

Scaling Properties in the Stochastic Network Calculus

A Dissertation

Presented to

the faculty of the School of Engineering and Applied Science

University of Virginia

In Partial Fulfillment

of the requirements for the Degree

Doctor of Philosophy

Computer Science

by

Florin Ciucu

August 2007

© Copyright August 2007

Florin Ciucu

All rights reserved

Approvals

This dissertation is submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy
Computer Science

Florin Ciucu

Approved:

Jörg Liebeherr (Co-Advisor)

John A. Stankovic (Chair)

Almut Burchard (Co-Advisor)

Stephen G. Wilson

Stephen D. Patek

Accepted by the School of Engineering and Applied Science:

James H. Aylor (Dean)

August 2007

Abstract

Modern networks have become increasingly complex over the past years in terms of control algorithms, applications and service expectations. Since classical theories for the analysis of telephone networks were found inadequate to cope with these complexities, new analytical tools have been conceived as of late. Among these, the stochastic network calculus has given rise to the optimism that it can emerge as an elegant mathematical tool for assessing network performance.

This thesis argues that the stochastic network calculus can provide new analytical insight into the scaling properties of network performance metrics. In this sense it is shown that end-to-end delays grow as $\Theta(H \log H)$ in the number of network nodes H , as opposed to the $\Theta(H)$ order of growth predicted by other theories under simplifying assumptions. It is also shown a comparison between delay bounds obtained with the stochastic network calculus and exact results available in some product-form queueing networks.

The main technical contribution of this thesis is a construction of a statistical network service curve that expresses the service given to a flow by a network as if the flow traversed a single node only. This network service curve enables the proof of the $\mathcal{O}(H \log H)$ scaling of end-to-end delays, and lends itself to explicit numerical evaluations for a wide class of arrivals. The value of the constructed network service curve becomes apparent by showing that, in the stochastic network calculus, end-to-end delay bounds obtained by adding single-node delay bounds grow as $\mathcal{O}(H^3)$.

Another technical contribution is the application of supermartingales based techniques in order to evaluate sample-path bounds in the stochastic network calculus. These techniques

are suitable to arrival processes with stationary and independent increments, and improve the performance bounds obtained with existing techniques.

Acknowledgments

I would like to express my deep gratitude to those who sustained the development of this dissertation with their professionalism, inspiration and encouragement.

This dissertation would never have been possible without the continuous help and support of my advisors Jörg Liebeherr and Almut Burchard. They were instrumental to the development of this dissertation from its inception to its final polished form. Their intellectual contribution is present throughout, as the majority of the results and views presented herein are the product of our collaboration. I will be forever indebted to Jörg and Almut for their professionalism in educating, mentoring and inspiring me. I feel fortunate to have been their student; their guidance, patience and adherence to high standards have played a central role in my evolution as a scholar. I can only hope that I will have a similarly great influence on others.

I would also like to thank Stephen D. Patek, John A. Stankovic and Stephen G. Wilson, whose thoughtful insights and suggestions were invaluable to this dissertation. I feel honored to have had them on my committee.

I am profoundly grateful to those professors who recognized my interest in science and encouraged my early intellectual evolution. In particular, I thank my math teachers Nicolae Grigorescu and Marin Toloși, and my informatics teacher Vasile Voinea. Furthermore, I thank my undergraduate advisor Tudor Bălănescu at the University of Bucharest, and my graduate advisors Brian L. Mark and Robert Simon at George Mason University.

Finally, I feel blessed for the love of my mom, dad and sister. Their inspiration and continuous support empower me to pursue my professional goals. I humbly dedicate this dissertation to them.

Contents

<i>List of Figures</i>	xi
<i>List of Tables</i>	xii
1 Introduction	1
1.1 Key Issues in Analyzing Data Networks	2
1.2 Theories for Network Analysis	4
1.3 Thesis Statement and Contributions	9
1.4 Thesis Structure	13
2 Background on the Deterministic Network Calculus	15
2.1 Network Model	15
2.2 Deterministic Envelope	17
2.3 Deterministic Service Curve	19
2.4 Single-Node Performance Bounds	25
2.5 Deterministic Network Service Curve	27
3 State of the Art in the Stochastic Network Calculus	33
3.1 The Need for a Probabilistic Extension of Network Calculus	33
3.2 Statistical Envelope	36
3.2.1 Statistical envelope as non-random function	36
3.2.2 Statistical envelope as random process	43
3.3 Statistical Service Curve	47

3.3.1	Statistical service curve as non-random function	47
3.3.2	Statistical service curve as random process	52
3.4	Single-Node Performance Bounds	54
3.4.1	With lower bound approximation	55
3.4.2	With Boole’s inequality	57
3.4.3	With a-priori busy period bounds and Boole’s inequality	60
3.4.4	With sample-path statistical envelope	61
3.5	The Problem of Statistical Network Service Curve	64
4	The Construction of a Statistical Network Service Curve	71
4.1	Statistical Envelope	72
4.2	Statistical Service Curve	75
4.3	Single-Node Performance Bounds	78
4.4	Statistical Network Service Curve	82
5	Scaling Properties of End-to-End Delay Bounds	88
5.1	Network with Fluid Service	95
5.2	Network with Packetized Service	101
5.2.1	Upper bound	105
5.2.2	Lower bound	109
5.3	Numerical Examples	118
5.3.1	Markov-modulated On-Off processes	118
5.3.2	Comparison between upper and lower bound	123
6	Accounting for Statistical Independence	126
6.1	Statistical Envelope	127
6.2	Statistical Service Curve	128
6.3	Single-Node Performance Bounds	132
6.4	Statistical Network Service Curve	135

6.5	The Special Case of Stationary Processes with Independent Increments	138
7	Comparison of Delay Bounds with Exact Results	144
7.1	The Single-Node Case	145
7.1.1	The M/M/1 queue	145
7.1.2	The M/M/1 queue with priorities	149
7.1.3	The M/D/1 queue	151
7.2	The Multi-Node Case	153
7.2.1	Independent arrivals and service	157
7.2.2	Correlated arrivals, independent service	159
7.2.3	Independent arrivals, correlated service	162
7.3	Numerical Examples	164
7.3.1	The single-node case	164
7.3.2	The multi-node case	165
8	Conclusions and Future Work	171
8.1	Conclusions	171
8.2	Future Work	173
	<i>Glossary of Notation</i>	175
	<i>Index</i>	179
	<i>Bibliography</i>	181

List of Figures

1.1	A network model	2
2.1	A network with cross traffic	16
2.2	An envelope function with two sample arrivals	18
2.3	A linear and time-invariant system	21
2.4	A service curve representation	22
2.5	Service curves in series	27
2.6	A network with cross traffic and leaky-bucket envelopes	29
3.1	Bernoulli traffic source with peak and average-rate envelopes	34
3.2	The cell and burst regions for Markov-modulated traffic	38
4.1	Statistical service curves in series	83
5.1	The z -quantile $w^{net}(z)$	90
5.2	A statistical packetizer	103
5.3	A Markov-modulated On-Off traffic model.	119
5.4	End-to-end delay bounds for Markov-modulated On-Off processes (low burstiness and equal traffic mix) as a function of the number of flows	120
5.5	End-to-end delay bounds for Markov-modulated On-Off processes (high burstiness and equal traffic mix) as a function of the number of flows	121
5.6	End-to-end delay bounds for Markov-modulated On-Off processes (low burstiness and different traffic mix) as a function of the number of flows	122

5.7	End-to-end delay bounds for Markov-modulated On-Off processes as a function of the number of nodes	122
5.8	End-to-end delay bounds in tandem networks as a function of the number of nodes	123
5.9	End-to-end delay bounds in tandem networks as a function of the utilization factor	124
7.1	Delay bounds in M/M/1 and M/D/1 queues	165
7.2	Expected delay bounds for the lowest priority flow in an M/M/1 queue as a function of the utilization factor	166
7.3	Comparison of network calculus bounds with exact results in an M/M/1 network by varying the number of nodes	167
7.4	Comparison of network calculus bounds with exact results in an M/M/1 network by varying the utilization factor	167
7.5	The impact of relaxing the statistical independence assumptions of arrivals and service in an M/M/1 network by varying the number of nodes	168
7.6	The impact of relaxing the statistical independence assumptions of arrivals and service in an M/M/1 network by varying the utilization factor	169
7.7	The impact of using a fluid service model in an M/M/1 network by varying the utilization factor	169

List of Tables

2.1	Operations in the usual and the $(\min, +)$ algebra	21
5.1	Parameters of a Markov-modulated On-Off Markov process.	120

Chapter 1

Introduction

For almost two decades *network calculus* has evolved as a new theory for the performance analysis of communication networks. The theory was conceived in 1991 by Cruz as a *deterministic* network calculus in two seminal works [37, 38]. One year later Kurose in [73], and Chang in [25] published the first extensions of Cruz's works in a probabilistic setting, that marked the debut of *stochastic* network calculus. Ever since, many researchers have contributed to the development of network calculus, in both its deterministic and probabilistic directions.

While network calculus is a relatively recent theory, the problem of network analysis that has motivated the calculus has a much longer history. Studies on network analysis date as far back as 1909 and 1917 when Erlang published his seminal works (see Brockmeyer *et al.* [19]) on the analysis of telephone networks. Erlang's work represented the foundation for *queueing theory*, which has become an important branch of applied mathematics. Among its many applications, queueing theory was instrumental in dimensioning telephone networks.

The first studies on data networks analysis were conducted by Kleinrock [67] and closely preceded the appearance of the Internet in the late 1960s. These studies undertook a queueing theoretical approach and were based on earlier results developed by Jackson [58]. Subsequent research in this direction led to the development of *queueing networks theory*, as an extension of queueing theory to multiple queues, and which has become an influential framework for

networks analysis (see Bertsekas and Gallager [10]).

Queueing networks theory is generally restricted to the Poisson traffic model which has been shown to accurately represent telephone networks traffic characterized by low variability. With the deployment of voice and video applications in the Internet, characterized by a high variability of traffic, the Poisson model was found to be inadequate to describe traffic in modern data networks. To improve the accuracy in predicting network performance metrics, new theories for network analysis were conceived in the 1980s and 1990s, such as the theory of effective bandwidth and the network calculus.

For the rest of this introduction, we first discuss some of the difficulties that arise in data network analysis. Then we give a brief overview of existing theories for network analysis. We then present the statement and the main contributions of this thesis. Finally we map out how the rest of the thesis will be organized.

1.1 Key Issues in Analyzing Data Networks

In contrast to telephone networks, there are two factors that significantly complicate the analysis of data networks. The first is that most data networks are based on a packet switching technology, as opposed to the circuit switching technology used in telephone networks. The second is that traffic in data networks is more complex than the simple traffic in telephone networks.

Let us consider the network model depicted in Figure 1.1. The network consists of nodes (packet switches) with fixed capacity. Flows carrying data traffic traverse the network, and each node may be transited by more than one flow.




Figure 1.1: A network model

With packet switching, the data of each flow is divided into packets, and each node serves the packets of the incoming flows according to some scheduling algorithm. For example, if the nodes implement static priority schedulers, then a node can serve packets belonging to a flow as long as there are no incoming packets belonging to flows of higher priorities. Different flows at a node may thus receive different service rates. Also, each flow may receive different service rates at different nodes. This sharing of the capacity of a node by the packets of multiple flows is usually called *statistical multiplexing* [10].

In comparison to other switching technologies, e.g., circuit switching which allocates a fixed service rate to each flow at a node, statistical multiplexing results in a better utilization of network resources. Indeed, since data flows usually transmit at their peak rate only a fraction of the time, statistical multiplexing can be up to 100 times more efficient than circuit switching (see Roberts [98]).

A consequence of packet switching is that network nodes require the availability of *buffers* to temporarily store packets when the amount of packets to be served exceeds the nodes' capacity. As effects of buffering, packets may experience different queueing delays in the buffers, and the performance of end-to-end flows may be highly variable. Buffering represents thus a key challenge in network performance analysis.

Another challenge in analyzing packet-switched data networks stems from the characteristics of traffic. Unlike traffic in telephone networks which can be adequately modelled with Poisson processes, traffic in data networks is more complex as it exhibits high variability or correlations. The high variability of traffic is usually referred to as traffic *burstiness*, and can be determined for instance by 'clustered' interarrival times followed by long idle periods. In general, analyzing traffic models for bursty traffic is more complicated than analyzing traffic described with Poisson models.

1.2 Theories for Network Analysis

In this section we provide an overview of analytical methods to analyze performance measures (e.g. delays, backlog) and statistical multiplexing in packet networks.

Queueing theory played an important role in the justification of packet-switching technology in the early 1960s. By that time, queueing theory was already a well established theory, and widely applied for analyzing circuit-switched telephone networks [46,47]. In the simplest form, an output link at a packet switch is modelled as an M/M/1 queue. The underlying assumption is that packet arrivals are governed by a Poisson process and packet sizes follow an exponential distribution. The flow of traffic through multiple nodes is modelled as a sequence of concatenated M/M/1 queueing systems. Jackson showed in [58] that in such a network, the queues behave as independent M/M/1 queues. The steady-state distribution of such a network can be described and exactly solved as a product of the steady-state distribution of each queue. A queueing network with this property is called a *product form queueing network*.

Modelling a packet network as a network of M/M/1 queues, however, requires independence assumptions on arrivals, service times and routing. In particular, the assumption of independence of service times means that, in the model, the size of a packet changes as the packet traverses multiple nodes. While this assumption does not hold in practice, the simplicity of the product form made queueing networks a popular tool for the quantitative analysis of packet networks. Work by Baskett et. al. [8] and Kelly [63] relaxed the assumptions on the service time distributions and routing, but maintained the assumption that external arrivals are Poisson and that service times distributions are independent.

The emergence of high-speed data networks in the 1980s has permitted the development of bandwidth demanding network applications such as voice or video. A particular characteristic of voice and video applications is that the transmitted traffic exhibits burstiness. Since Poisson models cannot capture burstiness, more complex traffic models have been proposed to analyze voice and video applications.

Markov-Modulated Fluid (MMF) models have been used by Anick, Mitra and Sondhi in [2] to derive exact solutions for the buffer overflow probability at a node fed by statistically independent flows. The analysis of statistical multiplexing for voice sources (see Daigle and Langford [42]), and video sources (see Maglaris *et. al.* [85]), is also based on MMF models. A fluid traffic model dispenses with the notion of packets (see Jagerman *et. al.* [59]), and is justified in scenarios where the number of packets is large relative to a chosen time scale. For example, a flow is described in [2] as a sequence of exponentially distributed ‘On’ and ‘Off’ periods; while in the ‘On’ state, the flow transmits at a constant rate, and is idle in the ‘Off’ state. Markov-Modulated Poisson Processes (MMPP) (see Heffes and Lucantoni [55]) is another model for the analysis of bursty traffic, such as voice and video [101]. This model is characterized by a support Markov chain; while in a state of the chain, a traffic source transmits as a Poisson process with a certain rate. Burstiness is captured with MMF and MMPP models by using different transmission rates for different states of the underlying Markov chains.

The discovery in the early 1990s that Internet traffic exhibits self-similarity and long-range dependence (LRD) (see Leland *et. al.* [75]) has led to an abandoning of Poisson traffic models (see Paxson and Floyd [93]). Other relevant studies that confirmed the existence of self-similarity and LRD include [93] for wide area networks traffic, Crovella and Bestavros [35] for world-wide-web traffic, and Beran *et. al.* [9] for variable bit rate video.

Self-similar or LRD traffic is fundamentally different from Markov-modulated traffic, and requires new analytical approaches. Traffic exhibits self-similarity if the corresponding rate process looks similar when plotting at multiple time-scale resolutions, ranging over several orders of magnitude; by contrast, the rate process corresponding to Markov-modulated traffic flattens out as the time-scale resolution is increased. Traffic exhibits LRD if it is characterized by correlations ‘at distance’; by contrast, Markov-modulated processes are characterized by short range dependence, meaning that they have a limited memory (Poisson processes have zero memory). In general, self-similarity is not equivalent to LRD. For example, Brownian

motion is self-similar but does not exhibit LRD [92].

Several traffic models, such as fractional Brownian motion (FBM) (see Norros [89]) or stable Lévy processes (see Mikosch *et. al.* [88]), have been proposed to formally capture self-similarity and LRD. The analysis of these models, using techniques such as large-deviations (see Duffield and O’Connell [43]), or extremal properties of Gaussian processes (see Mas-soulié and Simonian [86]), confirmed analytically that self-similarity is fundamentally different from Markov-modulated models. Indeed, self-similar traffic yields non-exponential queueing behavior [43, 86, 89], while Markov-modulated traffic is characterized by exponential queueing behavior [2, 55].

One of the most influential frameworks in analyzing statistical multiplexing in the 1990s is the *effective bandwidth* (see Hui [56], Gibbens and Hunt [51], or Guérin *et. al.* [54]). The effective bandwidth is associated to a flow, describing its minimum required bandwidth to meet certain service guarantees (e.g. buffer overflow probability); this bandwidth is a scalar between the average and peak rate of the flow. A common assumption in deriving effective bandwidths is the asymptotic representation of the steady-state buffer overflow probability $P(B > \sigma) \approx e^{-\theta\sigma}$, for some $\theta > 0$. This approximation is used for Markov-modulated arrivals, which are characterized by an exponential queueing decay. Then, the effective bandwidths of n flows A_j at a link with capacity C are represented by $\alpha_{A_j}(\theta)$, such that $\sum \alpha_{A_j}(\theta) < C$. One way to explicitly construct effective bandwidths $\alpha_{A_j}(\theta)$ is based on large deviations theory. Concretely, $\alpha_{A_j}(\theta) := \frac{\Lambda_j(\theta)}{\theta}$, where $\Lambda_j(\theta) = \lim_{t \rightarrow \infty} \frac{1}{t} \log E [e^{\theta A_j(t)}]$ is the asymptotic decay rate function of flow j (see Whitt [113]). For example, such a relationship has been established for two-states Markov-modulated processes (see Chang [26]), or more general Markov-modulated processes, including MMPP [113].

An attractive feature of effective bandwidth is that the effective bandwidth of an aggregate of flows can be represented as the sum of the individual flows’ effective bandwidth. A drawback, however, of effective bandwidths formulations based on the approximation with large buffer asymptotics is that the statistical multiplexing may not be accurately captured.

In this sense, Choudhury *et. al.* [32] point out that when multiplexing many flows which are more bursty than Poisson, the violation probabilities (in the asymptotic approximations) may be overestimated by several orders of magnitude. Consequently, the prediction of statistical multiplexing may be pessimistic.

An alternative approach for defining effective bandwidth was proposed by Kelly [64]. For stationary flows, effective bandwidths are defined as $\alpha_{A_j}(\theta, t) := \frac{1}{\theta t} \log E [e^{\theta A_j(t)}]$, and are known for a wide variety of arrivals (e.g. Markov-modulated, FBM) [64]. This definition of effective bandwidth is similar to the previous construction $\alpha_{A_j}(\theta)$. The difference is the time parameter t in the definition of $\alpha_{A_j}(\theta, t)$, which turned out to have a critical role in capturing statistical multiplexing (see Courcoubetis *et. al.* [34]).

In the early 1990s, Cruz proposed an entirely new approach for analyzing backlog and delays in networks [36], that later evolved into the *deterministic network calculus* (see Chang [29], Le Boudec and Thiran [16]). The novelty of network calculus is that arrivals and service are represented with envelope functions [36] and service curves (see Cruz [39]), respectively. Envelope functions set worst case descriptions of arrivals, and service curves set lower bounds on the amount of service received by flows. A consequence of these worst case representations is the worst case representation of performance bounds. Thus, the deterministic network calculus can be used for the analysis of network applications which require strict performance guarantees.

A fundamental feature of the deterministic network calculus is that the derivation of end-to-end performance bounds can be reduced to the single-node case. Indeed, using the $(\min, +)$ algebra formulation of deterministic network calculus (see Agrawal *et. al.* [1], Le Boudec [14], Chang [29]), the service given to a flow along a network path can be expressed using a network service curve, as if the flow traversed a single node only. A drawback, however, of the deterministic network calculus is that it cannot capture statistical multiplexing. The reason is that worst case descriptions of arrivals add, meaning that the envelope representation of an aggregate of a large number of flows may be too conservative, further

reflecting into overly pessimistic performance bounds.

The stochastic network calculus is an extension of the deterministic network calculus, motivated by the need to capture statistical multiplexing. The main idea is to extend the concepts of the deterministic network calculus into probabilistic frameworks. For example, a wide variety of studies concern with statistical representations of envelopes [5, 12, 26, 40, 73, 105, 116, 118], or service curves [22, 40, 96]. Statistical envelopes can be constructed from effective bandwidth representations (see Li *et. al.* [76]), which are already known for many types of arrivals. Also, statistical service curves can be constructed for several scheduling algorithms [76].

Using statistical characterizations for arrivals or service, the stochastic network calculus yields probabilistic performance bounds, that carry over relatively easy from the deterministic network calculus (see Burchard *et. al.* [22], Yin *et. al.* [118], Li *et. al.* [76]). By allowing for small violation probabilities in the derivation of performance measures, statistical multiplexing can be captured with the stochastic network calculus by using results from probability theory (e.g. Central Limit Theorem in Knightly [70] and Boorstyn *et. al.* [12], or large deviations tools [12, 26, 76, 110, 111]).

The single-node analysis with the stochastic network calculus has provided satisfactory results and interesting insights in network analysis. For example, it was shown that given some probabilistic delay constraints on flows belonging to several classes of arrivals, the number of admissible flows saturates the available capacity at high data rates [76, 77]. Moreover, the backlog and delay analysis with the network calculus yields probabilistic bounds which hold for all values (e.g. the backlog size), and not only in a log-asymptotic sense as predicted with other modern theories for networks analysis (e.g. effective bandwidth). An important insight provided with the calculus is that at high data rates, statistical multiplexing may dominate the effects of link scheduling; this means that simple scheduling algorithm may suffice in Internet routers (see Liebeherr [77]).

A significant challenge in the stochastic network calculus consists in formulating statis-

tical network service curves, that can carry the properties of deterministic network service curves in probabilistic settings (e.g. the derivation of end-to-end performance bounds). Statistical end-to-end performance bounds can also be derived by adding single-node bounds (see Yaron and Sidi [116]); however, the bounds obtained in this way tend to degrade rapidly in the number of nodes. The technical difficulties associated to the formulation of statistical network service curves led to the introduction of additional assumptions, such as the statistical independence of service (see Chang [29], Fidler [48]), additional requirements on the service curves (see Burchard *et. al.* [22]), or dropping policies at the nodes (see Li *et. al.* [76], Ayyorgun and Cruz [4]).

Therefore, unlike the single-node analysis which is quite well understood, the stochastic network calculus literature concerning the multi-node analysis left open questions. One is related to the construction of a statistical network service curve without relying on the additional assumptions mentioned above. Others concern the accuracy of end-to-end backlog and delay bounds obtained with the stochastic network calculus, or the impact on the performance bounds by assuming statistical independence. These fundamental questions have motivated this thesis.

