 221-254. 1986.

[Neuts 89] Marcel Neuts: *STRUCTURED STOCHASTIC MATRICES OF M/G/1 TYPE AND THEIR APPLICATION.* Dekker, New York, 1989.

[Norros 95] I. Norros: *On the Use of Fractional Brownian Motion in the Theory of Connectionless Networks.* IEEE JSAC **13**, no. 6, pp. 953-962. 1995.

[Padhye et al. 98] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose: *Modeling TCP Throughput: A Simple Model and its Empirical Validation.* ACM SIGCOMM, 1998.

[Pancha et al. 97] P. Pancha, M. Veeraraghavan, and S. Rai: *On MPEG video-based multiparty communication: Part I, Source Characterization and equivalent bandwidth computation.* Technical Report, Bell Laboratories, Holmdel. 1997.

[Paxson & Floyd 95] Vern Paxson and Sally Floyd: *Wide area traffic: The failure of Poisson Modeling.* IEEE/ACM Transactions on Networking **1**, No. 3. June, 1995.

[Ren & Kobayashi 95] Qiang Ren, Hisashi Kobayashi: *Transient Solutions for the Buffer Behaviour in Statistical Multiplexing.* Performance Evaluation, No. 23, pp. 65-87, 1995.

[Rimkus 99] Henning Rimkus: *Bewertung von Quellenmodellen in Telekommunikationssystemen.* Diplomarbeit, TU München. July 1999.

[Robert & Boudec 97] S. Robert and J.-Y. Le Boudec: *New models for self-similar traffic.* Performance Evaluation **30**, no. 1-2, pp. 57-68. July 1997

[Saad 96] Yousef Saad: *ITERATIVE METHODS FOR SPARSE LINEAR SYSTEMS.* PWS Publishing Company. Boston, 1996.

[Samorodnitsky & Taqqu 94] Gennady Samorodnitsky and Murad S. Taqqu: *STABLE NON-GAUSSIAN RANDOM PROCESSES: Stochastic Models with Infinite Variance.* Chapman and Hall, New York 1994.

[Schwefel 97] Hans-Peter Schwefel: *Modeling of Packet Arrivals Using Markov Modulated Poisson Processes with Power-Tail Bursts.* Diplomarbeit. Institut für Informatik, Technische Universität München. August 1997.

[Schwefel et al. 97] Hans-Peter Schwefel, Helmut Gogl, and Michael Greiner: *Model Calibration.* Techincal Report #8, Project 'Verkehrsstatistiken und Anwendungsprofile in ATM Netzen', Technische Universität München, December 1997.

[Schwefel99a] Hans-Peter Schwefel: *The Impact of Self-Similar Traffic on Performance in Stationary Switch-Models.* Technical Report #2, Project 'Pufferdimensionierung in ATM-Weitverkehrsnetzen mit selbstähnlichen Verkehrsströmen (...)', Technische Universität München, January 1999.

[Schwefel & Lipsky 99b] Hans-Peter Schwefel and Lester Lipsky: *Buffer Size Issues in the Presence of Self-Similar Traffic.* 3rd IFIP Workshop on Traffic Management and Design of ATM Networks. London, April 1999.

[Schwefel & Lipsky 99c] Hans-Peter Schwefel and Lester Lipsky: *Performance Results For Analytic Models of Traffic In Telecommunication Systems, Based on Multiple ON-OFF Sources with Self-Similar Behavior.* In P. Key and D. Smith (eds.), 'Teletraffic Engineering in a Competitive World, Vol 3A', pp. 55-66. Elsevier, 1999.

[Schwefel & Lipsky 99d] Hans-Peter Schwefel and Lester Lipsky: *Impact of aggregated, self-similar ON/OFF traffic on delay in stationary queueing models.* SPIE's International Symposium on Voice, Video, and Data Communications. Boston, September 1999.

[Schwefel 99e] Hans-Peter Schwefel: *Transient Behavior of Performance Models with Self-Similar Network Traffic.* Technical Report #3, Project 'Pufferdimensionierung in ATM-Weitverkehrsnetzen mit selbstähnlichen Verkehrsströmen (...)', Technische Universität München, July 1999.

[Schwefel & Lipsky 00] Hans-Peter Schwefel and Lester Lipsky: *Impact of aggregated, self-similar ON/OFF traffic on delay in stationary queueing models (extended version).* To appear in Performance Evaluation, November 2000.

[Schwefel 00a] Hans-Peter Schwefel: *How to avoid poor QoS for self-similar network traffic.* Final Report, Project 'Pufferdimensionierung in ATM-Weitverkehrsnetzen mit selbstähnlichen Verkehrsströmen (...)', Technische Universität München, February 2000.

[Schwefel 00b] Hans-Peter Schwefel: *Behavior of TCP-like elastic traffic at a buffered bottleneck router.* To appear in Proceedings of IEEE Infocom. 2001.

[Sericola 98] B. Sericola: *Transient Analysis of Stochastic Fluid Models.* Performance Evaluation, Nr. 32, pp. 245-263, 1998.

[Stevens 94] W. Stevens: *TCP/IP ILLUSTRATED, VOL.1, THE PROTOCOLS.* Addison-Wesley, 1994.

[Stewart et al. 92] Wiliam Stewart, B. Philippe, and Y. Saad: *Numerical Methods in Markov Chain Modelling.* Operations Research, Vol. 40, No. 6, pp. 1156-1179, 1992.

[Tanaka et al. 95] Takeshi Tanaka, On Hashida, Yukio Takahashi: *Transient Analysis of Fluid Models for ATM Statistical Multiplexer.* Performance Evaluation, No. 23, pp. 145-162, 1995.

[Tsybakov & Georganas 97] Boris Tsybakov and Nicolas Georganas: *On Self-Similar Traffic in ATM Queues: Definitions, Overflow Probability Bound, and Cell Delay Distribution.* IEEE/ACM Transactions on Networking, Vol. 5, No. 3, pp307-409, June 1997.

[Veres & Boda 00] Andras Veres and Miklos Boda: *The chaotic nature of TCP congestion control.* Proceedings of the IEEE Infocom, 2000.

[Vorst & Golub 97] Henk van der Vorst and Gene Golub: *150 Years old and still Alive: Eigenproblems.* In I.S. Duff and G.A. Watson (eds), 'The State of the Art in Numerical Analysis', pp. 93-119. Clarendon Press, Oxford, 1997.

[Willinger et al. 95] Walter Willinger, Murad Taqqu, Robert Sherman, and Daniel Wilson: *Self-Similarity Through High-Variability: Statistical Analysis of Ethernet LAN Traffic at the Source Level (Extended Version).* Proceedings of the ACM Sigcomm, 1995.