Before presenting the main contributions of this thesis, let us mention that the mathematical framework of the network calculus is expressed in terms of linear algebra, elementary calculus, and basic probability theory. We thus believe that the calculus is suitable to be employed by network engineers for performance evaluation purposes.

1.3 Thesis Statement and Contributions

As significant progress has been recently made in the area of the stochastic network calculus [5, 13, 22, 29, 33, 48, 76, 82, 105, 118], we share the vision of Liebeherr *et. al.* [78] who assert that: “stochastic network calculus can potentially lead to the development of *simple models* and fast computational methods for communication networks that are very different from the networks and protocols used today”. This thesis attempts to advance the stochastic

network calculus and demonstrate its applicability to analyze packet networks, as expressed in the following statement.

Thesis Statement: *The stochastic network calculus provides new analytical insights into the scaling behavior of network delays.*

The thesis makes contributions in the stochastic network calculus in three directions: theory, applications, and relationships with other theories.

Theory: We propose two formulations of a stochastic network calculus. The first, developed in conjunction with Burchard and Liebeherr (see [33]), is suitable for analyzing network scenarios where arrivals at each node are generally statistically independent, but arrivals and service across the nodes may be statistically correlated. In other words, statistical multiplexing gain is achieved at a single node, whereas arrivals and service across the network may conspire in creating adversarial events.

The literature contains other formulations of a stochastic network calculus for statistically correlated arrivals and service at the nodes (e.g. Yaron and Sidi [116], Cruz [40], or Li *et al.* [76]). Compared to these, the novelty of our formulation is the construction of a *statistical network service curve* that lends itself to explicit numerical evaluations for a wide class of commonly used traffic models. Moreover, our construction gives new insight into the scaling behavior of probabilistic end-to-end performance bounds in networks.

The second network calculus formulation combines the first calculus formulation mentioned above, and a calculus formulation due to Chang [29] and Fidler [48] that is suitable for independent arrivals and service. In this way the statistical independence of arrivals *or* service can be exploited, where available. A scenario where the resulting network calculus formulation is useful is a network with independent arrivals at the nodes, but correlated service times (e.g. a network with identical service times of packets at the nodes).

We apply the second calculus formulation to the class of arrivals processes having stationary and independent increments. To do so, we integrate in network calculus a technique used by Kingman to derive backlog bounds in GI/GI/1 queues (see [66]). The technique is

based on applying a maximal inequality to suitable constructed supermartingales (see [52], page 496). We show that with the maximal inequality, single-node performance bounds can be improved in the stochastic network calculus.

Applications: We consider the class of Exponentially-Bounded-Burstiness (EBB) arrivals (see Yaron and Sidi [116]) which includes many Markov-modulated processes and regulated arrivals. Also, we consider the class of EBB service curves which set lower bounds on the service whose violation probabilities are expressed with exponential functions. The service model is either fluid, i.e., a fraction of a packet becomes available for service as soon as processed upstream, or packetized, i.e., each packet becomes available for service as soon as fully processed upstream.

We apply our first calculus formulation to an abstract network scenario in which a flow with EBB arrivals traverses H nodes in series, each providing EBB service. For this scenario we demonstrate that end-to-end backlog and delay bounds of the flow grow as $\mathcal{O}(H \log H)$. We include examples of networks where end-to-end bounds scale in this fashion by considering both fluid and packetized service models; explicit end-to-end delay bounds are provided in each case.

The derivation of $\mathcal{O}(H \log H)$ end-to-end bounds is a consequence of our construction of statistical network service curves. To further reflect the importance of the network service curve, we show that the derivation of end-to-end bounds by using the alternative method of adding per-node bounds, as suggested for instance by Yaron and Sidi [116], yields results that grow as $\mathcal{O}(H^3)$. The difference between the two scaling behaviors of end-to-end bounds, established in a joint work with Burchard and Liebeherr (see [33]), provides strong evidence on the benefits of using a statistical network service curves in the stochastic network calculus. Similar benefits are known in the deterministic network calculus (see Le Boudec and Thiran [16]), or a stochastic network calculus with statistically independent arrivals and service (see Fidler [48]).

The significance of the $\mathcal{O}(H \log H)$ scaling behavior of end-to-end bounds is further

supported by a corresponding $\Omega(H \log H)$ scaling behavior, established in a joint work with Burchard and Liebeherr (see [21]). We derive this lower bound result for a specific network scenario with EBB arrivals and service. This is done in a tandem network with H nodes, Poisson arrivals and exponentially distributed packet sizes that are maintained across the nodes. The emerging $\Theta(H \log H)$ result clearly indicates that performance bounds in networks have a different scaling behavior than is predicted with other analytical tools. For example, queueing networks theory predicts a $\mathcal{O}(H)$ order of growth of end-to-end bounds, by making additional simplifying assumptions on the statistical independence of service times at the nodes.

Relationship to Existing Theories: One of the main concerns in using theories which express the arrivals and service in terms of bounds is whether the obtained backlog and delay bounds are accurate enough to be applied to practical problems. We attempt to provide insight into the accuracy of stochastic network calculus performance bounds by establishing a relationship with queueing networks theory. Concretely, we apply our second stochastic network calculus formulation in network scenarios which are amenable to an exact analysis with queueing networks theory, and compare the exact results with the bounds obtained with the network calculus approach.

In the single-node case we construct network calculus models for M/M/1, M/D/1 or M/M/1 queues with priorities. In these scenarios, the network calculus bounds closely match the exact results. For multi-node networks we derive network calculus bounds in M/M/1 queueing networks. When compared to the exact results, we find that the calculus bounds are reasonably accurate in scenarios with small amounts of cross traffic. By increasing the amount of cross traffic, the calculus bounds become more conservative.

The two network calculus formulations in this thesis permit the derivation of bounds in M/M/1 networks where arrivals and service at the nodes may be either statistically independent or correlated. The purpose of analyzing such scenarios is to quantify the impact of assuming statistical independence on end-to-end delays. We consider scenarios where statis-

tical correlations exist either among arrivals, service, or both. We also derive performance bounds by using either a fluid or packetized service model, thus providing with evidence on scenarios where the (approximative) fluid service model is justified.

1.4 Thesis Structure

The remaining part of the thesis is structured as follows.

In Chapter 2 we provide a background on the deterministic network calculus. We start with the description of arrivals and service by deterministic envelopes and service curves, respectively. Then we summarize existing results on single-node performance bounds, and finally discuss results on multi-node performance bounds obtained with network service curves, as opposed to adding per-node bounds.

In Chapter 3 we motivate the extension of the deterministic network calculus to a probabilistic setting and survey the literature on the stochastic network calculus. We review existing models of *statistical* envelopes and service curves, and then discuss the problem of constructing statistical network service curves.

In Chapter 4 we formulate a stochastic network calculus that is suitable to analyze network scenarios where arrivals and service at the nodes may be statistically correlated. Here we present the main result of the thesis, i.e., the construction of a statistical network service curve.

In Chapter 5 we use the statistical network service curve constructed in Chapter 4 to analyze how performance bounds scale with the number of nodes in the network. Specifically, in the case of networks with EBB arrivals and service, we compute explicit end-to-end delay bounds and demonstrate that they grow as $\mathcal{O}(H \log H)$. We also prove a corresponding $\Omega(H \log H)$ lower bound on end-to-end delays for a particular queueing model. We provide numerical examples to illustrate the difference between the end-to-end bounds obtained with the network service curves and by adding per-node bounds, and also the difference between the upper and lower bounds.

In Chapter 6 we formulate a stochastic network calculus that can exploit the statistical independence of arrivals or service, where available. We also consider the special case of arrival processes with stationary and independent increments, for which tighter performance bounds can be obtained.

In Chapter 7 we investigate the accuracy of performance bounds derived with the stochastic network calculus formulation from Chapter 6. For the single-node case we consider three common queueing models ($M/M/1$, $M/M/1$ with priorities and $M/D/1$), and for the multi-node case we consider $M/M/1$ networks. Using numerical examples, the bounds obtained using the calculus approach are compared with exact results. We also investigate the role of statistical independence assumptions of arrivals and service in network calculus. Finally, we discuss whether a fluid service model is justified to approximate the more realistic packetized service model.

In Chapter 8 we present conclusions and future work.

Chapter 2

Background on the Deterministic Network Calculus

In this chapter we first introduce notation and describe the network model considered in this thesis. Then we give background on the deterministic network calculus as it applies to topics in this thesis. The background includes discussions on envelope and service curve functions to characterize the arrivals and service, respectively. Then we discuss how to compute single-node and multi-node performance bounds on the backlog and delay of a flow.

2.1 Network Model

In this thesis we consider the network model depicted in Figure 2.1. An aggregate of *through* arrival flows (through traffic) traverses H nodes arranged in series, and each node is also transited by an aggregate of *cross* arrival flows (cross traffic). This network is referred to as a *network with cross traffic*. Each node has a fixed capacity C and the network is stable, i.e., the capacity C is greater than the average arrivals rate at each node. The performance measures of interest are bounds on the end-to-end backlog and delay processes corresponding to the through traffic.

The simplified network considered in Figure 2.1 corresponds to the *view* of a flow traversing a possibly larger network. Although we do not restrict the topology of this larger network




Figure 2.1: A network with cross traffic

(e.g. cycles are permitted), we make two critical assumptions regarding the flow's path. First, routing is always *fixed*, i.e., the flow's data follows the same path for the entire duration of its transmission. Second, and most importantly, we require descriptions of the cross traffic at each node on the path of the through flow. With the network calculus, such descriptions can be constructed for instance in networks with acyclic topologies.

We assume a continuous time model starting at time zero. We represent an incoming data flow at a node by an *arrival process* $A(t)$, and the corresponding outgoing data flow by a *departure process* $D(t)$, where t represents time. $A(t)$ and $D(t)$ represent the cumulative arrivals and departures, respectively, in the interval $[0, t]$. Unless otherwise specified, the data unit is taken to be as one bit. The processes $A(t)$ and $D(t)$ are left-continuous, nondecreasing and satisfy the causality condition $D(t) \leq A(t)$. Also, the arrival process satisfies the initial condition $A(0) = 0$. Occasionally, it is convenient to use the doubly-indexed arrival process $A(s, t)$ defined for all $0 \leq s \leq t$ as

$$A(s, t) \triangleq A(t) - A(s) .$$

Each node has a buffer to store excess data. A *backlog process* $B(t)$ models the amount of data in the buffer at any time $t \geq 0$. If $A(t)$ and $D(t)$ denote the arrivals and departures, respectively, at the node, then $B(t)$ is defined as

$$B(t) \triangleq A(t) - D(t) .$$

The buffer is infinitely sized such that the equation is well-defined.

Besides the backlog process, another measure of interest is the *delay* experienced by the data units of a flow at the node. The delay is defined by the process

$$W(t) \triangleq \inf \left\{ d : A(t-d) \leq D(t) \right\}, \quad (2.1)$$

for some arrivals $A(t)$ and departures $D(t)$. We note that $W(t)$ expresses the virtual delay experienced by a data unit departing at time t . If $A(t)$ is the only flow at the node, then $W(t)$ depends of the backlog process $B(t)$ and the node's rate. If there are additional flows at the node, then $W(t)$ also depends on the scheduling mechanism.

2.2 Deterministic Envelope

One of the original ideas pioneered in network calculus is that traffic is presumably *unknown*, but subject to some regularity constraints (see Cruz [37]). Concretely, the arrivals of a flow are *bounded* by a deterministic *envelope* function for all intervals of time.

In the case when the arrivals of a flow may violate the bounds set by the envelope, then mechanisms are needed to *shape* the surplus arrivals. For example, a *traffic regulator* shapes the arrivals by delaying those arrivals violating the envelope bounds. In contrast, a *traffic policer* simply drops the surplus arrivals. The output traffic resulted by either regulating or policing must satisfy with the envelope, and it should also be the maximum possible with this property. In other words, a shaper should output as much as possible, but within the bounds set by the envelope.

The definition of a deterministic envelope function is given by Cruz [37].

Definition 2.1 (DETERMINISTIC ENVELOPE) *A nonnegative and nondecreasing function $\mathcal{G}(t)$ is a deterministic envelope for an arrival process $A(t)$ if for all $0 \leq s \leq t$*

$$A(t) - A(s) \leq \mathcal{G}(t-s). \quad (2.2)$$

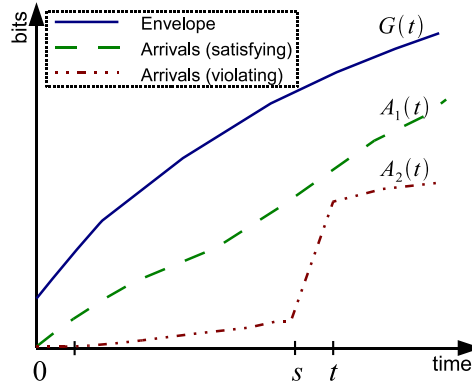


Figure 2.2: An example of an envelope function $\mathcal{G}(t)$ and two sample arrivals: $A_1(t)$ satisfying the envelope, and $A_2(t)$ violating the envelope in the interval $[s, t]$.

In other words, a traffic envelope sets an upper bound on the arrivals in any interval of time. This bound is invariant under time shift, i.e., $A(s, t)$ and $A(s + u, t + u)$ have the same bound for all $u \geq 0$.

Given an envelope $\mathcal{G}(t)$ there is an infinite number of sample arrivals $A(t)$ satisfying the constraint from Eq. (2.2). One example is the arrival process $A(t) = 0$ for all $t \geq 0$. If $\mathcal{G}(t)$ is a sub-additive function¹, then the arrival process defined by $A(0) = 0$ and $A(t) = \mathcal{G}(t)$ for all $t > 0$ also satisfies Eq. (2.2). The reason is that if $\mathcal{G}(t)$ is an envelope function for $A(t)$ then the sub-additive closure of $\mathcal{G}(t)$, i.e., the biggest sub-additive function smaller than $\mathcal{G}(t)$, is also an envelope function for $A(t)$ (see Chang [29], page 38). Figure 2.2 illustrates an envelope function $\mathcal{G}(t)$ together with two sample arrivals. The arrivals $A_1(t)$ satisfy the constraint of the envelope at all times. The arrivals $A_2(t)$ satisfy the envelope in the time interval $[0, s]$, but violate the envelope in the interval $[s, t]$.

Next we review three examples of traffic envelopes. The so-called *leaky-bucket* envelope is described by the function

$$\mathcal{G}(t) = rt + \sigma .$$

If $\mathcal{G}(t)$ is an envelope for the arrival process $A(t)$, then r is an upper bound on the long-term

¹A function $f(t)$ is sub-additive if $f(s + t) \leq f(s) + f(t)$ for all s, t .

average rate of the arrivals $A(t)$, i.e.,

$$r \geq \limsup_{t \rightarrow \infty} \frac{A(t)}{t} .$$

The parameter σ is an upper bound on the *instantaneous burst* of the arrivals, namely the amount of arrivals in a very short interval of time $\Delta t \rightarrow 0$.

The second example of an envelope is used in the specification of the IntServ architecture [18] of the Internet. Assuming that the data unit of arrivals is one packet, the envelope function is given by

$$\mathcal{G}(t) = \min \{rt + \sigma, Pt + L\} ,$$

where r and σ are defined as before. The parameter P sets an upper bound on the *peak rate*, i.e., the maximum arrival rate over any interval of time, and L sets an upper bound on packet sizes.

The third example of an envelope is the *multiple* leaky-bucket envelope proposed by Cruz [38]. For n leaky-buckets, the corresponding envelope takes the form

$$\mathcal{G}(t) = \min_{i=1, \dots, n} \{r_i t + \sigma_i\} .$$

where r_i are rates, and σ_i are bursts. Unlike the leaky-bucket envelope, the multiple leaky-bucket envelope captures the property that the rate of arrivals decreases over sufficiently large intervals. A generalization of the multiple-leaky bucket envelope model is the Deterministically Bounding INterval-length Dependent (D-BIND) model which allows for non-necessarily concave envelope functions (see Knightly and Zhang [72]).

2.3 Deterministic Service Curve

Network calculus represents service either explicitly with scheduling algorithms [37, 81, 90], or with *service curves* [1, 14, 28, 41] that offer an *unknown* representation of service, but sub-

ject to some regularity constraints. The advantage of using service curves is that scheduling can be *separated* from performance analysis, i.e., performance metrics are derived in the same fashion for many scheduling algorithms by first representing the properties of scheduling with service curves. In this thesis we consider the representation of service with service curves in a $(\min, +)$ algebra setting [1, 14, 28].

One way to understand the service curve concept is by analogy with linear-systems theory [74], as illustrated by Cruz and Okino [41], Le Boudec and Thiran [16], pp. *xiv*, and Liebeherr *et. al.* [79]. Let us first introduce the *convolution* operator ‘*’ of two functions $f(t)$ and $g(t)$ in the linear-systems theory

$$f * g(t) \triangleq \int_{-\infty}^{\infty} f(s)g(t-s)ds. \quad (2.3)$$

If $f(t) = 0$ and $g(t) = 0$ for all $t < 0$ then the integration in Eq. (2.3) is taken over the interval $[0, t]$.

For a linear and time-invariant (LTI) system, let us consider the *impulse-response* $h(t)$ of the system, i.e., the output signal produced by the system for the input signal $\delta(t)$ defined as

$$\delta(t) \triangleq \begin{cases} 0, & t \neq 0 \\ \text{undefined}, & t = 0 \end{cases}, \quad (2.4)$$

such that $\int_{-\infty}^{\infty} \delta(s)ds = 1$.

Then, for any input signal $u(t)$, the corresponding output signal $y(t)$ satisfies for all t

$$y(t) = u * h(t). \quad (2.5)$$

The equation is a consequence of the linearity and time-invariance properties of LTI systems. The relationship between the input and output signals is illustrated in Figure 2.3.

By analogy, in network calculus, the departures at a node are related to the corresponding




Figure 2.3: An LTI system: the output signal is represented as the convolution between the input signal and the impulse-response: $y(t) = u * h(t)$.

arrivals by a convolution operator. However, the convolution in network calculus is defined in a modified algebra, called the $(\min, +)$ algebra (see Baccelli *et. al.* [7]). In this algebra, the usual operations of addition and multiplication are replaced by the operations of infimum (minimum) and addition, respectively, as illustrated in Table 2.1

Operations in the usual algebra	Operations in the $(\min, +)$ algebra
$a + b$	$\min \{a, b\}$
$a \cdot b$	$a + b$

Table 2.1: Arithmetic operations in the usual and the $(\min, +)$ algebra.

The convolution operator ‘*’ in the $(\min, +)$ algebra is defined for all $t \geq 0$ as

$$f * g(t) \triangleq \inf_{0 \leq s \leq t} \{f(s) + g(t - s)\} . \quad (2.6)$$

We note that in both algebras we used the same symbol ‘*’ for convolution. For the rest of this thesis ‘*’ will stand for the $(\min, +)$ convolution operator.

The definition of a deterministic service curve is given by Cruz and Okino [41].

Definition 2.2 (DETERMINISTIC SERVICE CURVE) *A nonnegative, nondecreasing function $\mathcal{S}(t)$ is a deterministic service curve for an arrival process $A(t)$ if the corresponding departure process $D(t)$ satisfies for all $t \geq 0$*

$$D(t) \geq A * \mathcal{S}(t) \quad (2.7)$$




Figure 2.4: A service curve in network calculus. The departures are lower-bounded by the convolution between the arrivals and the service curve: $D(t) \geq A * S(t)$.

In other words, a (deterministic) service curve $S(t)$ sets a lower bound on the amount of service received by the arrivals $A(t)$ at the node. This relationship is illustrated in Figure 2.4.

Unlike Eq. (2.5) which holds with equality, Eq. (2.7) holds as an inequality. For this reason the system from Figure 2.4 is not a $(min, +)$ linear system. It becomes linear, however, when there exists a function $S(t)$ such that the relationship in Eq. (2.7) holds with equality for all pairs $(A(t), D(t))$ of arrivals and departures [79]. An example of such a function is the system's output when the input is the burst function

$$\delta(t) \triangleq \begin{cases} 0, & t = 0 \\ \infty, & t > 0 \end{cases},$$

that is the corresponding impulse function in the $(min, +)$ algebra of the input signal from Eq. (2.4).

One typical example of a service curve is the constant-rate (see Le Boudec and Thiran [16], pp. 18). It is represented by the function

$$\mathcal{R}(t) = rt,$$

and expresses the behavior of a node with constant rate. For example, if a node with constant rate r serves an arrival flow $A(t)$, then the corresponding departure process $D(t)$ satisfies for all $t \geq 0$

$$D(t) = A * \mathcal{R}(t).$$

Another example of a service curve is the latency-rate (see Stiliadis and Varma [106]).

To define it, let us now introduce the notation

$$[x]_+ \triangleq \max \{x, 0\}$$

for the positive part of a number x . A latency-rate service curve is a shifted version of the constant-rate service curve, i.e., is represented by a function

$$\mathcal{S}(t) = r [t - d]_+ ,$$

for some latency (delay) $d \geq 0$. This type of a service curve guarantees a maximum delay d for the first data unit seen in a *busy period of the flow* (a maximum period of time during which the average arrivals' rate of the flow is above r).

A special type of a service curve is the *strict service curve* (see Cruz and Okino [41]), that is a function $\mathcal{S}(t)$ setting a lower bound on the amount of departures in any (system) *busy period* of length t (a system busy period is an interval of time where the backlog process $B(t)$ is always positive). Formally, $\mathcal{S}(t)$ is a strict service curve if for any time interval $[s, t]$ during which the backlog is positive the following holds

$$D(t) - D(s) \geq \mathcal{S}(t - s) . \tag{2.8}$$

A strict service curve is also a service curve but the converse is not necessarily true. For this reason strict service curves provide more accurate characterizations of service at a node, and lead to improved bounds for the performance measures (see Le Boudec and Thiran [16], pp. 29). A drawback of strict service curves is that they are not closed under convolution, i.e., the convolution of strict service curves does not necessarily result in a strict service curve.

The deterministic network calculus provides constructions of service curves for several scheduling algorithms. Here we consider the case of a static priority (SP) scheduling algorithm which assigns priorities to the flows and selects for transmission a flow with a positive

backlog and the highest priority. The scheduler is workconserving, i.e., always active when the backlog is positive. When dealing with delay processes we further assume that the arrivals within a single flow are scheduled in a FIFO (First-In-First-Out) order; this is usually referred to as locally-FIFO.

Consider the construction of service curves for a flow, or an aggregate of flows, which receives the lowest priority at an SP scheduler. These service curves are suggestively referred to as *leftover service curves*, since they express the capacity left unused by the higher priority flows. Leftover service curves provide thus a worst-case description of service, and have the property that they are guaranteed by *any* workconserving scheduling mechanism. The next theorem (see Chang [29], pp. 60) provides such a construction.

Theorem 2.3 (LEFTOVER SERVICE CURVE) *Consider a workconserving node with fixed capacity C serving a tagged flow $A(t)$, and another flow $A_c(t)$ with an envelope function $\mathcal{G}_c(t)$. Then, a service curve given by the node to the tagged flow $A(t)$ is given by the function*

$$\mathcal{S}(t) = [Ct - \mathcal{G}_c(t)]_+ . \quad (2.9)$$

Let us next briefly sketch a proof of the theorem. Arguments from the proof will be used in later chapters.

Denote by $D_c(t)$ the departure process corresponding to $A_c(t)$. Because $A_c(t)$ has higher priority, the function $\mathcal{R}(t) = Ct$ is a constant-rate service curve for the cross flow satisfying Eq. (2.7) with equality, i.e.,

$$D_c(t) = A_c * \mathcal{R}(t) .$$

The function $\mathcal{R}(t)$ is also a service curve for the aggregate process $A(t) + A_c(t)$, such

that the output process $D(t)$ of the tagged flow can be written as follows

$$\begin{aligned}
D(t) &= (D(t) + D_c(t)) - D_c(t) \\
&= (A + A_c) * \mathcal{R}(t) - A_c * \mathcal{R}(t) \\
&\geq \inf_{0 \leq s \leq t} \{A(s) + A_c(s) + \mathcal{R}(t-s) - \min\{A_c(t), A_c(s) + \mathcal{R}(t-s)\}\} \\
&\geq \inf_{0 \leq s \leq t} \{A(s) + \max\{\mathcal{R}(t-s) - (A_c(t) - A_c(s)), 0\}\} \\
&\geq \inf_{0 \leq s \leq t} \{A(s) + \max\{C(t-s) - \mathcal{G}_c(t-s), 0\}\} \\
&= A * \mathcal{S}(t),
\end{aligned}$$

showing that the function $\mathcal{S}(t)$ from Eq. (2.9) is a leftover service curve for the tagged flow. The third line follows from $A_c * \mathcal{R}(t) \leq \min\{A_c(t), A_c(s) + \mathcal{R}(t-s)\}$. In the fourth line we reordered terms, and the fifth line follows from the definition of the envelope $\mathcal{G}_c(t)$ for the cross flow.

2.4 Single-Node Performance Bounds

The deterministic network calculus provides bounds on the performance measures of interest at the node, given an envelope description of the arrivals at a node, and a service curve relating the arrivals with the corresponding departures. Some of the bounds are concisely expressed with the *deconvolution* operator ‘ \oslash ’ defined for two functions $f(t)$ and $g(t)$ as

$$f \oslash g(t) = \sup_{s \geq 0} \{f(t+s) - g(s)\} .$$

The next theorem (from Le Boudec and Thiran [16], pp. 22-23) gives bounds on the backlog and delay processes of a flow at a node, and also constructs an output envelope for the flow.

Theorem 2.4 (PERFORMANCE BOUNDS) *Consider a flow at a node with arrivals and departures given by the processes $A(t)$ and $D(t)$, respectively. Assume that the arrivals are*

bounded by an envelope $\mathcal{G}(t)$, and that the node provides a service curve $\mathcal{S}(t)$ to the flow. Then following deterministic bounds hold.

1. **OUTPUT ENVELOPE:** The function $\mathcal{G} \circledast \mathcal{S}$ is a deterministic envelope for the departures $D(t)$, i.e., for all $0 \leq s \leq t$

$$D(t) - D(s) \leq \mathcal{G} \circledast \mathcal{S}(t - s) .$$

2. **BACKLOG BOUND:** A bound on the backlog process $B(t)$ is given for all $t \geq 0$ by

$$B(t) \leq \mathcal{G} \circledast \mathcal{S}(0) .$$

3. **DELAY BOUND:** A bound on the delay process $W(t)$ is given for all $t \geq 0$ by

$$W(t) \leq \inf \{d : \mathcal{G}(s) \leq \mathcal{S}(s + d) \text{ for all } s \geq 0\} .$$

If $\mathcal{G}(t)$ is sub-additive, $\mathcal{G}(0) = 0$ and $\mathcal{S}(0) = 0$, then the backlog and delay bounds are *tight*, in the sense that there exist arrivals that actually meet the bounds given in the theorem (see Le Boudec and Thiran [16], pp. 27).

To illustrate the simplicity of performance analysis with network calculus, let us next briefly sketch the proof for the backlog bound. Using the definition of the backlog process we can write

$$\begin{aligned} B(t) = A(t) - D(t) &\leq A(t) - A * \mathcal{S}(t) \\ &\leq \sup_{0 \leq s \leq t} \{A(t) - A(s) - \mathcal{S}(t - s)\} \\ &\leq \sup_{0 \leq s \leq t} \{\mathcal{G}(t - s) - \mathcal{S}(t - s)\} \\ &\leq \mathcal{G} \circledast \mathcal{S}(0) . \end{aligned} \tag{2.10}$$

The first line follows from the definition of the service curve. The convolution operator is

then expanded in the second line. The third line follows from the definition of an envelope, and is finally rewritten in terms of the deconvolution operator.

2.5 Deterministic Network Service Curve

A fundamental property of the network calculus is that service curves can be *concatenated*. If a flow is described with service curves at each node along a network path, then the flow can be described with a single service curve, called a *network service curve*, along the entire network path. The flow can thus be regarded as traversing a single node only, such that end-to-end performance bounds can be obtained by applying single-node results.




Figure 2.5: A flow with service curves at multiple nodes.

The next theorem (from Le Boudec and Thiran [16], pp. 28) formalizes the concatenation property of service curves.

Theorem 2.5 (DETERMINISTIC NETWORK SERVICE CURVE) *Consider a flow traversing H nodes in series, as in Figure 2.5. Assume that each node h provides a service curve $\mathcal{S}^h(t)$ to the flow. Then, the service given to the flow by the network as a whole can be expressed with the network service curve*

$$\mathcal{S}^{net}(t) = \mathcal{S}^1 * \mathcal{S}^2 * \dots * \mathcal{S}^H(t), \quad (2.11)$$

in the sense that for all $t \geq 0$

$$D^{net}(t) \geq A^{net} * \mathcal{S}^{net}(t).$$

The critical information used in the proof is that the departures at a node coincide with the arrivals at the next immediate node (i.e. $D^h = A^{h+1}$ for all $h = 1, \dots, H - 1$). Using this information, the proof is straightforward by applying the associativity property of the convolution operator ‘*’ in the $(\min, +)$ algebra. Indeed, one can write for all $t \geq 0$

$$\begin{aligned}
D^H(t) &\geq A^H * \mathcal{S}^H(t) \\
&= D^{H-1} * \mathcal{S}^H(t) \\
&\geq (A^{H-1} * \mathcal{S}^{H-1}) * \mathcal{S}^H(t) \\
&= D^{H-2} * (\mathcal{S}^{H-1} * \mathcal{S}^H)(t) \\
&\geq \dots \\
&\geq A^1 * (\mathcal{S}^1 * \mathcal{S}^2 * \dots * \mathcal{S}^H)(t) .
\end{aligned}$$

In general, the numerical complexity of the convolution operation from Eq. (2.11) is small because the functions \mathcal{S}^h are deterministic. Consider for example that all functions $\mathcal{S}^h(t)$ are constant-rate service curves with some rates r_h for $h = 1, \dots, H$. Then the corresponding network service curve is simply given for all $t \geq 0$ by

$$\mathcal{S}^{net}(t) = \inf \{r_1, \dots, r_H\} t ,$$

that is $\mathcal{S}^{net}(t)$ is a constant-rate service curve as well, and its rate is the minimum rate of the service curves in the convolution.

To illustrate the benefits of network service curves we next apply Theorem 2.5 to compute explicit end-to-end backlog and delay bounds in a particular network scenario, followed by the analysis of their scaling properties. We will then compare the obtained bounds with corresponding end-to-end bounds obtained using an alternative method of adding per-node bounds, i.e., without using a network service curve.

Let us first review the Landau notation for the asymptotic behavior of functions.

Definition 2.6 (LANDAU NOTATION) For two positive functions $f(t)$ and $g(t)$ we denote $f(t) = \mathcal{O}(g(t))$ (asymptotic upper bound) and $f(t) = \Omega(g(t))$ (asymptotic lower bound) if the fractions $f(t)/g(t)$ and $g(t)/f(t)$, respectively, are bounded as $t \rightarrow \infty$. Also, $f(t) = \Theta(g(t))$ (asymptotic equivalence) if both $f(t) = \mathcal{O}(g(t))$ and $f(t) = \Omega(g(t))$.




Figure 2.6: A network with cross traffic and leaky-bucket envelopes

We use the following notations and settings in the network from Figure 2.6. At each node $h = 1, \dots, H$, we denote the arrivals and departures of the through traffic by $A^h(t)$ and $D^h(t)$, respectively. Also, we denote the arrivals of the cross traffic at node h by $A_h(t)$. For simplicity we assume that the through and cross traffic are constrained by the same leaky-bucket envelope

$$\mathcal{G}(t) = rt + \sigma ,$$

with rate r and burst σ . Finally, we assume a stability condition, i.e., the capacity C at each node is greater than the total arrival rate $2r$ at each node. The performance measures of interest are bounds on the end-to-end backlog $B^{net}(t)$, and the end-to-end delay $W^{net}(t)$ corresponding to the through flow.

To derive end-to-end delay bounds we first invoke Theorem 2.3 yielding a leftover service curve $\mathcal{S}^h(t)$ for the through flow at each node h

$$\mathcal{S}^h(t) = [(C - r)t - \sigma]_+ . \quad (2.12)$$

Given the per-node leftover service curves $\mathcal{S}^h(t)$ from Eq. (2.12), Theorem 2.5 yields the

network service curve

$$\begin{aligned}
\mathcal{S}^{net}(t) &= \mathcal{S}^1 * \mathcal{S}^2 * \dots * \mathcal{S}^H(t) \\
&= \inf_{s_1 + \dots + s_H = t} \{ [(C - r)s_1 - \sigma]_+ + \dots + [(C - r)s_H - \sigma]_+ \} \\
&\geq [(C - r)t - H\sigma]_+ .
\end{aligned}$$

In the last equation we used that $[x]_+ + [y]_+ \geq [x + y]_+$ for any numbers x, y . To simplify notation we let $\mathcal{S}^{net}(t) = [(C - r)t - H\sigma]_+$ to be the network service curve; we can do so because, in general, if a function $\mathcal{S}(t)$ is a service curve for some arrivals $A(t)$ then any function smaller than $\mathcal{S}(t)$ is a service curve as well.

The single-node results from Theorem 2.4 now yields the end-to-end backlog bound

$$\begin{aligned}
B^{net}(t) &\leq \sup_{u \geq 0} \{ ru + \sigma - (C - r)u + H\sigma \} \\
&\leq (H + 1)\sigma .
\end{aligned}$$

The corresponding end-to-end delay bound is

$$W^{net}(t) \leq (H + 1) \frac{\sigma}{C - r} .$$

These end-to-end bounds grow according to

$$B^{net}(t) = \mathcal{O}(H), \quad W^{net}(t) = \mathcal{O}(H) . \quad (2.13)$$

For comparison, we now turn to the derivation of end-to-end bounds using the method of adding per-node bounds. To derive per-node bounds at a node h we first need an envelope description for the intermediary arrival process $A^h(t)$. Applying the output envelope bound from Theorem 2.4 with the service curve derived in Eq. (2.12), and using that $D^h(t) =$

$A^{h+1}(t)$, we get the envelope description

$$\begin{aligned} A^2(t) - A^2(s) &\leq \mathcal{G} \otimes \mathcal{S}^1(t-s) \\ &\leq \sup_{u \geq 0} (r(t-s+u) + \sigma - (C-r)u + \sigma) \\ &\leq r(t-s) + 2\sigma, \end{aligned}$$

at the second node. Then, inductively, the envelopes descriptions for the through flow at each node h are given by

$$A^h(t) - A^h(s) \leq r(t-s) + h\sigma. \quad (2.14)$$

Having an envelope and service curve description for the through flow at each node h (Eqs. (2.14) and (2.12)), Theorem 2.4 yields the per-node backlog bounds

$$\begin{aligned} B^h(t) &\leq \sup_{u \geq 0} (ru + h\sigma - (C-r)u + \sigma) \\ &\leq (h+1)\sigma. \end{aligned}$$

Similarly, the per-node delay bounds are given by

$$W^h(t) \leq (h+1) \frac{\sigma}{C-r}.$$

Finally, a bound on the end-to-end backlog $B^{net}(t)$ is obtained by adding the per-node bounds $B^h(t)$ for $h = 1, \dots, H$, i.e., for all $t \geq 0$

$$B^{net}(t) \leq \frac{H(H+3)}{2} \sigma. \quad (2.15)$$

Similarly, the end-to-end delay bound is given for all $t \geq 0$ by

$$W^{net}(t) \leq \frac{H(H+3)}{2} \frac{\sigma}{C-r}.$$

From the last two equations we conclude that the method of adding per-node bounds yields end-to-end backlog and delay bounds characterized by a quadratic growth in the number of network nodes H , i.e.,

$$B^{net}(t) = \mathcal{O}(H^2), \quad W^{net}(t) = \mathcal{O}(H^2), \quad (2.16)$$

as opposed to the linear growth in H observed in Eq. (2.13). We also remark that the bounds from Eq. (2.13) are always smaller than the bounds from Eq. (2.16), i.e., the improvement of the network service curve method over the method of adding per-node bounds is not only asymptotic, but it holds for all the values of the traffic descriptions.

The property of a network service curve to yield $\mathcal{O}(H)$ end-to-end bounds is related to the so-called ‘*pay-bursts-only-once*’ property (see Le Boudec and Thiran [16], pp. 28, where a comparison between the two methods is carried out for two nodes). This property suggests that the burst of a flow contributes to the end-to-end bound as if the flow traversed a single node only. In our example, the contribution of the through flow to the end-to-end backlog bound from Eq. (2.13) is a *single* burst σ ; the remaining burst $H\sigma$ stems from the individual bursts of each of the cross flows. In contrast, the contribution of the through flow to the end-to-end backlog bound from Eq. (2.15) is in the order of $\mathcal{O}(H^2)$ bursts, obtained by adding the linearly increasing bursts of the through flow at each traversed node.

Another interpretation on the improved performance of the method of using network service curves over the method of adding per-node bounds is provided by Chang (see [29], pp. 87). The former method accounts for the fact that the maximum delay at each node may be not experienced by the *same* data unit. In contrast, the latter method simply adds the worst-case delays at each node.

Chapter 3

State of the Art in the Stochastic Network Calculus

In this chapter we review prior work in the stochastic network calculus. The main concepts and results of the calculus are discussed by closely following the structure used in the previous chapter. Let us first discuss the motivation of extending the deterministic network calculus to a probabilistic setting.

3.1 The Need for a Probabilistic Extension of Network Calculus

The deterministic network calculus is a theory for the worst-case performance analysis of networks. Because the arrivals are represented with *deterministic* envelopes, the calculus may yield overly pessimistic performance bounds as argued in the following.

A drawback of using deterministic envelopes is that they cannot accurately capture the statistical properties of arrivals. Consider for example a Bernoulli traffic source. At discrete instants of time, equally spaced, the source emits P data units with probability p , and is idle with probability $1 - p$. We call P the peak rate, and denote $r = pP$ as the average rate. The

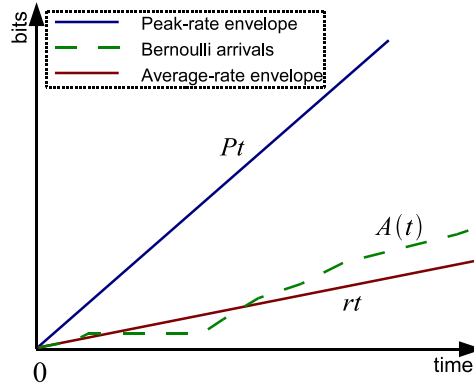


Figure 3.1: Bernoulli traffic source $A(t)$ with peak-rate envelope Pt and average-rate envelope rt .

smallest deterministic envelope describing the Bernoulli source is the function

$$\mathcal{G}(t) = Pt, \quad (3.1)$$

for the source may potentially transmit at each instance of time. However, the probability of such an event is zero. Moreover, the law of large numbers gives that the corresponding cumulative arrival process $A(t)$ behaves asymptotically as

$$A(t) \sim rt,$$

for large values of time t . In other words, the cumulative arrival process behaves asymptotically as if the source transmits at r data units at each instance of time. Therefore, especially for small probabilities p , the envelope $\mathcal{G}(t)$ from Eq. (3.1) is too conservative to model a Bernoulli source over long periods of time; for illustration see Figure 3.1. It then follows that backlog and delay bounds may be too conservative since they increase with the envelope function (see Theorem 2.4).

A closely related limitation of deterministic envelopes is that they cannot account for statistical multiplexing. Consider N arrival flows represented by the processes $A_1(t), A_2(t), \dots, A_N(t)$, and assume that each flow is bounded by the same leaky-bucket envelope with rate r and burst σ . Then, the aggregate arrivals $A(t) = \sum_{i=1}^N A_i(t)$ of all flows

is bounded by the leaky-bucket envelope

$$\mathcal{G}(t) = Nrt + N\sigma .$$

The key problem with the aggregate envelope $\mathcal{G}(t)$ is that the corresponding burst $N\sigma$ may overestimate the actual burst of the aggregate arrivals. Although all flows may simultaneously experience the same burst σ , the probability of such an event is small under some mild statistical independence assumptions on the arrivals, and when N is large enough. In fact, the Central Limit Theorem implies that the aggregate burst of the arrivals is in the order of $\mathcal{O}(\sqrt{N})$, when the number of flows N is large. Consequently, the corresponding performance bounds are likely to grow as $\mathcal{O}(\sqrt{N})$, which is much smaller than the $\mathcal{O}(N)$ order of growth predicted with the deterministic network calculus (see for instance the bounds computed in Section 2.5, properly scaled by N).

It is thus possible to capture statistical multiplexing gain when accounting for the statistical properties of arrivals, e.g., statistical independence. The cost of capturing the statistical multiplexing is that the predicted performance bounds may be violated with some probabilities ε . For instance, a buffer may overflow with probability ε , if its size is set to a value σ satisfying

$$Pr\left(B(t) > \sigma\right) \leq \varepsilon .$$

In practice, the violation probabilities corresponding to performance bounds are chosen to be negligible, i.e., in the order of 10^{-6} to 10^{-9} . Nonetheless, substantial statistical multiplexing gain can be achieved for such extreme choices of the violation probabilities. This observation has partly motivated the extension of the deterministic network calculus in a probabilistic setting.

3.2 Statistical Envelope

Traffic is generally described in the stochastic network calculus by statistical envelopes, which are probabilistic extensions of deterministic envelopes. There are several formulations of statistical envelopes which can be classified with regard to whether the envelope functions are either non-random or random. For each envelope we discuss construction methods and expose the class of arrivals covered by the envelope.

3.2.1 Statistical envelope as non-random function

A statistical envelope sets bounds on the arrivals. The bounds may be violated with some probabilities specified by an *error function*.

The *Exponentially Bounded Burstiness (EBB)* model (Yaron and Sidi [116]) defines statistical envelopes as non-random functions.

Definition 3.1 (EBB ENVELOPE) *An arrival process $A(t)$ is bounded by an EBB envelope with rate r if there exist the constants $M, \theta > 0$ such that for all $0 \leq s \leq t$ and $\sigma \geq 0$*

$$\Pr\left(A(t) - A(s) > r(t - s) + \sigma\right) \leq M e^{-\theta\sigma} . \quad (3.2)$$

The function $\mathcal{G}(t) = rt$ is the *EBB statistical envelope*, and $\varepsilon(\sigma) = M e^{-\theta\sigma}$ is the corresponding *error function*. The constant M is usually referred to as *prefactor*. In general, both the rate r and the prefactor M depend on θ . The constant θ is called the *exponential decay rate* and determines the shape of the error function. In turn, the shape of the error function is closely related the traffic models covered by the envelope.

The EBB envelope model relates to the linear envelope process model proposed by Chang [26]. The linear envelope model can be expressed using the following effective bandwidth characterization for a stationary process $A(t)$, for all $t, \theta > 0$ [64]

$$\alpha_A(\theta, t) = \frac{1}{\theta t} \log E \left[e^{\theta A(t)} \right] . \quad (3.3)$$

A stationary arrival process $A(t)$ is bounded by a linear envelope process with rate r and burst σ , for a choice of $\theta > 0$, if for all $t \geq 0$

$$t\alpha_A(\theta, t) \leq rt + \sigma . \quad (3.4)$$

In general, both the rate r and the burst σ depend on θ .

If a linear envelope model has a rate r and burst σ , then it reduces to the EBB model with the same rate r , prefactor $M = e^{\theta\sigma}$, and decay rate θ ; this follows from the Chernoff bound (see Eq. (3.32)). The converse is also true, i.e., the EBB envelope model reduces to a linear envelope model; this follows from Lemma 1 in Li *et. al.* [76].

Yaron and Sidi propose two methods to construct EBB envelopes for an arrival process [116]. The first method assumes a bound on the moment generating function of the arrivals. This method is equivalent to the above construction of an EBB envelope given a linear envelope model. The second method is based on the relationship between an arrival process and the corresponding backlog process: if the arrival process is EBB then the backlog process has an exponentially decaying rate, and vice-versa. Therefore, any arrival process for which the backlog process has an exponential decaying rate can be expressed with an EBB envelope. Markov-modulated processes are typical examples of processes with an exponential decay rate of the backlog.

To express arrival processes for which the decay rate of the corresponding backlog process is not necessarily exponential, the EBB model was generalized in several ways. One such generalization is the *Stochastically Bounded Burstiness* (SBB) envelope model (Starobinski and Sidi [105]). The SBB model generalizes the EBB model in that the error function $\varepsilon(\sigma)$ is now required to only be n -fold integrable, i.e.,

$$\underbrace{\int \dots \int}_{n \text{ times}} \varepsilon(u) du^n < \infty ,$$

and not necessarily an exponential.

When compared to the EBB model, one advantage of the SBB model is that it covers a broader class of traffic. In particular, SBB envelopes describe FBM arrivals with an error function given by $\varepsilon(\sigma) = e^{-\theta\sigma^\alpha}$, for some decay rate $\theta > 0$ and a parameter $0 < \alpha < 1$ that relates to the Hurst parameter of the FBM. This type of error function has a decay rate that is slower than the exponential decay rate corresponding to the EBB model, yet it satisfies the n -fold integrability condition (see Yin *et. al.* [118]).

Another advantage of SBB envelopes is that they can give a more accurate representations of arrivals represented by Markov-modulated processes than EBB envelopes [105]. By allowing the error function to be a sum of exponentials, i.e., $\varepsilon(\sigma) = \sum_i M_i e^{-\theta_i \sigma}$, the SBB model can capture the property that the backlog process associated to some Markovi-modulated arrivals experience a loss probability with two regions.

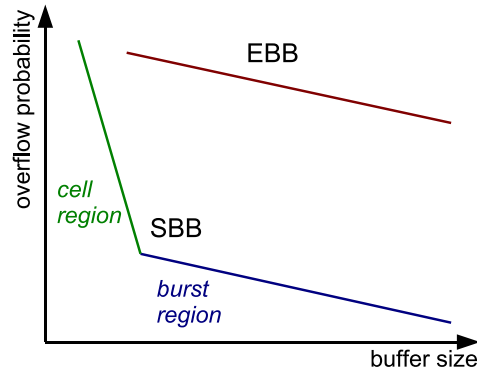


Figure 3.2: The *cell* and *burst* regions corresponding to Markov-modulated traffic. The buffer overflow probability is represented on a logarithmic scale (from [105])

Figure 3.2 illustrates the two regions by showing the buffer overflow probability, on a logarithmic scale, as a function of the buffer size. The following explanation of the two regions is provided by Schwartz [101]. For small buffer sizes, i.e., the *cell region*, the loss probability decays fast. The reason is that the dynamic of the system is mainly driven by the arrival rates in some states that may generate high bursts over short interval of times. For bigger buffer sizes, i.e., the *burst region*, the decay of the loss probability slows down as arrival bursts generated over short periods are absorbed by the buffer, and the probability that

the underlying Markov chain changes states increases.

The (*local*) *effective envelope* model is another statistical envelope model which extends the scope of the EBB model, and was proposed by Boorstyn *et. al.* [13]. An arrival process $A(t)$ is bounded by an effective envelope $\mathcal{G}(t, \varepsilon)$, for some violation probability $\varepsilon > 0$, if for all $0 \leq s \leq t$

$$Pr\left(A(t) - A(s) > \mathcal{G}(t - s, \varepsilon)\right) \leq \varepsilon. \quad (3.5)$$

Next we review two constructions of effective envelopes. The first construction is given for an aggregate of regulated flows. Let us first introduce the concept of *rate variance envelope* (see Knightly [70, 71]) for a stationary arrival process $A(t)$ as

$$RV(t) \triangleq Var\left(\frac{A(t)}{t}\right), \quad (3.6)$$

where $Var(X) \triangleq E[(X - E[X])^2]$ denotes the variance of a random variable X . Given N deterministically-regulated flows, each with the rate variance $RV(t)$, an effective envelope is obtained in [13] using the Central Limit Theorem

$$\mathcal{G}(t, \varepsilon) = Nrt + z\sqrt{Nt}\sqrt{RV(t)}, \quad (3.7)$$

where $z \approx \sqrt{|\log(2\pi\varepsilon)|}$, and $r = \lim_{s \rightarrow \infty} \frac{A(t, t+s)}{s}$ is the upper bound on the long-term arrival rate of $A(t)$. The effective envelope from Eq. (3.7) includes the first and second moments of the arrivals. The envelope captures statistical multiplexing since the underlying burst determined by the variance scales as $\mathcal{O}(\sqrt{N})$ in the number of flows.

A second construction of an effective envelope is proposed by Li *et. al.* [76] using the effective bandwidth of an arrival flow from Eq. (3.3). Suppose that an arrival process $A(t)$ has the effective bandwidth $\alpha_A(\theta, t)$, and let some violation probability $\varepsilon > 0$. Then, the

corresponding statistical envelope $\mathcal{G}(t, \varepsilon)$ is given by [76]

$$\mathcal{G}(t, \varepsilon) = \inf_{\theta > 0} \left\{ t\alpha_A(\theta, t) - \frac{\log \varepsilon}{\theta} \right\} .$$

The importance of this result is that it enables the applicability of network calculus to a wide class of arrivals, for which effective bandwidths are available (e.g. Markov-modulated processes, deterministically regulated or FBM). We point out that the method from [76] can also be used to the construction of SBB envelopes.

The multiplexing gain obtained using effective envelopes constructed as above is numerically evaluated for regulated traffic by Boorstyn *et. al.* [13], MPEG video traces by Liebeherr [77], and Markov-modulated and FBM traffic by Li *et. al.* [76].

The statistical envelope models considered so far (EBB, SBB, effective envelope) have in common the idea of bounding probabilities involving only *two points of the history* of an arrival processes; we denoted these points above by s and t . Next we review statistical envelopes which bound probabilities involving the *entire past history* of an arrival process; such envelopes are usually referred to as *sample-path envelopes* (see Burchard *et. al.* [23]).

The *generalized Stochastically Bounded Burstiness* (gSBB) model (see Yin *et. al.* [118]) is an example of sample-path envelopes. An arrival process $A(t)$ is bounded by a gSBB envelope with upper rate r and error function $\varepsilon(\sigma)$ if for all $t \geq 0$ and $\sigma \geq 0$

$$Pr \left(\sup_{0 \leq s \leq t} (A(t) - A(s) - r(t - s)) > \sigma \right) \leq \varepsilon(\sigma) . \quad (3.8)$$

This envelope model appeared first in Cruz [40], and later in Ayyorgun and Feng [5], under different names.

An advantage of the gSBB model is that it extends the class of traffic models covered by the SBB model. In particular, the gSBB model captures heavy-tailed traffic characterized by a power law decay of the error function, i.e., $\varepsilon(\sigma) = \sigma^{-\alpha}$ with $1 < \alpha \leq 2$.

Although the gSBB model appears more stringent than the SBB model, they are closely

related. Indeed, the gSBB model reduces to the SBB model, subject to the condition that the corresponding error function $\varepsilon(\sigma)$ is n -fold integrable. Conversely, suppose that an arrival process $A(t)$ is bounded by an SBB envelope $\mathcal{G}(t) = rt$ with error function $\varepsilon(\sigma)$. Then, for any choice of $\delta > 0$, the arrival process $A(t)$ is described using the gSBB envelope

$$Pr \left(\sup_{0 \leq s \leq t} (A(t) - A(s) - (r + \delta)(t - s)) > \sigma \right) \leq \frac{1}{\delta} \int_{\sigma}^{\infty} \varepsilon(u) du .$$

The gSBB envelope has a bigger rate $r + \delta$, and the corresponding error function is obtained by integrating the error function corresponding to the SBB envelope [118].

gSBB envelopes can also be constructed following a procedure proposed by Jiang and Emstad [62]. Consider a node with capacity C and some arrival flows with *unknown* statistical envelopes. If a backlog bound with some error function $\varepsilon(\sigma)$ is available, then the arrival process can be described with gSBB envelope $\mathcal{G}(t) = Ct$ and error function $\varepsilon(\sigma)$.

A second example of a sample-path envelope model is the *sample-path effective envelope* model (see Burchard *et. al.* [23]). An arrival process $A(t)$ is bounded by a sample-path effective envelope $\mathcal{G}(t, \varepsilon)$, for some violation probability $\varepsilon > 0$, if for all $t \geq 0$

$$Pr \left(\sup_{0 \leq s \leq t} (A(t) - A(s) - \mathcal{G}(t - s, \varepsilon)) > 0 \right) \leq \varepsilon . \quad (3.9)$$

This model generalizes the effective envelope model in a similar way that the gSBB model generalizes the SBB model. The gSBB and sample-path effective envelopes enable the derivation of single-node performance bounds (e.g. backlog, delay, output envelopes) which resemble to the corresponding bounds in the deterministic network calculus (for more details see Section 3.4.4).

The third example of a sample-path envelope model is the *global effective envelope model* proposed by Boorstyn *et. al.* [13]. An arrival process $A(t)$ is bounded by a global effective envelope sample-path $\mathcal{G}(t, \beta, \varepsilon)$ for an interval $I_{\beta} = [u, u + \beta)$ of length β , and for some

violation probability $\varepsilon > 0$, if

$$\Pr \left(\sup_{u \leq s \leq t < u + \beta} (A(t) - A(s) - \mathcal{G}(t - s, \beta, \varepsilon)) > 0 \right) \leq \varepsilon. \quad (3.10)$$

Unlike the gSBB and sample-path effective envelopes, the global effective envelope sets probabilistic bounds on interarrivals in a *fixed* interval of time I_β . The attribute “global” is justified since the global effective envelope model poses stronger requirements on the interarrivals than the sample-path effective envelope model, by letting two free variables (i.e. s and t in Eq. (3.10)) compared to a single free variable (i.e. u in Eq. (3.9)).

Global effective envelopes $\mathcal{G}(s, \beta, \varepsilon)$ can be constructed from statistical envelopes $\mathcal{G}(s, \varepsilon)$ satisfying Eq. (3.5), for any interval of time I_β (see Boorstyn *et. al.* [13]). The actual construction implies that the violation probability ε in Eq. (3.10) grows with the length β of the interval I_β . Consequently, global effective envelopes are relevant only for finite time intervals that yield violation probabilities less than one in Eq. (3.10). Busy periods are typical examples of such intervals (see [13]). For stationary arrival processes, if Eq. (3.10) holds for a single interval of length I_β , then it holds for all intervals of same length [13].

An advantage of the effective envelope and global effective envelope models is that they allow the derivation of schedulability conditions in the stochastic network calculus for several scheduling algorithms [13], in a manner that resembles the corresponding schedulability conditions in a deterministic context (see Liebeherr *et. al.* [81]). Since the global effective envelope model is stronger than the effective envelope model, the former model leads to more conservative schedulability conditions than the latter.

To further support the claim from the beginning of the section that statistical envelopes lead to statistical multiplexing gain, let us refer next to the works of Knightly [70, 71], Boorstyn *et. al.* [13], and Liebeherr [77]. The authors derive schedulability conditions for admitting a maximum number of flows at a node, under some delay constraints and some violation probabilities. The conditions involve the statistical envelope for an aggregate of flows,

and are extensions of the corresponding schedulability conditions concerning deterministic delay guarantees (see Liebeherr *et. al.* [81]). For example, let the statistical envelope $\mathcal{G}(t, \varepsilon)$, in the sense of Eq. (3.5), for an aggregate of N flows. Then, a FIFO node with capacity C can guarantee delays less than d for the flows with a violation probability less than ε , if the following condition given by Boorstyn *et. al.* [13]

$$\sup_{t \geq 0} \{\mathcal{G}(t, \varepsilon) - Ct\} \leq Cd$$

holds (i.e. the N flows can be admitted under the service requirements). Similar conditions are available for SP or EDF (Earliest Deadline First) scheduling as well. The results obtained using the schedulability conditions from [13, 70, 71, 77] closely match the results obtained from simulations or based on average-rate utilizations, indicating that a network calculus with statistical envelopes accounts for most of the available statistical multiplexing gain.

3.2.2 Statistical envelope as random process

One can also use random processes to serve as statistical envelopes. Unlike statistical envelopes defined as non-random functions, the statistical envelopes defined as random processes do not necessarily require error functions to capture the violation probabilities of the bounds.

Statistical envelopes may be defined in two ways, depending on the ordering relationship used between arrivals and envelopes. The first type of envelope is based on the notion of stochastic ordering [107]. A random variable X is stochastically smaller than a random variable Y , and we write this as $X \leq_{st} Y$, if for all real z

$$Pr(X > z) \leq_{st} Pr(Y > z) .$$

Using stochastic ordering, Kurose [73] defines a statistical envelope as follows. A non-negative random process $G(t)$ is a statistical envelope for an arrival process $A(t)$ if for all

$$0 \leq s \leq t$$

$$A(t) - A(s) \leq_{st} G(t - s) . \quad (3.11)$$

We will use the convention that the capital letter G denotes envelopes defined as random processes.

To give an example of a statistical envelope, in a discrete-time setting, consider the Bernoulli source described in Section 3.1 transmitting P data units with probability p in a time unit, and being idle with probability $1 - p$. Then, the corresponding arrival process $A(t)$ is bounded according to

$$A(t) - A(s) \leq_{st} P \cdot B(t - s, p) , \quad (3.12)$$

where $B(t, p)$ is a binomial random variable (counting the number of successes in t independent trials, where the probability of a success is p). Compared to the deterministic envelope from Eq. (3.1), the statistical envelope from Eq. (3.12) captures the long-term average rate of the Bernoulli source by invoking the expectation of binomial random variables. Similar examples of statistical envelopes defined with binomial distributions are provided in [73, 120].

In a continuous time setting, Zhang and Knightly [120] consider a Markov-modulated fluid model with two states (an underlying Markov chain with two states determines, based on the current state, whether a source transmits at some positive rates r_1 or r_2). The statistical envelope bounding the corresponding arrivals is then given by

$$G(t) = \int_0^t \{r_1 I_{s(u)=1} + r_2 I_{s(u)=2}\} du , \quad (3.13)$$

where $s(u)$ denotes the state of the Markov chain at time u , and I_x denotes the indicator function

$$I_x = \begin{cases} 1, & x = \text{true} \\ 0, & x = \text{false} \end{cases} ,$$

for some logical clause x .

The second type of statistical envelope is based on the notion of almost surely ordering.

A random variable X is almost surely smaller than a random variable Y , and we write this as $X \leq Y$ *a.s.*, if

$$\Pr(X > Y) = 0 .$$

A.s. ordering implies stochastic ordering but the converse is not necessarily true.

Using *a.s.* ordering, statistical envelopes can be defined as follows.

Definition 3.2 (STATISTICAL ENVELOPE with *a.s.* Ordering) *A nonnegative, doubly-indexed random process $G(s, t)$ is a statistical envelope for an arrival process $A(t)$ if for all $0 \leq s \leq t$*

$$A(t) - A(s) \leq G(s, t) \text{ a.s. .} \quad (3.14)$$

The random process $G(s, t)$ is assumed to be decreasing in s , increasing in t , and to satisfy $G(s, t) = G(s, u) + G(u, t)$ for all $0 \leq s \leq u \leq t$. To derive performance bounds one usually needs the availability of bounds on the moment generating function (MGF) of $G(s, t)$, i.e., bounds on $E [e^{\theta G(s, t)}]$ for some $\theta > 0$ [29,48]. The linear envelope model from Eq. (3.4) provides with such bounds; see also Eq. (3.37) below concerning with the derivation of backlog bounds.

The definition closely resembles to the definition of the deterministic envelope from Eq. (2.2). The difference is that the envelope is now defined with a doubly-indexed random process. Also, the statistical envelope with *a.s.* ordering belongs to the class of the statistical envelopes with stochastic ordering from Eq. (3.11), because *a.s.* ordering implies stochastic ordering. The two examples of statistical envelopes from Eqs. (3.12) and (3.13), properly rewritten with two indexes, are also examples of statistical envelopes with *a.s.* ordering. For example, in the case of the Bernoulli source, the corresponding statistical envelope with *a.s.* ordering is

$$G(s, t) = P \cdot B(s, t, p) ,$$

where $B(s, t, p)$ is a binomial random variable that counts the number of successes in the interval $(s, t]$.

An immediate construction of statistical envelopes with *a.s.* ordering is the following. Suppose that the cumulative arrivals $A(t)$ are described with a random process $X(t)$ (e.g. a Poisson process with some arrival rate). Then, the corresponding statistical envelope is defined as

$$G(s, t) = X(t) - X(s) ,$$

such that Eq. (3.14) holds with equality. In other words, the random process $X(t)$ is an envelope for itself. Providing an envelope $G(s, t)$ with *a.s.* ordering for an arrival process $A(s, t)$ is particularly useful when there are available bounds on the MGF of $G(s, t)$, but no available immediate bounds on the MGF of $A(s, t)$. If bounds do exist for the MGF of $A(s, t)$ then one should view $A(s, t)$ as the envelope itself (this implicit view is adopted by Chang [29] and Fidler [48]).

Let us next clarify why statistical envelopes with *a.s.* ordering need to be defined with two indexes. Note that in the case of the Bernoulli source, the realizations of the binomial random variables $B(s, t, p)$ and $B(s + u, t + u, p)$ may be *different* for all $u \neq 0$. In other words, the interarrivals $A(s, t)$ and $A(s + u, t + u)$ are not correlated. Instead, if statistical envelopes with *a.s.* ordering were defined with a single index, e.g., in the case of the Bernoulli source

$$G(t) = P \cdot B(t, p) ,$$

then the interarrivals would be subject to correlations. Concretely, if for some sample-path the realization of the random variable $G(t)$ were *small*, then $A(t + s) - A(s)$ would have to be *small* as well for all $s \geq 0$.

Since the statistical envelopes reviewed in this section are defined as random processes, they capture the statistical properties of arrivals. More importantly, they can account for statistical multiplexing, as pointed out in [73, 120] for the case of statistical envelopes defined with stochastic ordering, and in [48] for the case of statistical envelopes defined with *a.s.* ordering. It is not yet studied which type of statistical envelopes, i.e., defined as either non-

random functions or random processes, better capture statistical multiplexing.

3.3 Statistical Service Curve

Here we review extensions of deterministic service curves to a probabilistic setting. The resulting *statistical service curves* set *probabilistic* bounds on the service received by a flow, or an aggregate of flows, at a node.

Similar to statistical envelopes, statistical service curves can be either defined as non-random functions, or as random processes. We next discuss these statistical service curves, and then show the construction of statistical leftover service curves.

3.3.1 Statistical service curve as non-random function

Similar to the statistical envelopes as non-random functions from Section 3.2.1, statistical service curves as non-random functions are defined with error functions that capture the violation probabilities of the bounds set on the service.

The following statistical service curve is introduced by Cruz in [40].

Definition 3.3 (STATISTICAL SERVICE CURVE) *A nonnegative, nondecreasing function $\mathcal{S}(t)$ is a statistical service curve with error function $\varepsilon(\sigma)$ for an arrival process $A(t)$ if the corresponding departure process $D(t)$ satisfies for all $t \geq 0$ and $\sigma \geq 0$*

$$\Pr\left(D(t) + \sigma < A * \mathcal{S}(t)\right) \leq \varepsilon(\sigma). \quad (3.15)$$

In the definition from [40], $\mathcal{S}(t)$ is called a *service curve with deficit profile* $\varepsilon(\sigma)$. If $\mathcal{S}(t)$ is a linear function and $\varepsilon(\sigma)$ is an exponential function, then $\mathcal{S}(t)$ is an EBB statistical service curve [21].

Another formulation of a statistical service curve is given by Burchard *et. al.* in [22], and is called *effective service curve*. A nonnegative, nondecreasing function $\mathcal{S}(t)$ is a statistical

service curve with violation probability ε for an arrival process $A(t)$ if the corresponding departure process $D(t)$ satisfies for all $t \geq 0$

$$\Pr\left(D(t) < A * \mathcal{S}(t)\right) \leq \varepsilon . \quad (3.16)$$

An effective service curve differs from the statistical service curve from Definition 3.3 in that by fixing t , Eq. (3.16) sets a single violation probability, whereas Eq. (3.15) sets multiple violation probabilities depending on σ . On the other hand, the statistical service curve from Definition 3.3 reduces to an effective service curve. Indeed, if $\mathcal{S}'(t)$ satisfies Eq. (3.15) with an error function $\varepsilon'(\sigma)$, then the function

$$\mathcal{S}(t) = \mathcal{S}'(t) - \sigma$$

is an effective service curve with violation probability ε if $\varepsilon'(\sigma) = \varepsilon$.

A modified definition of an effective service curve is given by Burchard *et. al.* [22] by using a modified definition of the convolution operator. The modified convolution operator ' $*_t$ ' is defined for all $l > 0$ by

$$A *_t \mathcal{S}(l) = \min \left\{ \mathcal{S}(l), B(t) + \inf_{x \leq l} (A(t, t + l - x) + \mathcal{S}(x)) \right\} , \quad (3.17)$$

where $B(l)$ denotes the backlog process at the node. The key difference from the usual convolution operator ' $*$ ' from Eq. (2.6) is that the definition of $A *_t \mathcal{S}(l)$ dispenses with the past arrivals $A(0, t)$, but takes into account the backlog at the node at time t . Note that by setting $t = 0$, the modified convolution operator ' $*_t$ ' reduces to the usual convolution operator.

With the modified convolution operator, the function $\mathcal{S}(t)$ is called an *l-adaptive effective service curve* for intervals of length l , with violation probability ε_l , if the departure process

at the node satisfies for all $t \geq 0$ and $0 \leq x \leq l$

$$Pr\left(D(t, t+x) < A *_t \mathcal{S}(x)\right) \leq \varepsilon_l. \quad (3.18)$$

This definition is relaxed in [23] where an *adaptive statistical service curve* is defined as in Eq. (3.18), but only for $x = l$.

A stronger formulation of an adaptive effective service curve, called *strong adaptive effective service curve*, is defined in [22] by modifying Eq. (3.18) to

$$Pr\left(\forall [t, t+x] \in I_l : D(t, t+x) < A *_t \mathcal{S}(x)\right) \leq \varepsilon_l,$$

where I_l is any interval of length l . The advantage of strong adaptive effective service curves is that their concatenation results in network service curves, in an exact manner that the concatenation of deterministic service curves results in deterministic network service curves [22].

For the rest of the section we review constructions of statistical leftover service curves. In this sense, consider a workconserving node with fixed capacity C serving a tagged flow $A(t)$, and a cross flow $A_c(t)$. Assume SP scheduling at the node, and that the cross flow receives higher priority. As pointed out in Section 2.3, the service curve construction for the low priority flow $A(t)$ holds for all workconserving scheduling mechanisms.

Consider the case when the cross flow is bounded by a global effective envelope $\mathcal{G}_c(t, T, \varepsilon_1)$ for some $T > 0$ and violation probability ε_1 , according to Eq. (3.10). Assume also that T sets a bound on the busy period at the node satisfying for all $t \geq 0$

$$Pr\left(t - \underline{t} > T\right) \leq \varepsilon_2, \quad (3.19)$$

where \underline{t} denotes the beginning of the last busy period containing t . Then, Liebeherr *et.*

al. [80] construct the following statistical leftover service curve for the tagged flow

$$\mathcal{S}(t) = Ct - \mathcal{G}_c(t, T, \varepsilon_1) \quad (3.20)$$

with violation probability $\varepsilon_1 + \varepsilon_2$, that complies with Eq. (3.16).

A drawback of the construction from Eq. (3.20) is that the bound T on the busy period requires the availability of a deterministic envelope for the aggregate arrivals $A(t) + A_c(t)$. Unless the deterministic envelope provides with an accurate characterization for the aggregate arrivals, the busy period bound T may be too pessimistic. Moreover, since the violation probability ε_1 corresponding to the statistical envelope $\mathcal{G}_c(t, T, \varepsilon_1)$ is usually proportional with T , the violation probability $\varepsilon_1 + \varepsilon_2$ corresponding to $\mathcal{S}(t)$ may be too large.

Consider now the case when the cross flow is bounded by an effective envelope $\mathcal{G}_c(t, \varepsilon_1)$ with violation probability ε_1 , complying with Eq. (3.5). Assume also that T is a bound on the busy period satisfying Eq. (3.19). Then, a statistical leftover service curve for the tagged flow, complying with Eq. (3.16), is given by

$$\mathcal{S}(t) = Ct - \mathcal{G}_c(t, \varepsilon_1) , \quad (3.21)$$

with an error function that is proportional to the busy period bound T (see Li *et. al.* [76]). Unlike the construction from [80], the busy period bound obtained in [76] does not require the availability of deterministic envelopes for the aggregate arrivals.

If the cross flow is bounded by a sample-path effective envelope $\mathcal{G}_c(t, \varepsilon_1)$, complying with Eq. (3.9), then the function $\mathcal{S}(t)$ in Eq. (3.21) is an adaptive effective service curve for any interval of length l with error function ε (see Burchard *et. al.* [23]).

For the last construction of a statistical leftover service curve as a non-random function, consider the case when the cross flow is described with gSBB envelopes. The next theorem from Liu *et. al.* [82] provides such as construction.

Theorem 3.4 (STATISTICAL LEFTOVER SERVICE CURVE *as non-random function*) *As-*

sume that the cross flow $A_c(t)$ is bounded by a gSBB envelope $\mathcal{G}_c(t) = rt$ with error function $\varepsilon(\sigma)$ according to Eq. (3.8). Then, a statistical service curve given by the node to the tagged flow $A(t)$ is given by the non-random function

$$\mathcal{S}(t) = [Ct - \mathcal{G}_c(t)]_+ , \quad (3.22)$$

with error function $\varepsilon(\sigma)$, according to Definition 3.3.

Let us next sketch the proof. Fix $t, \sigma \geq 0$ and assume that on a particular sample-path the inequalities

$$A_c(t) - A_c(s) \leq \mathcal{G}_c(t - s) + \sigma \quad (3.23)$$

hold for all $0 \leq s \leq t$.

When constructing the *deterministic* leftover service curve in the proof sketch of Theorem 2.3 we showed the inequality

$$D(t) \geq \inf_{0 \leq s \leq t} (A(s) + [C(t - s) - (A_c(t) - A_c(s))]_+)$$

Using Eq. (3.23) it immediately follows that

$$D(t) + \sigma \geq A * \mathcal{S}(t) .$$

Therefore

$$\begin{aligned} Pr\left(D(t) + \sigma < A * \mathcal{S}(t)\right) &\leq Pr\left(\text{Eq. (3.23) fails}\right) \\ &\leq \varepsilon(\sigma) , \end{aligned}$$

i.e., $\mathcal{S}(t)$ is a statistical leftover service curve for the tagged flow.

Although statistical leftover service curves provide worst-case probabilistic bounds on the service, they lead to similar performance bounds, at high data rates, to the performance

bounds obtained with statistical service curves constructed for EDF and GPS (Generalized Processor Sharing) scheduling; corresponding constructions are provided in [76]. It thus appears that scheduling is dominated by statistical multiplexing (see Liebeherr [77], Li *et al.* [76]).

3.3.2 Statistical service curve as random process

To define statistical service curves as random processes we first define a convolution operator for two doubly-indexed functions f and g as

$$f * g(u, t) \triangleq \inf_{u \leq s \leq t} \{f(u, s) + g(s, t)\} .$$

for all $0 \leq u \leq t$. If $u = 0$, we make the convenient notation

$$f * g(t) \triangleq f * g(0, t) .$$

A definition of a statistical service curve is given by Chang [29].

Definition 3.5 (STATISTICAL SERVICE CURVE with *a.s. Ordering*) *A nonnegative, doubly-indexed random process $S(s, t)$ is a statistical service curve for an arrival process $A(t)$ if the corresponding departure process $D(t)$ satisfies for all $t \geq 0$*

$$D(t) \geq A * S(t) \text{ a.s. .} \tag{3.24}$$

In the terminology from [29], the process $S(s, t)$ is called a dynamic F-server.

We will use the convention that the capital letter S denotes service curves defined as random processes. The random process $S(s, t)$ is decreasing in s , increasing in t , and satisfies $S(s, t) = S(s, u) + S(u, t)$ for all $0 \leq s \leq u \leq t$. As in the case of statistical envelopes as random processes, statistical service curves as random processes usually require the availability

of bounds on the moment generating function of $S(s, t)$, in particular bounds on $E [e^{-\theta S(s,t)}]$ for some $\theta > 0$ [29, 48].

The representation with two indexes is necessary to capture the variability of service over intervals of same lengths. This is similar to the representation of the statistical envelopes from Definition 3.2 with doubly-indexed random processes $G(s, t)$ (see the discussion concerning doubly-indexed vs. single-indexed processes from Section 3.2.2). We point out that a representation of statistical service curves with single-indexed random processes $S(t)$ is introduced by Qiu *et. al.* [95, 96] using stochastic ordering.

The next theorem (from Fidler [48]) provides a construction of leftover service curves as random processes.

Theorem 3.6 (STATISTICAL LEFTOVER SERVICE CURVE *as random process*) *Assume that the cross flow $A_c(t)$ is bounded by a statistical envelope $G_c(s, t)$, according to Definition 3.2. Then, a service curve given by the node to the tagged flow $A(t)$ is represented by the doubly-indexed random process*

$$S(s, t) = [C(t - s) - G_c(s, t)]_+ . \quad (3.25)$$

The obtained service curve $S(s, t)$ complies with the formulation from Definition 3.5. It is generally useful in applications when $A(t)$ and $G_c(s, t)$ are statistically independent. Similar constructions are available when the arrivals are described with single-indexed random processes and stochastic ordering [96].

We point out that the statistical representation of service leads to statistical multiplexing gain (see Qiu *et. al.* [95, 96] for the case of statistical service curves defined as random processes, and Liebeherr *et. al.* [80] for the case of statistical service curves defined as non-random functions).

3.4 Single-Node Performance Bounds

The stochastic network calculus provides probabilistic performance bounds for three main scenarios that depend on whether statistical descriptions are given for arrivals, service, or both.

The first scenario describes arrivals with statistical envelopes, while the representation of service is either implicitly given by a scheduling algorithm (see Kurose [73], Yaron and Sidi [116,117], Chang [26], Zhang and Knightly [120], Zhang *et. al.* [121], Knightly [69–71], Starobinski and Sidi [105], Boorstyn *et. al.* [13], Yin *et. al.* [118], Liebeherr [77], Yu *et. al.* [119]), or is explicitly given with deterministic service curves (see Qiu *et. al.* [95,96], Ayyorgun and Feng [5,6], Liu *et. al.* [82], Jiang and Emstad [62]). The second scenario describes arrivals with deterministic envelopes and service with statistical service curves (see Burchard *et. al.* [22,23]); this resembles much to the single-node deterministic network calculus. The third scenario describes arrivals with statistical envelopes, service with statistical service curves, and it is a relatively simple extension at a single node of the first scenario (see Qiu *et. al.* [95,96], Chang [29], Li *et. al.* [76], Jiang *et. al.* [60,61], Fidler [48]).

Probabilistic performance bounds can also be obtained in scenarios with deterministic envelopes and deterministic service. The key idea is to account for the statistical multiplexing gain characteristic to multiplexing multiple independent regulated flows. For this reason, this scenario reduces to the first scenario above. Some authors address the statistical multiplexing for several scheduling algorithms such as SP (see Knightly [71], Boorstyn *et. al.* [13]), GPS (see Elwalid and Mitra [44]), or EDF (see Sivaraman *et. al.* [103,104]). Kesidis and Konstantopoulos [65] derive probabilistic backlog bounds for independent regulated flows at constant rate servers. These bounds are improved by Chang *et. al.* [30]. Further improvements of the bounds, and extensions to servers described by service curves are provided by Vojnovic and Le Boudec [111,112]. Analysis of statistical multiplexing for regulated sources at bufferless links is investigated by Elwalid *et. al.* [45], Lo Presti *et. al.* [94], or Reisslein *et. al.* [97].

This section deals with the scenario when the arrivals are represented by statistical en-

velopes, and service are represented by deterministic service curves. This scenario is representative for exposing the major difficulties in the stochastic network calculus, and the corresponding techniques for the derivation of single-node performance bounds.

Li *et. al.* [76] isolate the key problem in the stochastic network calculus analysis of bounds. Suppose that an arrival process $A(t) = \sum_i A_i(t)$ is served at a node with capacity C , i.e., the node can be represented by a constant-rate deterministic service curve $\mathcal{R}(t) = Ct$. Reich's equation gives the expression of the backlog process $B(t)$ for all $\sigma \geq 0$

$$Pr\left(B(t) > \sigma\right) = Pr\left(\sup_{0 \leq s \leq t} (A(s, t) - C(t - s)) > \sigma\right). \quad (3.26)$$

The difficulty in evaluating the right-hand side of Eq. (3.26) is that the value s^* attaining the supremum is a random variable. This is a fundamental reason for which the stochastic network calculus is considered to be hard.

Next we review four network calculus techniques to evaluate the backlog bound from Eq. (3.26), in a discrete-time setting.

3.4.1 With lower bound approximation

Boorstyn *et. al.* [13] discuss the following approximation to evaluate Eq. (3.26)

$$Pr\left(\sup_{0 \leq s \leq t} (A(s, t) - C(t - s)) > \sigma\right) \approx \sup_{0 \leq s \leq t} Pr\left(A(s, t) - C(t - s) > \sigma\right). \quad (3.27)$$

The right-hand side of Eq. (3.27) is a *lower bound* on the corresponding left-hand side, rather than an upper bound. Since we are generally interested in the upper bound, we refer to Eq. (3.27) as an upper bound approximation with a lower bound. Choe and Shroff [31] showed using simulations that this approximation is accurate for Gaussian arrival processes.

Next we review several works in the stochastic network calculus literature that use the approximation from Eq. (3.27). Kurose [73], and Zhang and Knightly [120] use Eq. (3.27)

to derive delay bounds for FIFO and SP scheduling, respectively. In these works the arrivals $A_i(t)$ are bounded with statistical envelopes $G_i(t)$ in the sense of stochastic ordering, according to Eq. (3.11). An approximative upper bound for the backlog takes the expression

$$Pr(B(t) > \sigma) \approx \sup_{0 \leq s \leq t} Pr\left(\sum_i G_i(s, t) - C(t - s) > \sigma\right). \quad (3.28)$$

Knightly [69] proposes a slight improvement to Eq. (3.28) by restricting the range of s in the supremum to the interval $[t - T, t]$, where T is a bound on the busy period. Such a bound, however, requires a deterministic envelope $\mathcal{G}(t)$ for the aggregate arrivals $A(t)$, and can be computed as in [26]

$$T = \inf \{s : \mathcal{G}(t) \leq Ct\}. \quad (3.29)$$

A problem with Eq. (3.28) is that the evaluation of the right-hand side probability, either using convolutions or the Fast Fourier Transform, can be costly when the number of arrival flows $A_i(t)$ is large. Assuming statistical independence for the arrivals, Knightly [70, 71] proposes to evaluate Eq. (3.28) by approximating the sum $\sum_i G_i(t)$ with the Central Limit Theorem and using the rate variance envelopes defined in Eq. (3.6). A similar method is used by Qiu *et. al.* in [95, 96].

The approximation from Eq. (3.27) is discussed by Boorstyn *et. al.* [13] to derive admission regions for several scheduling algorithms. Different from the previously mentioned works, the arrivals $A_i(t)$ are bounded in [13] using statistical envelopes as non-random functions, i.e., effective envelopes $\mathcal{G}_i(t, \varepsilon_i)$ complying with Eq. (3.5). For this reason, the derivation of performance bounds does not encounter the problem of evaluating sums of random variables, as it appears in Eq. (3.28).

We next sketch the argument from Boorstyn *et. al.* [13] for the derivation of a backlog bound. Suppose that for a *fixed* s with $s \leq t$ the inequalities

$$A_i(s, t) \leq \mathcal{G}_i(t - s, \varepsilon_i) \quad (3.30)$$

hold for all i . Set $\mathcal{G}(t) = \sum_i \mathcal{G}_i(t, \varepsilon_i)$ and $\sigma = \mathcal{G} \circ \mathcal{R}(0)$ in Eq. (3.27); recall that $\mathcal{R}(t) = Ct$. It immediately follows that

$$\begin{aligned} Pr\left(B(t) > \mathcal{G} \circ \mathcal{S}(0)\right) &\approx Pr\left(\text{Eq. (3.30) fails for some } i\right) \\ &\leq \sum_i \varepsilon_i . \end{aligned} \quad (3.31)$$

The last equation follows from the definition of the statistical envelopes $\mathcal{G}_i(t, \varepsilon_i)$. It is important to note that this argument requires Eq. (3.30) to hold for a *single* s , because of the approximation from Eq. (3.27). An immediate consequence of dispensing with this approximation is that a proof for the backlog bound would require Eq. (3.30) to hold for *all* the values of s , leading to further complications. These will be addressed in the next section.

3.4.2 With Boole's inequality

A rigorous way to evaluate Eq. (3.26) relies on Boole's inequality, used in conjunction with the Chernoff bound. Let us first state these two results from probability theory. Given n nonnegative random variables X_1, X_2, \dots, X_n , Boole's inequality states that

$$E\left[\sup_{i=1, \dots, n} X_i\right] \leq \sum_{i=1}^n E[X_i] .$$

Another form of the inequality which we frequently use states that for all σ

$$Pr\left(\sup_{i=1, \dots, n} X_i \geq \sigma\right) \leq \sum_{i=1}^n Pr(X_i \geq \sigma) .$$

Given a random variable X , the Chernoff bound states that for any positive θ and all $x > 0$

$$Pr\left(e^{\theta X} \geq x\right) \leq \frac{1}{x} E\left[e^{\theta X}\right] . \quad (3.32)$$

Let us also state the following useful result which gives a bound for series of terms that appear in the violation probabilities of performance bounds. Let a nonnegative, nonincreasing

and integrable function $f(t)$. Then for all t

$$\sum_{s=t+1}^{\infty} f(s) \leq \int_t^{\infty} f(u) du, \quad (3.33)$$

In particular

$$\sum_{s \geq 1} e^{-as} \leq \frac{1}{a}, \quad (3.34)$$

for any $a > 0$. The left-hand side of Eq. (3.34) is a geometric series that converges to

$$\sum_{s \geq 1} e^{-as} = \frac{e^{-a}}{1 - e^{-a}}.$$

In general, we prefer the bound from Eq. (3.34) that lends to simpler formulas than the exact result.

Using Boole's inequality, Eq. (3.26) becomes

$$Pr(B(t) > \sigma) \leq \sum_{s=0}^{t-1} Pr(A(s, t) > C(t-s) + \sigma). \quad (3.35)$$

The right-hand side of Eq. (3.35) can be further evaluated depending on the types of statistical envelopes used for the aggregate arrivals $A(t)$. Yaron and Sidi [116] evaluate the right-hand side of Eq. (3.35) for the special case when the arrivals $A(t)$ are bounded by an EBB envelope with rate $r < C$ and error function $\varepsilon(\sigma) = Me^{-\theta\sigma}$, as follows

$$\begin{aligned} Pr(B(t) > \sigma) &\leq \sum_{s=0}^{t-1} Pr(A(s, t) > r(t-s) + (C-r)(t-s) + \sigma) \\ &\leq \sum_{s \geq 1} Me^{-\theta(C-r)s} e^{-\theta\sigma} \\ &\leq \frac{M}{\theta(C-r)} e^{-\theta\sigma}. \end{aligned} \quad (3.36)$$

The first equation represents the capacity C as $r+(C-r)$, which permits the direct application of the EBB definition in the second equation. The last equation uses the inequality from

Eq. (3.34).

A similar derivation for the backlog bound, i.e., as in Eq. (3.36), can be given for the case when the arrivals are bounded with SBB envelopes. The obtained bound is always finite because of the integrability condition on the corresponding error function $\varepsilon(\sigma)$ (see Starobinski and Sidi [105]).

A related derivation of the backlog bound is given by Chang [26]. Suppose that the arrivals $A(t)$ are bounded by a linear envelope process with rate $r < C$ and burst σ' , for a choice of $\theta > 0$, according to Eq. (3.4). Then, a bound on the backlog process can be derived from Eq. (3.26) as follows

$$\begin{aligned}
 Pr(B(t) > \sigma) &\leq Pr(e^{\theta \sup_{0 \leq s < t} (A(s,t) - C(t-s))} > e^{\theta \sigma}) \\
 &\leq E[e^{\theta \sup_{0 \leq s < t} (A(s,t) - C(t-s))}] e^{-\theta \sigma} \\
 &\leq \sum_{s \geq 1} E[e^{\theta(A(s) - Cs)}] e^{-\theta \sigma} \\
 &\leq \frac{e^{\theta \sigma'}}{\theta(C - r)} e^{-\theta \sigma}, \tag{3.37}
 \end{aligned}$$

for all $\sigma \geq 0$. The first equation uses that $\theta > 0$, and then the Chernoff bound and Boole's inequality are applied. The rest follows as in Eq. (3.36). The same bound can be obtained by first representing the linear envelope process with an EBB envelope, as shown in Section 3.2.1, and then invoking Eq. (3.36).

The steps illustrated in Eq. (3.37) for the derivation of a backlog bound can be reproduced in the case of statistical envelopes $G(s, t)$ defined as random processes with *a.s.* ordering (Eq. (3.14)), subject to the availability of bounds on $E[e^{\theta G(s,t)}]$ for some choice of $\theta > 0$ (see Chang [29] and Fidler [48]). The derivation from Eq. (3.37) can also be reproduced in the case when service is described with statistical service curves $S(s, t)$ with *a.s.* ordering (Eq. (3.24)), subject to the availability of bounds on $E[e^{-\theta S(s,t)}]$ for some choice of $\theta > 0$ (see Chang [29] and Fidler [48]); the term $C(t - s)$ in the first line of Eq. (3.37) is to be replaced by $S(s, t)$. In the case of statistical envelopes $G(s, t)$ and service curves $S(s, t)$ as

random processes, these are further required to be statistically independent.

3.4.3 With a-priori busy period bounds and Boole's inequality

Suppose here that the arrivals $A(t)$ are bounded by an effective envelope $\mathcal{G}(t, \varepsilon)$ with violation probability ε , according to Eq. (3.5). Following the argument from Li *et. al.* [76], we first show that without a-priori bounds on the busy period the direct use of Boole's inequality to derive backlog bounds may lead to infinite estimates of the violation probabilities.

Assume that for a particular *sample-path* the following inequalities

$$A(s, t) \leq \mathcal{G}(t - s, \varepsilon) \quad (3.38)$$

hold for all $0 \leq s \leq t$. Then Reich's equation (Eq. (3.26)) yields

$$B(t) \leq \mathcal{G} \circ \mathcal{R}(0),$$

such that we arrive at

$$\begin{aligned} Pr\left(B(t) > \mathcal{G} \circ \mathcal{R}(0)\right) &\leq Pr\left(\text{Eq. (3.38) fails}\right) \\ &\leq \sum_{s \leq t} \varepsilon. \end{aligned} \quad (3.39)$$

Recall that $\mathcal{R}(t) = Ct$ is a deterministic service curve for a node with capacity C . The last equation follows from Boole's inequality. By taking $t \rightarrow \infty$, the estimate for the violation probability in Eq. (3.39) becomes unbounded making the corresponding backlog bound not useful.

A solution to the problem of obtaining infinite violation probabilities is provided by Li *et. al.* [76]. The idea is to use an a-priori bound T on the busy period satisfying Eq. (3.19) with some violation probability ε_2 . Then, to obtain a backlog bound with *finite* violation

probability we can follow the steps from Eqs. (3.38)-(3.39) and replace Eq. (3.38) with

$$A(s, t) \leq \mathcal{G}(t - s, \varepsilon) \text{ for all } t - T \leq s \leq t.$$

The backlog bound takes now the expression

$$Pr(B(t) > \mathcal{G} \circledast R(0)) \leq T\varepsilon + \varepsilon_2,$$

and the violation probability $T\varepsilon + \varepsilon_2$ is always finite. The term $T\varepsilon$ stems from applying Boole's inequality for T times, rather than for t times as in Eq. (3.39). The term ε_2 stems from the a-priori bound T on the busy period. The value of ε_2 is 0 when a deterministic envelope $\mathcal{G}(t)$ for the arrivals $A(t)$ is available (see the construction of T from Eq. (3.29)). Otherwise, the value of ε_2 can be obtained as the *finite* sum of a convergent series (see Lemma 1 in [76]).

The main reason for the statistical envelopes $\mathcal{G}(t, \varepsilon)$ require a-priori busy period bounds is that the envelopes are defined for a *single* violation probability ε . In contrast, EBB or SBB envelopes do not necessarily require such a-priori bounds since they are defined for *multiple* violation probabilities $\varepsilon(\sigma)$. This flexibility leads to convergent series when applying Boole's inequality to derive backlog bounds (see for instance Eq. (3.36)).

3.4.4 With sample-path statistical envelope

Perhaps the most straightforward way to derive performance bounds in a calculus with statistical envelopes and deterministic service curves is by assuming the existence of gSBB [5, 118] or sample-path effective envelopes [23].

First we consider the case of gSBB envelopes. The next theorem, restating results from [5, 118] in our notations, gives probabilistic bounds for the output envelope, backlog and delay processes at a node.

Theorem 3.7 (PROBABILISTIC PERFORMANCE BOUNDS - *from gSBB envelopes*) *Consider a flow at a node with arrivals and departures given by the processes $A(t)$ and $D(t)$, respec-*

tively. Assume that the arrivals are bounded by the gSBB envelope $\mathcal{G}(t) = rt$ with error function $\varepsilon(\sigma)$, according to Eq. (3.8). If the node provides the deterministic service curve $\mathcal{S}(t)$ to the flow, then the following probabilistic bounds hold.

1. **OUTPUT ENVELOPE:** The function $\mathcal{G} \circledast \mathcal{S}$ is an SBB statistical envelope for the departures $D(t)$, i.e., for all $0 \leq s \leq t$ and $\sigma \geq 0$

$$\Pr\left(D(t) - D(s) > \mathcal{G} \circledast \mathcal{S}(t - s) + \sigma\right) \leq \varepsilon(\sigma) ,$$

subject to the condition that the error function is n -fold integrable.

2. **BACKLOG BOUND:** A probabilistic bound on the backlog process $B(t)$ is given for all $t, \sigma \geq 0$ by

$$\Pr\left(B(t) > \mathcal{G} \circledast \mathcal{S}(0) + \sigma\right) \leq \varepsilon(\sigma) .$$

3. **DELAY BOUND:** A probabilistic bound on the delay process $W(t)$ is given for all $t, \sigma \geq 0$ by

$$\Pr\left(W(t) > d(\sigma)\right) \leq \varepsilon(\sigma) ,$$

where

$$d(\sigma) = \inf \{d : \mathcal{G}(s) + \sigma \leq \mathcal{S}(s + d) \text{ for all } s \geq 0\} .$$

The theorem hold for more general expressions of $\mathcal{G}(t)$. Also, as shown in Section 3.2.1, a gSBB output envelope can be further derived from the obtained SBB output envelope. The obtained probabilistic bounds in the theorem are similar to the corresponding deterministic bounds from Theorem 2.4. This indicates that, in the single-node case, the analysis with the deterministic network calculus carries over to a probabilistic setting.

Following [118] we next sketch the proof for the backlog bound for a general expression of $\mathcal{G}(t)$. This is similar to the corresponding proof for the backlog with the deterministic network calculus (see Eq. (2.10)). Fix t, σ and assume that for a particular sample-path the

following inequalities

$$A(t) - A(s) \leq \mathcal{G}(t - s) + \sigma \quad (3.40)$$

hold for all $0 \leq s \leq t$. Then, Theorem 2.4 yields the (deterministic) backlog bound

$$B(t) \leq \mathcal{G} \circledast \mathcal{S}(0) + \sigma ,$$

such that

$$\begin{aligned} Pr\left(B(t) > \mathcal{G} \circledast \mathcal{S}(0) + \sigma\right) &\leq Pr\left(\text{Eq. (3.40) fails}\right) \\ &\leq Pr\left(\sup_{0 \leq s \leq t} (A(t) - A(s) - \mathcal{G}(t - s)) > \sigma\right) \\ &\leq \varepsilon(\sigma) . \end{aligned}$$

We note that the obtained backlog bound simply follows from the definition of the gSBB envelope $\mathcal{G}(t)$ (see Eq. (3.8)). Moreover, if $\mathcal{G}(t) = rt$ with $r \leq C$, the backlog bound becomes [5]

$$Pr\left(B(t) > \sigma\right) \leq \varepsilon(\sigma) .$$

The apparent simplicity of the bounds from Theorem 3.7 comes at the cost of requiring the availability of gSBB envelopes, which can be constructed from SBB envelopes (see Section 3.2.1). We point out that the performance bounds obtained from gSBB envelopes, constructed from SBB envelopes, match exactly the performance bounds obtained from SBB envelopes (e.g. as in Eq. (3.36)).

Let us next suppose that a flow is bounded by a sample-path effective envelope $\mathcal{G}(t, \varepsilon_g)$ satisfying for some $\varepsilon_g > 0$

$$Pr\left(\sup_{0 \leq s \leq t} (A(t) - A(s) - \mathcal{G}(t - s)) > 0\right) \leq \varepsilon_g ,$$

according to Eq. (3.9). Also, assume that a node offers the flow a statistical adaptive service

curve $\mathcal{S}(t)$ with violation probability ε_l for intervals of length l (according to Eq. (3.18)). If the parameter l satisfies the condition

$$\mathcal{G}(l) < \mathcal{S}(l) ,$$

then one can derive single-node performance bounds for the flow (see Burchard *et. al.* [23]).

For instance, a probabilistic bound on the backlog is given by

$$\Pr\left(B(t) > \mathcal{G} \circ \mathcal{S}(0)\right) \leq \varepsilon_g + \varepsilon_l \frac{\mathcal{S}(l)}{\mathcal{S}(l) - \mathcal{G}(l)} .$$

3.5 The Problem of Statistical Network Service Curve

In this section we discuss existing solutions for the problem of formulating statistical network service curves from statistical service curves descriptions. We consider both cases of statistical service curves defined as non-random functions and random processes. As in a deterministic context, these network service curve are useful to reduce the end-to-end network analysis to a single-node analysis.

The search for statistical network service curves has been motivated by the need to reproduce the benefits of deterministic network service curves (see Section 2.5) in a probabilistic setting. Unlike in the deterministic context, the formulation of *statistical network service curves* from service curves defined as non-random functions was shown to be a difficult problem (see Li *et. al.* [76], Burchard *et. al.* [23]). Let us first illustrate the technical challenge of the problem by describing the argument presented in [23, 76].

Consider a through flow crossing a network with H nodes as in Figure 2.5. We denote the arrivals and departures at nodes $h = 1, 2, \dots, H$ by the processes $A^h(t)$ and $D^h(t)$, respectively. Each node h offers the flow a statistical service curve $\mathcal{S}^h(t)$ according to Eq. (3.16), i.e.,

$$\Pr\left(D^h(t) < A^h * \mathcal{S}^h(t)\right) \leq \varepsilon$$

for some $\varepsilon > 0$.

The straightforward application of Boole's inequality to derive a statistical network service curve leads to *unbounded* violation probabilities. Consider the derivation at the first two nodes only; the derivation can then be inductively extended to the entire network. Assume that for some $t \geq 0$

$$D^2(t) \geq A^2 * \mathcal{S}^2(t), \quad (3.41)$$

and that for a particular sample-path the following inequalities

$$D^1(s) \geq A^1 * \mathcal{S}^1(s) \quad (3.42)$$

hold for all $0 \leq s \leq t$.

Expanding the convolution in Eq. (3.41) and then inserting Eq. (3.42) yields

$$D^2(t) \geq \inf_{0 \leq s \leq t} \{A^2(s) + \mathcal{S}^2(t-s)\} \quad (3.43)$$

$$\geq \inf_{0 \leq s \leq t} \{A^1 * \mathcal{S}^1(s) + \mathcal{S}^2(t-s)\} \quad (3.44)$$

$$\geq A^1 * \mathcal{S}^1 * \mathcal{S}^2(t). \quad (3.45)$$

Since we started with Eqs. (3.41)-(3.42) we arrive at

$$\begin{aligned} Pr(D^2(t) < A^1 * \mathcal{S}^1 * \mathcal{S}^2(t)) &\leq Pr(\text{Eq. (3.41) fails}) + Pr(\text{Eq. (3.42) fails}) \\ &\leq \varepsilon + t\varepsilon. \end{aligned}$$

The term $t\varepsilon$ in the violation probability stems from applying the definition of the service curve $\mathcal{S}^1(t)$ for t times, as required by the assumption from Eq. (3.42). By taking $t \rightarrow \infty$ the violation probability becomes unbounded. Consequently, the obtained statistical network service curve has a practical value only for bounded intervals of time.

The key reason for obtaining unbounded estimates for the violation probability is that the value of s attaining the infimum in Eq. (3.43) is a random variable. Since this random

variable can take any values in the interval $[0, t]$, the continuation of Eq. (3.43) (as shown in Eqs. (3.44)-(3.45)) requires the sample-path inequalities from Eq. (3.42), whose violation probabilities' sum is unbounded. Although the same argument is used to derive a *deterministic* service curve, the difference is that each of the inequalities in Eq. (3.42) are never violated in a deterministic context.

Based on the above argument, Burchard *et. al.* [23] conclude that the expression for a deterministic network service curve from Eq. (2.11) does not simply carry over in a statistical setting. In the following we review three solutions for formulating statistical network service curves as non-random functions.

One way to solve the problem of unbounded estimates of violation probabilities for statistical network service curves is proposed by Li *et. al.* [76]. The authors assume the existence of an a-priori bound T on the busy period at the first node satisfying Eq. (3.19) with some violation probability ε_2 . Then, replacing Eq. (3.42) with

$$D^1(s) \geq A^1 * \mathcal{S}^1(s) \text{ for all } t - T \leq s \leq t ,$$

and following the steps from Eqs. (3.43)-(3.45) we arrive at

$$Pr (D^2(t) \leq A^1 * \mathcal{S}^1 * \mathcal{S}^2(t)) \leq \varepsilon + T\varepsilon + \varepsilon_2 .$$

The violation probability for the network service curve is now finite for all the values of t . This idea is similar to the idea of obtaining finite violation probabilities for backlog bounds (see Section 3.4.3).

A limitation of the technique with a-priori busy period bounds is that it cannot be directly extended to more than two nodes. Although busy periods bounds can be easily derived at the first node, as mentioned in Section 3.4.3, the derivation of busy period bounds at the downstream nodes requires additional assumptions. In particular, the authors of [76] introduce a delay threshold d^* , such that the fraction of the flow's data exceeding this threshold at each

node h is dropped. With this assumption, the arrivals at node h satisfy for all $s \leq t$

$$A^h(t) - A^h(s) \leq A^1(t) - A^1([s - (h - 1)d^*]_+) .$$

Therefore, busy period bounds at each node h can be obtained as for the first node, but from a *shifted* version of the arrivals $A^1(t)$. Consequently, statistical network service curves can be derived with finite violation probabilities.

Another solution for obtaining statistical network service curves with finite violation probabilities is proposed by Burchard *et. al.* [23]. Rather than relying on assumptions such as a-priori bounds on some performance measures, the authors of [23] use a modified definition of a statistical service curve, i.e., the adaptive statistical service curve from Eq. (3.18), based on a modified definition of the convolution operator from Eq. (3.17).

An important property of the adaptive statistical service curve is that it leads to network service curves expressed as in the deterministic network calculus. In this sense, consider the multi-node scenario from Figure 2.5. Assume that each node h provides the flow a service curve $\mathcal{S}^h(t)$ satisfying Eq. (3.18) with some violation probability ε_l , for intervals of length $l > 0$. Then, the service given to the flow by the network as a whole can be expressed using the network service curve [23]

$$\mathcal{S}^{net}(t) = \mathcal{S}^1 * \mathcal{S}^2 * \dots * \mathcal{S}^H(t) ,$$

satisfying Eq. (3.18) as follows

$$Pr\left(D^{net}(t, t + l) < A^{net} *_t \mathcal{S}^{net}(l)\right) \leq \varepsilon_l(1 + (H - 1)l) .$$

The violation probability above is bounded for all the values of t . The choice for the parameter l is usually made in conjunction with the derivation of performance bounds.

A third formulation of a statistical network service curve, given with a non-random func-

tion, is proposed by Ayyorgun and Cruz [4]. The formulation is based on the concept of a *service curve with loss* that is defined next.

Consider a flow whose arrivals $A(t)$ count the number of packets in the time interval $[0, t]$. Also, let a function $\mathcal{S}(t)$. For the k 'th packet of the flow, we assign the deadline

$$D_k = \inf \{t : A * \mathcal{S}(t) \geq k\} \quad (3.46)$$

Then $\mathcal{S}(t)$ is a service curve curve with loss parameter α if at least an α fraction of the flow's packets meet the deadlines assigned in Eq. (3.46). Moreover, the fraction $1 - \alpha$ of the flow's packets which do not meet their assigned deadlines are dropped.

We recall that the idea of introducing the deadlines for dropping some of the traffic is also used in [76]. In terms of implementation, the dropping policy from [76] appears to be simpler since a single deadline is assigned for all packets, whereas Eq. (3.46) assigns different deadlines for different packets.

To extend a service curve with loss to the multi-node case, consider the network scenario from Figure 2.5. Assume that $\mathcal{S}^h(t)$ is a service curve with loss parameter α_h at node h . Then, the function

$$\mathcal{S}^{net}(t) = \mathcal{S}^1 * \mathcal{S}^2 * \dots * \mathcal{S}^H(t)$$

is a network service curve with loss parameter $\prod \alpha_h$ that is clearly bounded [4].

A limitation of the model is that each node has to implement a scheduling algorithm to guarantee the dropping requirements of the model. Such an algorithm is given in [4] and resembles the SCED (Service Curve Earliest Deadline) scheduling algorithm [100].

Next we review the construction of statistical network service curves from statistical service curves defined as random processes, in particular the model with *a.s.* ordering from Definition 3.5. For such service curves, the construction of network service curves is much simpler than in the case of statistical service curves defined as non-random functions.

Suppose that each node h in the network from Figure 2.5 offers a statistical service curve

$S^h(s, t)$ to the flow. Then, the construction of the corresponding network service curve is straightforward, i.e.,

$$S^{net}(s, t) = S^1 * S^2 * \dots * S^H(s, t). \quad (3.47)$$

This result extends the corresponding result from the deterministic network calculus (see Eq. (2.11)). As in a deterministic setting, the proof of Eq. (3.47) follows from the associativity property of the convolution operator.

We note that the construction from Eq. (3.47) is possible for the service curve model (i.e. Definition 3.5) dispenses with error functions. Unlike in the case of statistical service curve models defined with error functions, the expansion of the convolution in Eqs. (3.43)-(3.45) for service curves defined without error functions does not lend, by default, to unbounded violation probabilities.

The application of the statistical network service curve from Eq. (3.47) to derive end-to-end performance bounds for a flow requires the statistical independence of the service curves S^h for all $h = 1, \dots, H$ (S^h should be also independent of the flow's arrivals). Also, the moment generating functions of the service curves must be finite (this assumption prevents the application of the calculus to heavy-tailed arrivals). Under these assumptions, Fidler [48] provides a calculus based on moment generating functions that yields explicit end-to-end bounds. For the network from Figure 2.1, the end-to-end bounds obtained in [48] grow as $\mathcal{O}(H)$. Therefore, the 'pay-burst-only-once' property, shown in the deterministic network calculus, holds in a probabilistic setting as well.

A long-standing problem in the stochastic network calculus concerns the formulation of statistical network service curves without relying on a-priori bounds on the busy periods [76], or modified definitions of statistical service curves [4, 23], or statistical independence of arrivals/service [29, 48]. In the next chapter we will provide a solution which was developed with Burchard and Liebeherr in [33]. This is the central result of this thesis.

Finally we point out that as in the deterministic network calculus, the stochastic network calculus provides an alternative method to derive end-to-end performance bounds by adding

per-node bounds (see [73, 82, 105, 116]). The derivation of per-node bounds requires the iterative construction of output statistical envelope description for the through flow at each traversed node. However, as we have seen in Section 2.5, the problem with these iterative constructions is that the burstiness of the through flow appears in the expressions for each of the output envelopes, further reflecting into the single-node backlog and delay bounds. Consequently, end-to-end bounds are obtained by adding the burstiness of the through flow for multiple times, and are likely to be overly pessimistic. In fact, we will show in Section 5.1 that for case of the network with cross traffic from Figure 2.1 and EBB envelopes, end-to-end delay bounds scale as $\mathcal{O}(H^3)$. This is even more pessimistic than the $\mathcal{O}(H^2)$ scaling established in Section 2.5 for leaky-bucket envelopes.

Chapter 4

The Construction of a Statistical Network Service Curve

In this chapter we present a stochastic network calculus formulation which is generally suitable to analyze network scenarios where arrivals and service at network nodes may be statistically correlated; nonetheless, statistical multiplexing can be accounted for within aggregates of independent flows. The results from this chapter were developed in a joint work with Burchard and Liebeherr (see [33]).

The novelty of the presented formulation is the construction of a statistical network service curve. Compared to the deterministic network service curve, the constructed statistical network service curve has a slightly reduced rate. This rate relaxation implies bounded error functions corresponding to the network service curves. Accordingly, we claim that the constructed statistical network service curve is a solution to the problem left open at the end of the previous chapter.

Applications of this network calculus with a statistical network service curve will be provided in Chapter 5, where the scaling properties of end-to-end delay bounds in the network with cross traffic from Figure 2.1 are discussed.

For the rest of this chapter we closely follow the structure used in the previous two chapters: we first describe the models for statistical envelopes and service curves, then we show how to derive single-node performance bounds, and finally we provide a construction of a

statistical network service curve.

4.1 Statistical Envelope

We use the network model introduced in Chapter 2. Additionally, we introduce a parameter τ_0 for discretizing time in convolution operations.

To model the arrivals of a flow at a node, we use the following envelope model.

Definition 4.1 (STATISTICAL ENVELOPE) *A nonnegative and nondecreasing function $\mathcal{G}(t)$ is a statistical envelope with error function $\varepsilon(\sigma)$ for an arrival processes $A(t)$ if for all $0 \leq s \leq t$ and all σ*

$$Pr\left(A(t) - A(s) > \mathcal{G}(t - s) + \sigma\right) \leq \varepsilon(\sigma). \quad (4.1)$$

Note that if $\mathcal{G}(t) = rt$ and $\varepsilon(\sigma) = Me^{-\theta\sigma}$ for some $r, M, \theta > 0$, then we have an EBB envelope (see also Definition 3.1).

The error function $\varepsilon(\sigma)$ is nonnegative and nonincreasing. We do not restrict the sign of σ , but we assume without loss of generality that

$$\varepsilon(\sigma) \geq 1 \quad \text{for all } \sigma < -\mathcal{G}(0). \quad (4.2)$$

Moreover, when modelling arrivals in the network from Figure 2.1, we usually require the following integrability condition

$$\int_0^\infty \int_\sigma^\infty \varepsilon(u) du d\sigma < \infty. \quad (4.3)$$

This condition is necessary for instance when deriving end-to-end performance bounds in the network from Figure 2.1. Concretely, if the cross traffic is described as in Definition 4.1, then the derivation of statistical leftover service curves requires the error functions to be integrable. Furthermore, the convolution of the statistical leftover service curves requires the

integrals of the error functions to be integrable themselves, whence the doubly integrability condition from Eq. (4.3).

The statistical envelope model generalizes the EBB and SBB models in that it allows for envelope functions which are not necessarily linear in time. This generalization is particularly useful when constructing statistical envelopes from effective bandwidth expressions which depend on the time parameter, as it is the case for FBM or multiplexed regulated arrivals (see Li *et. al.* [76]). For example, a statistical envelope corresponding to FBM arrivals with Hurst parameter $1/2 < H < 1$ can be written as

$$\mathcal{G}(t) = rt + Kt^H ,$$

where r is the long term rate and K is a constant. The corresponding error function takes the form $\varepsilon(\sigma) = e^{-\sigma^\alpha}$, where $0 < \alpha < 1$ depends on H .

For an integrable error function $\varepsilon(\sigma)$ and a positive number a it is convenient to introduce the function

$$\tilde{\varepsilon}_a(\sigma) = \frac{1}{a} \int_{\sigma}^{\infty} \varepsilon(u) du, \quad (4.4)$$

as an upper bound (see Eq. (3.33)) for the discrete sum

$$\sum_{j=1}^{\infty} \varepsilon(\sigma + ja) .$$

The next lemma makes a connection between the statistical envelope from Definition 4.1 and the sample-path statistical envelope from Eq. (3.8). The result will be used to the construction of statistical leftover service curves and to the derivation of single-node performance bounds.

Lemma 4.2 (SAMPLE-PATH STATISTICAL ENVELOPE) *Assume that $\mathcal{G}(t)$ is a statistical envelope for an arrival process $A(t)$ with an integrable error function $\varepsilon(\sigma)$. Then for any choice*

of $\tau_0 > 0$ and $\delta > 0$

$$Pr \left(A(t) > \inf_{0 \leq u \leq s} \{A(u) + \mathcal{G}(t + \tau_0 - u) + \delta(s + \tau_0 - u) + \sigma\} \right) \leq \tilde{\varepsilon}_{\delta\tau_0}(\sigma), \quad (4.5)$$

for all $0 \leq s \leq t$ and all σ .

The parameter τ_0 is used for discretizing the probability event in Eq. (4.5). The explicit use of τ_0 in Eq. (4.5) will later simplify formulas. By letting $s = t$ and $\tau_0 = 0$ in Eq. (4.5), the result reduces to the corresponding result from [118] for the construction of sample-path statistical envelopes (see also Section 3.2.1). The advantage of this generalization is that it results in output envelopes which have smaller rates than those constructed for the particular case when $s = t$ in Eq. (4.5) (for further technical details see Theorem 4.6).

PROOF. Fix $\delta, \tau_0 > 0$, $0 \leq s \leq t$, and $\sigma \geq 0$. We estimate the event in Eq. (4.5) using Boole's inequality. Since the infimum is taken for continuous-time u , we first discretize the event with step τ_0 .

For $0 \leq u \leq s$ we let $j = \lfloor \frac{s-u}{\tau_0} \rfloor$ be the integer part of $\frac{s-u}{\tau_0}$, so that

$$[s - (j + 1)\tau_0]_+ < u \leq s - j\tau_0.$$

Furthermore, since $A(t)$ and $\mathcal{G}(t)$ are nondecreasing functions, we can bound the probability in Eq. (4.5) as follows

$$\begin{aligned} & P \left(A(t) > \inf_{0 \leq u \leq s} \{A(u) + \mathcal{G}(t + \tau_0 - u) + \delta(s + \tau_0 - u) + \sigma\} \right) \\ & \leq Pr \left(A(t) > \inf_{j=0, \dots, \lfloor \frac{s}{\tau_0} \rfloor} \{A([s - (j + 1)\tau_0]_+) + \mathcal{G}(t - [s - (j + 1)\tau_0]_+) \right. \\ & \quad \left. + \delta(j + 1)\tau_0 + \sigma\} \right) \\ & \leq \sum_{j=0}^{\lfloor \frac{s}{\tau_0} \rfloor} Pr \left(A(t) - A([s - (j + 1)\tau_0]_+) > \mathcal{G}(t - [s - (j + 1)\tau_0]_+) \right. \\ & \quad \left. + \delta(j + 1)\tau_0 + \sigma \right). \end{aligned}$$

In the last equation we applied Boole's inequality. Using the definition of a statistical envelope and the monotonicity of $\varepsilon(\sigma)$, we can apply Eq. (3.33) and bound the sum by

$$\sum_{j=0}^{\infty} \varepsilon(\sigma + (j+1)\delta\tau_0) \leq \frac{1}{\delta\tau_0} \int_{\sigma}^{\infty} \varepsilon(u) du = \tilde{\varepsilon}_{\delta\tau_0}(\sigma), \quad (4.6)$$

which completes the proof. \square

4.2 Statistical Service Curve

To model the service received by a flow at a node, we use the following statistical service curve formulation.

Definition 4.3 (STATISTICAL SERVICE CURVE) *Let $\tau_0 > 0$. A nondecreasing function $\mathcal{S}(t)$ is a statistical service curve with error function $\varepsilon(\sigma)$ for an arrival process $A(t)$ if the corresponding departure process $D(t)$ satisfies for all $t \geq 0$ and all σ*

$$Pr\left(D(t) < A * [\mathcal{S} - \sigma]_+(t + \tau_0)\right) \leq \varepsilon(\sigma). \quad (4.7)$$

The error function $\varepsilon(\sigma)$ is nonnegative and nonincreasing. As in the case of error functions for statistical envelopes, we do not restrict the sign of σ , but we assume without loss of generality that

$$\varepsilon(\sigma) \geq 1 \quad \text{for all } \sigma < \mathcal{S}(\tau_0). \quad (4.8)$$

Moreover, when modelling the service in the network from Figure 2.1, we usually require the integrability condition

$$\int_0^{\infty} \varepsilon(u) du < \infty. \quad (4.9)$$

This condition relates to the doubly-integrability condition from Eq. (4.3) in that error functions corresponding to statistical service curves may be obtained by integrating error functions corresponding to statistical envelopes (see also Theorem 4.5 below).

Definition 4.3 extends Cruz's service model from Definition 3.3 in two ways. First, we impose a positivity constraint in order to simplify the analysis of scenarios with negative service curve functions. Second, we explicitly use a discretization parameter τ_0 and let the convolution span the extended time interval $[0, t + \tau_0]$; this will simplify formulas. We point out that because we extended the interval in the convolution, the formulation from Eq. (4.7) appears stronger than Cruz's formulation; the next lemma shows that there is no loss of generality in doing so.

Lemma 4.4 *Consider an arrival process $A(t)$, the corresponding departure process $D(t)$, and an error function $\varepsilon(\sigma)$. If a nondecreasing function $\hat{\mathcal{S}}(t)$ satisfies*

$$\Pr\left(D(t) < A * \left[\hat{\mathcal{S}} - \sigma\right]_+(t)\right) \leq \varepsilon(\sigma), \quad (4.10)$$

and $\varepsilon(\sigma) \geq 1$ for all $\sigma < \hat{\mathcal{S}}(0)$, then the function

$$\mathcal{S}(t) = \hat{\mathcal{S}}(t - \tau_0)$$

is a statistical service curve in the sense of Definition 4.3, for some $\tau_0 > 0$.

PROOF. Using the positivity of $A(t)$, the proof immediately follows from the relationship

$$A * [\mathcal{S} - \sigma]_+(t + \tau_0) \leq A * \left[\hat{\mathcal{S}} - \sigma\right]_+(t), \quad (4.11)$$

for all $t, \tau_0 \geq 0$, and all σ . Note that the condition from Eq. (4.8) on the error function $\varepsilon(\sigma)$ is satisfied because $\mathcal{S}(\tau_0) = \hat{\mathcal{S}}(0)$. \square

The next theorem provides a construction for statistical leftover service curves. The result will be used to analyze end-to-end performance bounds in the network with cross traffic from Figure 2.1.

First, let us introduce a convenient notation. For a real function $f(t)$ and a real number δ we define the function

$$f_\delta(t) \triangleq f(t) + \delta t . \quad (4.12)$$

The key property of the function $f_\delta(t)$ is that its rate is obtained by adding δ to the rate of the function $f(t)$. We note that the sign of δ is not restricted. In our calculus, the function f can be either an envelope or a service curve. In the former case the parameter δ is usually positive, and in the latter case the parameter δ is usually negative.

Theorem 4.5 (STATISTICAL LEFTOVER SERVICE CURVE) *Consider a node with capacity C serving two arrival processes $A(t)$ and $A_c(t)$, whose corresponding departure processes are $D(t)$ and $D_c(t)$, respectively. Assume that $\mathcal{G}(t)$ is a statistical envelope for $A_c(t)$ with an integrable error function $\varepsilon(\sigma)$. Then for any choice of a discretization parameter τ_0 and $\delta > 0$, the function*

$$\mathcal{S}(t) = Ct - \mathcal{G}_\delta(t) \quad (4.13)$$

is a statistical service curve for $A(t)$ with error function $\tilde{\varepsilon}_{\delta\tau_0}(\sigma - C\tau_0)$.

If the error function $\varepsilon(\sigma)$ satisfies the stronger integrability condition from Eq. (4.3), then the function $\tilde{\varepsilon}_{\delta\tau_0}(\sigma)$ in the theorem satisfies the integrability condition from Eq. (4.9).

PROOF. Fix $\delta, \tau_0 > 0$, $t \geq 0$, and some σ . Assume that for a particular sample-path the following inequality

$$A_c(t) \leq \inf_{0 \leq s \leq t} \{A_c(s) + \mathcal{G}_\delta(t + \tau_0 - s) + \sigma\} . \quad (4.14)$$

holds. We recall from the proof of Theorem 2.3 that

$$D(t) \geq \inf_{0 \leq s \leq t} \{A(s) + [C(t - s) - (A_c(t) - A_c(s))]_+\} . \quad (4.15)$$

Inserting Eq. (4.14) into Eq. (4.15) yields

$$\begin{aligned} D(t) &\geq \inf_{0 \leq s \leq t} \{A(s) + [C(t-s) - \mathcal{G}_\delta(t + \tau_0 - s) - \sigma]_+\} \\ &\geq A * [\mathcal{S} - \sigma]_+(t + \tau_0). \end{aligned}$$

Since we started by assuming Eq. (4.14) we arrive at

$$\begin{aligned} Pr(D(t) < A * [\mathcal{S} - \sigma - C\tau_0]_+(t + \tau_0)) &\leq Pr(\text{Eq. (4.14) fails}) \\ &\leq \tilde{\varepsilon}_{\delta\tau_0}(\sigma). \end{aligned}$$

In the last equation we applied Lemma 4.2 with $s = t$. Furthermore we can replace σ with $\sigma - C\tau_0$ and obtain that $\mathcal{S}(t)$ is a statistical service curve with error function $\tilde{\varepsilon}_{\delta\tau_0}(\sigma - C\tau_0)$.

Lastly, we need to prove that $\tilde{\varepsilon}_{\delta\tau_0}(\sigma - C\tau_0) \geq 1$ whenever $\sigma < \mathcal{S}(\tau_0)$. For such a value of σ we can write

$$\begin{aligned} \tilde{\varepsilon}_{\delta\tau_0}(\sigma - C\tau_0) &\geq \tilde{\varepsilon}_{\delta\tau_0}(\mathcal{S}(\tau_0)) \geq \tilde{\varepsilon}_{\delta\tau_0}(-\mathcal{G}(\tau_0) - \delta\tau_0) \\ &\geq \frac{1}{\delta\tau_0} \int_{-\mathcal{G}(\tau_0) - \delta\tau_0}^{-\mathcal{G}(\tau_0)} \varepsilon(u) du \geq \inf_{-\mathcal{G}(\tau_0) - \delta\tau_0 \leq u < -\mathcal{G}(\tau_0)} \varepsilon(u) \\ &\geq 1. \end{aligned}$$

In the first line we used the monotonicity of $\tilde{\varepsilon}_{\delta\tau_0}(\sigma)$. In the second line we used the monotonicity of $\mathcal{G}(t)$ and reduced the domain of the integration. In the last line we used Eq. (4.2), i.e., $\varepsilon(u) \geq 1$ for all $u < -\mathcal{G}(0)$. \square

4.3 Single-Node Performance Bounds

Here we show how to derive single-node performance bounds for a flow described with a statistical envelope and service curve.

Theorem 4.6 (PROBABILISTIC PERFORMANCE BOUNDS) *Consider a flow at a node with arrivals and departures given by the processes $A(t)$ and $D(t)$, respectively. Assume that $\mathcal{G}(t)$ is a statistical envelope for $A(t)$ with an integrable error function $\varepsilon^a(\sigma)$. Also, the service available to the flow is given by a statistical service curve $\mathcal{S}(t)$ with error function $\varepsilon^s(\sigma)$. Fix a discretization parameter $\tau_0 > 0$, $\delta > 0$, and define for all σ*

$$\varepsilon(\sigma) = \tilde{\varepsilon}_{\delta\tau_0}^a * \varepsilon^s(\sigma) . \quad (4.16)$$

Then the following probabilistic bounds hold.

1. **OUTPUT ENVELOPE:** *The function $\mathcal{G} \circledast \mathcal{S}_{-\delta}$ is a statistical envelope for the departures $D(t)$ with error function $\varepsilon(\sigma)$, i.e., for all $0 \leq s \leq t$ and all σ*

$$Pr\left(D(t) - D(s) > \mathcal{G} \circledast \mathcal{S}_{-\delta}(t - s) + \sigma\right) \leq \varepsilon(\sigma) , \quad (4.17)$$

2. **BACKLOG BOUND:** *A probabilistic bound on the backlog process $B(t)$ is given for all $t \geq 0$ and all σ by*

$$Pr\left(B(t) > \mathcal{G}_\delta \circledast \mathcal{S}(0) + \sigma\right) \leq \varepsilon(\sigma) . \quad (4.18)$$

3. **DELAY BOUND:** *A probabilistic bound on the delay process $W(t)$ is given for all $t, \sigma \geq 0$ by*

$$Pr\left(W(t) > d(\sigma)\right) \leq \varepsilon(\sigma) , \quad (4.19)$$

where

$$d(\sigma) = \inf \{d : \mathcal{G}_\delta(s) + \sigma \leq \mathcal{S}(s + d) \text{ for all } s \geq 0\} . \quad (4.20)$$

The bounds on the backlog and delay processes $B(t)$ and $W(t)$, respectively, do not depend on the time parameter t . Assuming that the system converges to a steady-state, it then follows that the bounds hold as well for the steady-state backlog and delay processes $\lim_{t \rightarrow \infty} B(t)$ and $\lim_{t \rightarrow \infty} W(t)$, respectively.

The literature contains similar results on single-node performance bounds, which are obtained using other formulations of statistical envelopes and service curves (see for instance [23, 82]). The technical contribution of our result is that the output envelope from Eq. (4.17) has a smaller rate than the rates of the output envelopes derived in the literature, usually taking the form $\mathcal{G}_\delta \circ \mathcal{S}(t)$. This improvement stems from our particular characterization of sample-path statistical envelopes in Lemma 4.2.

PROOF. Fix $\delta, \tau_0 > 0$, $0 \leq s \leq t$ and σ . Also, let σ^a and σ^s such that $\sigma^a + \sigma^s = \sigma$.

Assume for the moment that $\sigma^s \geq \mathcal{S}(\tau_0)$. Also, assume that for a particular sample-path the following inequalities

$$A(t) - A(u) \leq \mathcal{G}(t + \tau_0 - u) + \delta(s + \tau_0 - u) + \sigma^a, \quad (4.21)$$

hold for all $0 \leq u \leq s$, and that

$$D(s) \geq A * [\mathcal{S} - \sigma^s]_+(s + \tau_0). \quad (4.22)$$

To derive the output envelope bound we write

$$\begin{aligned} D(t) - D(s) &\leq A(t) - A * [\mathcal{S} - \sigma^s]_+(s + \tau_0) \\ &\leq \sup_{0 \leq u \leq s} \{A(t) - A(u) - [\mathcal{S}(s + \tau_0 - u) - \sigma^s]_+\} \\ &\leq \sup_{0 \leq u \leq s} \{\mathcal{G}(t + \tau_0 - u) - \mathcal{S}_{-\delta}(s + \tau_0 - u) + \sigma^a + \sigma^s\} \\ &\leq \mathcal{G} \circ \mathcal{S}_{-\delta}(t - s) + \sigma. \end{aligned}$$

In the first line we applied Eq. (4.22). In the next line we could restrict the range of the convolution because $\sigma^s \geq \mathcal{S}(\tau_0)$. Then we applied Eq. (4.21) and the definition of the deconvolution operator.

Since we started by assuming Eqs. (4.21) and (4.22), we arrive at

$$\begin{aligned} Pr(D(t) - D(s) > \mathcal{G} \circ \mathcal{S}_{-\delta}(t - s) + \sigma) &\leq Pr(\text{Eqs. (4.21) or (4.22) fail}) \\ &\leq \tilde{\varepsilon}_{\delta\tau_0}^a(\sigma^a) + \varepsilon^s(\sigma^s). \end{aligned} \quad (4.23)$$

In the last equation we applied Lemma 4.2.

The derivation for the backlog bound is similar. Assume that Eqs. (4.21) and (4.22) hold for $s = t$, such that we can write

$$\begin{aligned} B(t) &= A(t) - D(t) \\ &\leq A(t) - A * [\mathcal{S} - \sigma^s]_+(t + \tau_0) \\ &\leq \sup_{0 \leq u \leq t} \{A(t) - A(u) - [\mathcal{S}(t + \tau_0 - u) - \sigma^s]_+\} \\ &\leq \sup_{0 \leq u \leq t} \{\mathcal{G}_\delta(t + \tau_0 - u) - \mathcal{S}(t + \tau_0 - u) + \sigma^a + \sigma^s\} \\ &\leq \mathcal{G}_\delta \circ \mathcal{S}(0) + \sigma. \end{aligned}$$

Since we started by assuming Eqs. (4.21) and (4.22), we arrive at

$$\begin{aligned} Pr(B(t) > \mathcal{G}_\delta \circ \mathcal{S}(0) + \sigma) &\leq Pr(\text{Eqs. (4.21) or (4.22) fail}) \\ &\leq \tilde{\varepsilon}_{\delta\tau_0}^a(\sigma^a) + \varepsilon^s(\sigma^s). \end{aligned} \quad (4.24)$$

In the last equation we applied Lemma 4.2 for $s = t$.

Consider now the case when $\sigma^s < \mathcal{S}(\tau_0)$. From the properties of statistical service curves (see Eq. (4.8)) we get that $\varepsilon^s(\sigma^s) \geq 1$, implying that Eqs. (4.23) and (4.24) still hold. This proves the output envelope and backlog bounds from Eqs. (4.17) and (4.18), respectively.

For the output envelope we also need to prove that $\varepsilon(\sigma) \geq 1$ whenever $\sigma < -\mathcal{G} \circ \mathcal{S}_{-\delta}(0)$. For such a value of σ we have from the definition of the deconvolution operator and the monotonicity of $\mathcal{S}(t)$ that

$$\sigma^a + \sigma^s < -\mathcal{G}(0) + \mathcal{S}(\tau_0).$$

It then follows that either $\sigma^a < -\mathcal{G}(0)$ or $\sigma^s < \mathcal{S}(\tau_0)$. Using Eqs. (4.2) and (4.8) we obtain that either $\varepsilon^a(\sigma^a) \geq 1$ or $\varepsilon^s(\sigma^s) \geq 1$, further yielding $\varepsilon(\sigma) \geq 1$.

Last, to prove the delay bound, we fix $\sigma \geq 0$. Also, fix d satisfying Eq. (4.20), and let $t \geq d$. We assume that Eq. (4.21) holds for $s = t - d$, and for t replaced by $t - d$. Also, assume that Eq. (4.22) holds for $s = t$. We can write

$$\begin{aligned} D(t) - A(t - d) &\geq \inf_{0 \leq u \leq t + \tau_0} \{A(u) - A(t - d) + [\mathcal{S}(t + \tau_0 - u) - \sigma^s]_+\} \\ &\geq \min \left\{ \inf_{0 \leq u \leq t - d} \{-\mathcal{G}_\delta(t - d + \tau_0 - u) + \mathcal{S}(t + \tau_0 - u) - \sigma\}, 0 \right\} \\ &\geq 0. \end{aligned}$$

In the second line we used that the values of the infimum for $u > t - d$ are always nonnegative.

In the last line we used the property of d from Eq. (4.20).

Since we started by assuming Eqs. (4.21) and (4.22), we arrive at

$$\begin{aligned} Pr(W(t) > d) &= Pr(A(t - d) > D(t)) \leq Pr(\text{Eqs. (4.21) or (4.22) fail}) \\ &\leq \tilde{\varepsilon}_{\delta\tau_0}^a(\sigma^a) + \varepsilon^s(\sigma^s). \end{aligned}$$

In the last equation we applied Lemma 4.2 for $s = t - d$. The proof is completed by minimizing after d , σ^a and σ^s .

4.4 Statistical Network Service Curve

Here we present the main result of this chapter, that is the formulation of a statistical network service curve.

Theorem 4.7 (STATISTICAL NETWORK SERVICE CURVE) *Consider a flow traversing H nodes in series as in Figure 4.1. For all $h = 1, \dots, H$ assume that $\mathcal{S}^h(t)$ is a statistical envelope for the flow at node h with error function function $\varepsilon^h(\sigma)$ satisfying the integrability*

condition from Eq. (4.9). Then, for every choice of $\delta > 0$, the function

$$\mathcal{S}^{net} = \mathcal{S}^1 * \mathcal{S}_{-\delta}^2 * \dots * \mathcal{S}_{-(H-1)\delta}^H \quad (4.25)$$

is a statistical network service curve with an error function given by

$$\varepsilon^{net} = \tilde{\varepsilon}_{\delta\tau_0}^1 * \dots * \tilde{\varepsilon}_{\delta\tau_0}^{H-1} * \varepsilon^H \quad (4.26)$$




Figure 4.1: A flow with statistical service curves at multiple nodes.

Let us make some remarks about our result. The formula for the statistical network service curve closely resembles to the formula for the deterministic network service curve given in Theorem 2.5. The main difference is that the rates of the service curve functions $\mathcal{S}^h(t)$ are relaxed by $(h-1)\delta$ for $h = 2, \dots, H$. These relaxations are critical to the derivation of a bounded error function for the network service curve.

The free parameters δ and τ_0 that appear in the theorem can be chosen to optimize single-node performance bounds. For example, the derivation of an end-to-end delay bound can be reduced to

$$Pr(W^{net}(t) > d) \leq \varepsilon(d),$$

where the value of $\varepsilon(d)$ depends on δ and τ_0 , which can be optimized. In practice, optimizing the parameter δ can significantly improve on the bounds, whereas optimizing the parameter τ_0 does not result in significantly smaller bounds. A procedure to optimize these parameters will be provided in Section 5.1.

The proof of the theorem uses the next lemma that constructs sample-path descriptions for statistical service curves. With these descriptions the statistical network service curve will then be derived using a similar argument as for the derivation of a deterministic network ser-

vice curve. In Section 3.4.4 we have shown a similar derivation of single-node performance bounds based on sample-path descriptions for statistical envelopes.

Lemma 4.8 (SAMPLE-PATH STATISTICAL SERVICE CURVE) *Consider a statistical service curve $\mathcal{S}(t)$ with an error function $\varepsilon(\sigma)$ that satisfy the integrability condition from Eq. (4.9). Then for any choice of $\tau_0 > 0$, $\delta > 0$, and all $t \geq 0$ and σ , we have*

$$Pr \left(\sup_{0 \leq s \leq t} \{A * [\mathcal{S} - \delta(t + \tau_0 - s) - \sigma]_+(s) - D(s)\} > 0 \right) \leq \tilde{\varepsilon}_{\delta\tau_0}(\sigma). \quad (4.27)$$

PROOF. Fix $\delta, \tau_0 > 0$, $t \geq 0$ and σ . We evaluate the event in Eq. (4.27) using Boole's inequality. Since the supremum is taken for continuous-time s , we discretize the event with step τ_0 .

For $0 \leq s \leq t$ we let $j = \lfloor \frac{t-s}{\tau_0} \rfloor$ be the integer part of $\frac{t-s}{\tau_0}$. This gives

$$[t - (j + 1)\tau_0]_+ < s \leq t - j\tau_0. \quad (4.28)$$

Furthermore, since $A(t)$, $D(t)$, and $\mathcal{S}(t)$ are nondecreasing functions, we can bound the probability event from Eq. (4.27) as follows

$$\begin{aligned} & Pr \left(\sup_{0 \leq s \leq t} \{A * [\mathcal{S} - \delta(t + \tau_0 - s) - \sigma]_+(s) - D(s)\} > 0 \right) \\ & \leq Pr \left(\sup_{j=0, \dots, \lfloor \frac{t}{\tau_0} \rfloor} \left\{ A * [\mathcal{S} - \delta(j + 1)\tau_0 - \sigma]_+(t - j\tau_0) - D([t - (j + 1)\tau_0]_+) \right\} > 0 \right) \\ & \leq \sum_{j=0}^{\lfloor \frac{t}{\tau_0} \rfloor} P \left(A * [\mathcal{S} - \delta(j + 1)\tau_0 - \sigma]_+(t - j\tau_0) - D([t - (j + 1)\tau_0]_+) > 0 \right). \end{aligned}$$

In the last equation we applied Boole's inequality. Using the definition of a statistical service curve and the monotonicity of $\varepsilon(\sigma)$, we can bound the sum by

$$\sum_{j=0}^{\infty} \varepsilon(\sigma + (j + 1)\delta\tau_0) \leq \frac{1}{\delta\tau_0} \int_{\sigma}^{\infty} \varepsilon(u) du = \tilde{\varepsilon}_{\delta\tau_0}(\sigma), \quad (4.29)$$

which completes the proof. \square

PROOF OF THEOREM 4.7. Fix $\delta, \tau_0 > 0, t \geq 0$, and σ . Let $A^h(t)$ and $D^h(t)$ denote the arrivals and departures, respectively, at node h . Also, denote $A^{net} := A^1$ and $D^{net} := D^H$.

Let σ^h for $h = 1, 2, \dots, H$ such that $\sum_{h=1}^H \sigma^h = \sigma$.

Assume for the moment that $\sigma^H \geq \mathcal{S}^H(\tau_0)$. Also, assume that for a particular sample-path the following inequalities

$$D^h(s) \geq A^h * [\mathcal{S}^h - \delta(t + \tau_0 - s) - \sigma^h]_+(s) \quad (4.30)$$

hold for all $0 \leq s \leq t$ and all $h = 1, 2, \dots, H - 1$, and that

$$D^H(t) \geq A^H * [\mathcal{S}^H - \sigma^H]_+(t + \tau_0). \quad (4.31)$$

Expanding the convolution in Eq. (4.31) yields

$$\begin{aligned} D^H(t) &\geq \inf_{0 \leq s \leq t + \tau_0} \left\{ A^H(s) + [\mathcal{S}^H(t + \tau_0 - s) - \sigma^H]_+ \right\} \\ &\geq \inf_{0 \leq s \leq t} \left\{ A^H(s) + [\mathcal{S}^H(t + \tau_0 - s) - \sigma^H]_+ \right\} \end{aligned} \quad (4.32)$$

The restriction of the infimum in Eq. (4.32) is justified by the assumption that $\mathcal{S}^H(\tau_0) - \sigma^H \leq 0$, and the monotonicity properties of $A^H(t)$, $D^H(t)$ and $\mathcal{S}^H(t)$. Furthermore, by inserting the inequality from Eq. (4.31) for $h = H - 1$, we obtain

$$\begin{aligned} D^H(t) &\geq \inf_{0 \leq u \leq s \leq t} \left\{ A^{H-1}(u) + [\mathcal{S}^{H-1}(s - u) - \delta(t + \tau_0 - s) - \sigma^{H-1}]_+ \right. \\ &\quad \left. + [\mathcal{S}^H(t + \tau_0 - s) - \sigma^H]_+ \right\} \\ &\geq \inf_{0 \leq u \leq s \leq t} \left\{ A^{H-1}(u) + [\mathcal{S}^{H-1}(s - u) + \mathcal{S}_{-\delta}^H(t + \tau_0 - s) - \sum_{h=H-1}^H \sigma^h]_+ \right\}. \end{aligned}$$

In the last equation we used the inequality $[x]_+ + [y]_+ \geq [x + y]_+$ for all real numbers x and y .

After iterating the procedure of inserting the inequalities from Eq. (4.31) for $h = H - 2, \dots, 1$, we eventually obtain

$$D^H(t) \geq \inf_{0 \leq x^1 \leq \dots \leq x^H \leq t} \left\{ A^1(x^1) + [\mathcal{S}^1(x^2 - x^1) + \mathcal{S}_{-\delta}^2(x^3 - x^2) + \dots + \mathcal{S}_{-(H-1)\delta}^H(t + \tau_0 - x^H) - \sum_{h=1}^H \sigma^h]_+ \right\}.$$

The infimum in the last equation can be extended to $[0, t + \tau_0]$, such that we can contract the infimum into a convolution

$$D^{net}(t) \geq A^{net} * [\mathcal{S}^{net} - \sigma]_+(t + \tau_0).$$

Since we started by assuming Eqs. (4.30) and (4.31), we arrive at

$$\begin{aligned} & Pr \left(D^{net}(t) < A^{net} * [\mathcal{S}^{net} - \sigma]_+(t + \tau_0) \right) \\ & \leq \sum_{h=1}^{H-1} Pr(\text{Eq. (4.30) fails for some } s \leq t) + Pr(\text{Eq. (4.31) fails}) \\ & \leq \sum_{h=1}^{H-1} \tilde{\varepsilon}_{\delta\tau_0}^h(\sigma^h) + \varepsilon^H(\sigma^H). \end{aligned} \quad (4.33)$$

In the last equation we applied Lemma 4.8, and the definition of a statistical service curve.

Assume now that $\sigma^H < \mathcal{S}^H(\tau_0)$. From the properties of statistical service curves (see Eq. (4.8)) we get that $\varepsilon^H(\sigma^H) \geq 1$, implying that Eq. (4.33) still holds. This proves that $\mathcal{S}^{net}(t)$ is a statistical network service curve with error function $\varepsilon^{net}(\sigma)$.

Lastly, we need to prove that $\varepsilon^{net}(\sigma) \geq 1$ whenever $\sigma < \mathcal{S}^{net}(\tau_0)$. Let us assume by contradiction that $\sigma^h \geq \mathcal{S}^h(\tau_0)$ for all $h = 1, \dots, H$. It then follows that

$$\sigma = \sum_{h=1}^H \sigma^h \geq \sum_{h=1}^H \mathcal{S}^h(\tau_0) \geq \mathcal{S}^{net}(\tau_0),$$

which contradicts with the choice of $\sigma < \mathcal{S}^{net}(\tau_0)$. It then follows that there exists h such that $\sigma^h < \mathcal{S}^h(\tau_0)$. This implies that $\varepsilon^h(\sigma^h) \geq 1$, and further that $\varepsilon^{net}(\sigma) \geq 1$, which completes the proof. \square

Chapter 5

Scaling Properties of End-to-End Delay Bounds

In this chapter we study the scaling properties of probabilistic end-to-end delay bounds computed using the network calculus with statistical network service curve from Chapter 4. The results presented herein were derived in joint works with Burchard and Liebeherr (see [33] and [21]).

The contribution of this chapter is a set of results demonstrating that, under suitable assumptions on arrivals and service, end-to-end delays in the network with cross traffic from Figure 2.1 grow as $\Theta(H \log H)$. The main assumption is that the arrivals of the through flow are described by an EBB envelope, and its service at the nodes is described by EBB service curves. We first present an abstract result (Theorem 5.1) which establishes the $\mathcal{O}(H \log H)$ upper bound on end-to-end delays. This result is then used to derive end-to-end delay bounds in networks with both a fluid-flow service model (Section 5.1), and a packetized service model (Section 5.2). For a network with packetized service we also establish a corresponding $\Omega(H \log H)$ lower bound on end-to-end delays (Section 5.2.2).

The $\Theta(H \log H)$ scaling behavior of end-to-end bounds established in this chapter is different from the $\Theta(H)$ scaling behavior predicted by other theories for network analysis, such as the deterministic network calculus or queueing networks theory. The difference of the two scaling behaviors stems from the different assumptions on the traffic models and the statistical independence of traffic and service. For example, in the deterministic network

calculus, traffic is described with deterministic envelopes which set worst-case bounds on the arrivals. As a consequence, a given flow receives some worst-case guarantees along a network path. Since these guarantees hold with probability one, end-to-end performance bounds for the flow grow as $\Theta(H)$. In contrast, traffic described with EBB envelopes can be arbitrarily large, but with exponentially decaying probabilities. It is the rare events of some very high bursts which contribute to the extra ‘ $\log H$ ’ factor in the $\Theta(H \log H)$ scaling of end-to-end bounds established in this chapter.

In M/M/1 queueing networks, end-to-end (exact) delays grow as $\Theta(H)$ [68]. The reason is that traffic and service at the queues are assumed to be statistically independent, such that the end-to-end delays can be obtained by adding independent random variables (i.e. the per-node (exact) delays at the queues). We point out that the M/M/1 queueing model falls within the class of EBB envelopes and service curves.

A recent development of the stochastic network calculus, pursued by Fidler [48] using techniques based on moment generating functions proposed by Chang [29], established the $\Theta(H)$ scaling of end-to-end delays. The critical assumption there is that traffic and service at the nodes are statistically independent; also, the arrivals must have bounded MGFs. In contrast, the $\Theta(H \log H)$ result established in this chapter does not rely on the statistical independence assumptions of traffic or service at the nodes.

We mention that, in general, the service of a flow in a packet network at different nodes is correlated because each packet has the same size at the traversed nodes. More exactly, service times at downstream nodes are correlated with packet sizes. One way to enforce the statistical independence of service, as done in M/M/1 (packet) queueing networks, is by assuming that each packet has statistically independent sizes at each traversed node [67]. Another way is by using a fluid-flow service model which dispenses with the notion of a packet, i.e., the service unit is infinitesimal. In contrast, the stochastic network calculus formulation from Chapter 4 does not require the statistical independence of service, and can be thus applied to realistic network scenarios with correlated service.

Assume that the arrival processes in the network from Figure 2.1 are stationary, and that the network is stable, i.e., the average rate of the aggregated traffic at each node is smaller than the service capacity. We also assume the existence of a steady-state W^{net} for the end-to-end delay process $W^{net}(t)$. At a single node the existence of a steady-state is guaranteed under some conditions of stationarity and ergodicity of the arrival processes (see Loynes [84]). In general, the existence of a steady-state W^{net} is not guaranteed, especially because the through and cross traffic may be correlated. We point out that the assumption on the existence of a steady-state does not imply a loss of generality of our results since we derive bounds on $W^{net}(t)$ for all times t ; as these bounds do not depend on t , it is convenient to dispense with the index t in notation.

We will analyze W^{net} through the quantiles of its distribution. For $0 < z < 1$, the z -quantile of the distribution of W^{net} is defined as

$$w^{net}(z) = \inf \{ w : Pr(W^{net} > w) \leq 1 - z \}. \quad (5.1)$$

Figure 5.1 illustrates the z -quantile $w^{net}(z)$. In general we choose values of z that are close to one, i.e., $z = 1 - \varepsilon$ where typical values of ε are 10^{-6} or 10^{-9} .




Figure 5.1: The z -quantile $w^{net}(z)$.

Next we present the main result of this chapter, i.e., an upper bound on end-to-end delays.

Let us first recall that an arrival process $A(t)$ is bounded by an EBB envelope $\mathcal{G}(t) = rt$ if

$$Pr(A(t) - A(s) > \mathcal{G}(t - s) + \sigma) \leq Me^{-\theta\sigma} ,$$

for some values M and θ (see Definition 4.1). Also, $\mathcal{S}(t) = rt$ is an EBB service curve if

$$Pr(D(t) < A * [\mathcal{S} - \sigma]_+(t + \tau_0)) \leq Me^{-\theta\sigma} ,$$

for some values M , θ , and τ_0 (see Definition 4.3).

Theorem 5.1 ($\mathcal{O}(H \log H)$ END-TO-END DELAY BOUNDS). *Consider a flow traversing H nodes in series. Assume that $\mathcal{G}(t) = r_a t$ is an EBB statistical envelope for the flow with error function $\varepsilon^a(\sigma) = M_a e^{-\theta_a \sigma}$, where $r_a, M_a, \theta_a > 0$. Also, for all $h = 1, \dots, H$, assume that $\mathcal{S}(t) = r_s t$ is an EBB statistical service curve for the flow at node h with error function $\varepsilon^s(\sigma) = M_s e^{-\theta_s \sigma}$, where $r_s > r_a$ and $M_s, \theta_s > 0$. Then for each $0 < z < 1$ there exist γ_1 and γ_2 depending on H , where γ_1 is uniformly bounded in H and $\gamma_2 = \mathcal{O}(H)$, such that*

$$Pr\left(W^{net} > \gamma_1 H \log(\gamma_2 H)\right) \leq 1 - z .$$

Explicit values for γ_1 and γ_2 are provided in the proof. We assume the same rate r_s for the service curves at the nodes in order to simplify notation; using different rates does not change the growth behavior of γ_1 and γ_2 . The condition $r_s > r_a$ is used for stability.

The definition of the z -quantile $w(z)$ gives

$$w^{net}(z) \leq \gamma_1 H \log(\gamma_2 H) ,$$

implying the following result.

Corollary 5.2 *The z -quantile $w^{net}(z)$ of the distribution of the steady-state end-to-end*

delay W^{net} satisfies

$$w^{net}(z) = \mathcal{O}(H \log H) .$$

The proof of the theorem uses the following result concerning extreme values of convex functions.

Lemma 5.3 *Let the positive numbers M_j, θ_j for $j = 1, \dots, n$. Then for any $\sigma \geq 0$,*

$$\inf_{\sigma_1 + \dots + \sigma_n = \sigma} \sum_{j=1}^n M_j e^{-\theta_j \sigma_j} = \prod_{j=1}^n (M_j \theta_j w)^{\frac{1}{\theta_j w}} e^{-\frac{\sigma}{w}} ,$$

where $w = \sum_{j=1}^n \frac{1}{\theta_j}$.

PROOF. Fix $\sigma \geq 0$ and let us define the functions

$$f(\sigma_1, \dots, \sigma_n) = \sum_{j=1}^n M_j e^{-\theta_j \sigma_j}$$

and $g(\sigma_1, \dots, \sigma_n) = \sigma_1 + \dots + \sigma_n$. We solve the problem of minimizing the function $f(\sigma_1, \dots, \sigma_n)$ subject to the condition $g(\sigma_1, \dots, \sigma_n) = \sigma$ using the method of Lagrange multipliers. The gradient condition of the Lagrange multipliers method yields

$$\nabla \{f(\sigma_1, \dots, \sigma_n) + \lambda(g(\sigma_1, \dots, \sigma_n) - \sigma)\} = 0 ,$$

where $\lambda > 0$ is to be determined from the equation. We obtain that

$$\lambda = M_j \theta_j e^{-\theta_j \sigma_j} , \tag{5.2}$$

for all $j = 1, \dots, n$.

Let us introduce now $p_j = \frac{1}{\theta_j w}$ for $j = 1, \dots, n$. Since $\sum_{j=1}^n p_j = 1$ we can write

$$\begin{aligned} \lambda &= \prod_{j=1}^n \lambda^{p_j} = \prod_{j=1}^n (M_j \theta_j)^{p_j} e^{-\sum_{j=1}^n p_j \theta_j \sigma_j} \\ &= \prod_{j=1}^n (M_j \theta_j)^{p_j} e^{-\frac{\sigma}{w}}. \end{aligned} \quad (5.3)$$

Using Eqs. (5.2) and (5.3) we finally obtain that

$$\inf_{g(\sigma_1, \dots, \sigma_n) = \sigma} f(\sigma_1, \dots, \sigma_n) = \sum_{j=1}^n \frac{\lambda}{\theta_j} = \prod_{j=1}^n (M_j \theta_j w)^{\frac{1}{\theta_j w}} e^{-\frac{\sigma}{w}}.$$

The solution is unique by the strict convexity of f . \square

PROOF OF THEOREM 5.1. The proof has two steps. We first derive a statistical network service curve for the flow, and then apply single-node results.

Fix $\delta, \tau_0 > 0$ and $t, \sigma \geq 0$. Using Theorem 4.7 we obtain the statistical network service curve

$$\begin{aligned} \mathcal{S}^{net}(t) &= \mathcal{S} * \mathcal{S}_{-\delta} * \dots * \mathcal{S}_{-(H-1)\delta}(t) \\ &= \inf_{x_h \geq 0, x_1 + \dots + x_H = t} \{r_s x_1 + (r_s - \delta)x_2 + \dots + (r_s - (H-1)\delta)x_H\} \\ &= (r_s - (H-1)\delta)t. \end{aligned}$$

The corresponding error function is given by

$$\varepsilon^{net}(\sigma) = \underbrace{\tilde{\varepsilon}_{\delta\tau_0}^s * \dots * \tilde{\varepsilon}_{\delta\tau_0}^s}_{H-1 \text{ times}} * \varepsilon^s(\sigma), \quad (5.4)$$

where

$$\tilde{\varepsilon}_{\delta\tau_0}^s(\sigma) = \frac{1}{\delta\tau_0} \int_{\sigma}^{\infty} M_s e^{-\theta_s u} du = \frac{M_s}{\theta_s \delta\tau_0} e^{-\theta_s \sigma}.$$

Having the statistical envelope $\mathcal{G}(t)$ and the statistical service curve $\mathcal{S}^{net}(t)$, we can now

derive a bound on the end-to-end delay $W(t)$. Let us choose δ such that

$$\delta \leq \frac{r_s - r_a}{H}. \quad (5.5)$$

Then, Eq. (4.20) from Theorem 4.6 yields

$$\begin{aligned} d(\sigma) &= \inf \{d : \mathcal{G}_\delta(s) + \sigma \leq \mathcal{S}^{net}(s+d) \text{ for all } s \geq 0\} \\ &= \inf \{d : (r_a + \delta)(s) + \sigma \leq (r_s - (H-1)\delta)(s+d) \text{ for all } s \geq 0\} \\ &= \frac{\sigma}{r_s - (H-1)\delta}. \end{aligned}$$

It then follows that a delay bound is given by

$$Pr \left(W^{net}(t) > \frac{\sigma}{r_s - (H-1)\delta} \right) \leq \tilde{\varepsilon}_{\delta\tau_0}^a * \varepsilon^{net}(\sigma), \quad (5.6)$$

where

$$\tilde{\varepsilon}_{\delta\tau_0}^a(\sigma) = \frac{1}{\delta\tau_0} \int_\sigma^\infty M_a e^{-\theta_a u} du = \frac{M_a}{\theta_a \delta\tau_0} e^{-\theta_a \sigma}.$$

Expanding the convolution in the right-hand side of Eq. (5.6) yields

$$\begin{aligned} &Pr \left(W^{net}(t) > \frac{\sigma}{r_s - (H-1)\delta} \right) \\ &\leq \inf_{\sigma_1 + \dots + \sigma_{H+1} = \sigma} \left\{ \frac{M_a}{\theta_a \delta\tau_0} e^{-\theta_a \sigma_1} + \sum_{j=2}^H \frac{M_s}{\theta_s \delta\tau_0} e^{-\theta_s \sigma_j} + M_s e^{-\theta_s \sigma_{H+1}} \right\}. \end{aligned}$$

We next compute the infimum with Lemma 5.3. Denoting for convenience

$$\alpha = H\theta_a + \theta_s,$$

replacing σ with $d(r_s - (H-1)\delta)$, and letting $t \rightarrow \infty$ yields

$$Pr(W^{net} > d) \leq \frac{\alpha}{\theta_a} M_a^{\frac{\theta_s}{\alpha}} M_s^{\frac{H\theta_a}{\alpha}} \left(\frac{1}{\theta_s \delta\tau_0} \right)^{\frac{\alpha - \theta_a}{\alpha}} e^{-\frac{\theta_a \theta_s}{\alpha} (r_s - (H-1)\delta)d}. \quad (5.7)$$

For $0 < z < 1$, let us define γ_1, γ_2 as

$$\begin{aligned}\gamma_1 &= \frac{\alpha}{H\theta_a\theta_s(r_s - (H-1)\delta)}, \\ \gamma_2 &= \frac{1}{1-z} \frac{\alpha}{H\theta_a} M_a^{\frac{\theta_s}{\alpha}} M_s^{\frac{H\theta_a}{\alpha}} \left(\frac{1}{\theta_s\delta\tau_0} \right)^{\frac{\alpha-\theta_a}{\alpha}}.\end{aligned}$$

From Eq. (5.5) we obtain that γ_1 is uniformly bounded in H , and $\gamma_2 = \mathcal{O}(H)$.

Finally, by equating the right-hand side of Eq. (5.7) with $1 - z$, we arrive at

$$Pr\left(W^{net} > \gamma_1 H \log(\gamma_2 H)\right) \leq 1 - z,$$

which completes the proof. □

5.1 Network with Fluid Service

In this section we compute explicit end-to-end delay bounds in the network with cross traffic from Figure 2.1, and analyze their scaling properties. We use the method of first deriving a statistical network service curve and then applying single-node results, as in Theorem 5.1. For comparison, we also use a second method based on first deriving per-node bounds and then adding the single-node results.

To derive simple formulas we assume a *fluid service model* throughout this section. In the context of a packet network, the fluid service model pretends that each fraction of a packet becomes available for service as soon as processed upstream. In other words, a packet can be in service at multiple nodes at the same time. The fluid model is thus an approximative model which is generally justified at high data rates (numerical results illustrating the impact of assuming a fluid service model will be presented in Chapter 7).

The main finding of this section is that the method of using a statistical network service curve yields end-to-end delay bounds which grow as $\mathcal{O}(H \log H)$, whereas the method of adding per-node bounds yields end-to-end delay bounds which grow as $\mathcal{O}(H^3)$. The differ-

ence in the two orders of growths provides conclusive evidence that statistical network service curves lend to significantly smaller network bounds than the corresponding network bounds obtained by adding per-node bounds. Moreover, the bounds obtained with the network service curve are always smaller than the bounds obtained by adding single-node results. We recall similar advantages of using network service curves in the deterministic network calculus (see Section 2.5), and in the stochastic network calculus where arrivals and service at the nodes are statistically independent (see Section 3.5).

In the following we compute end-to-end delay bounds using the method of a statistical network service curve. Let us first give a result which will be frequently used for optimization purposes throughout the rest of the thesis.

Lemma 5.4 *Let a positive number a . Then*

$$\inf_{x>0} \frac{e^{ax}}{x} = ea, \quad (5.8)$$

and the extremum is attained at $x = 1/a$.

The proof follows immediately from the convexity of the minimized function.

Consider now the network with cross traffic from Figure 2.1, and assume that the through flow $A(t)$ is bounded by an EBB statistical envelope

$$\mathcal{G}(t) = rt,$$

with error function $\varepsilon(\sigma) = e^{-\theta\sigma}$, where $r, \theta > 0$. Also, at each node h , assume that the cross flow $A_h(t)$ is bounded by an EBB statistical envelope

$$\mathcal{G}_c(t) = r_c t,$$

with error function $\varepsilon_h(\sigma) = e^{-\theta\sigma}$, where $r_c > 0$. Note that all the cross flows have identical envelopes, in order to keep the notation simple. Also, we assume for stability that the capacity

of each node is $C > r + r_c$.

Having the statistical envelopes for the cross flows at each node, Theorem 4.5 gives that for all $\delta, \tau_0 > 0$ the function

$$\mathcal{S}(t) = (C - r_c - \delta)t, \quad (5.9)$$

is a statistical (EBB) service curve at each node h with error function

$$\varepsilon^s(\sigma) = \frac{1}{\delta\tau_0} \int_{\sigma - C\tau_0}^{\infty} e^{-\theta u} du = \frac{e^{\theta C\tau_0}}{\theta\delta\tau_0} e^{-\theta\sigma}.$$

Next, using the EBB statistical envelope and the EBB statistical service curves for the flow just derived, we obtain bounds on the (steady-state) end-to-end delay W^{net} from Theorem 5.1 with the parameters

$$\begin{aligned} r_a &= r, \quad M_a = 1, \quad \theta_a = \theta, \quad \text{and} \\ r_s &= C - r_c - \delta, \quad M_s = \frac{e^{\theta C\tau_0}}{\theta\delta\tau_0}, \quad \theta_s = \theta. \end{aligned}$$

Plugging these values into Eq. (5.7) we obtain for all $d \geq 0$ the delay bound

$$Pr(W^{net} > d) \leq (H + 1) \left(\frac{e^{\frac{\theta}{2}C\tau_0}}{\theta\delta\tau_0} \right)^{\frac{2H}{H+1}} e^{-\frac{\theta}{H+1}(C - r_c - H\delta)d}, \quad (5.10)$$

provided that we choose δ according to Eq. (5.5), i.e.,

$$\delta \leq \frac{C - (r + r_c)}{H + 1}.$$

With Lemma 5.4, the values of τ_0 and δ which minimize the delay bound from Eq. (5.10) are

$$\begin{aligned} \tau_0 &= \frac{2}{\theta C}, \quad \text{and} \\ \delta &= \min \left\{ \frac{2}{\theta d}, \frac{C - (r + r_c)}{H + 1} \right\}. \end{aligned} \quad (5.11)$$

With these values, the end-to-end delay bound becomes

$$Pr(W^{net} > d) \leq (H + 1) \left(\frac{eC}{2\delta} \right)^{\frac{2H}{H+1}} e^{-\frac{\theta}{H+1}(C-r_c-H\delta)d}. \quad (5.12)$$

Setting the right-hand side of Eq. (5.12) equal to $1 - z$ for some $0 < z < 1$, and solving for d , the z -quantile $w^{net}(z)$ of the distribution of W^{net} is bounded by

$$d = \frac{H + 1}{\theta(C - r_c - H\delta)} \log \left(\frac{H + 1}{1 - z} \left(\frac{eC}{2\delta} \right)^{\frac{2H}{H+1}} \right). \quad (5.13)$$

It remains to choose δ . Since the optimal value of δ from Eq. (5.11) depends on d , the value of d in Eq. (5.13) can be determined for numerical purposes using an iterative method. We first let $\delta = \frac{C-(r+r_c)}{H+1}$ in Eq. (5.13) yielding a value d_0 . We then let $d = d_0$ in Eq. (5.11) yielding a value δ which is finally used in Eq. (5.13).

We recall that the order of growth of $w^{net}(z)$ is given by Theorem 5.1, i.e.,

$$w^{net}(z) = \mathcal{O}(H \log H).$$

For the rest of this section we compute for comparison end-to-end delay bounds using the method of adding per-node results.

To derive bounds on the delay process $W^h(t)$ at nodes $h = 1, \dots, H$ with Theorem 4.6, we first need the statistical envelope descriptions for the through flow at each node. We show by induction that a statistical envelope at node h is given by the function

$$\mathcal{G}^h(t) = rt,$$

with error function

$$\varepsilon^{g,h}(\sigma) = \left(\frac{Ce}{\delta} \right)^{\frac{h^2+h-2}{2h}} e^{-\frac{\theta}{h}\sigma},$$

where the parameter $\delta > 0$ satisfies the condition

$$\delta \leq \frac{C - (r + r_c)}{2}.$$

The case $h = 1$ is satisfied by the EBB description of the through flow. Assume now that we have the statistical envelope $\mathcal{G}^h(t)$ at node h and let us derive the statistical envelope at node $h + 1$. Having the statistical service curve $\mathcal{S}(t)$ at node h from Eq. (5.9) and invoking Theorem 4.6, we obtain that

$$\begin{aligned} \mathcal{G}^{h+1}(t) &= \mathcal{G}^h \circledast \mathcal{S}_{-\delta}(t) \\ &= \sup_{s \geq 0} \{r(t+s) - (C - r_c - 2\delta)s\} \\ &= rt \end{aligned}$$

is a statistical envelope at node $h + 1$ with error function

$$\begin{aligned} \varepsilon^{g,h+1}(\sigma) &= \tilde{\varepsilon}_{\delta\tau_0}^{g,h} * \varepsilon^s(\sigma) \\ &= \inf_{\sigma_1 + \sigma_2 = \sigma} \left\{ \frac{h}{\theta\delta\tau_0} \left(\frac{Ce}{\delta} \right)^{\frac{h^2+h-2}{2h}} e^{-\frac{\theta}{h}\sigma_1} + \frac{e^{\theta C\tau_0}}{\theta C\tau_0} e^{-\theta\sigma_2} \right\}. \end{aligned}$$

Applying Lemma 5.3 with $\theta_1 = \theta/h$ and $\theta_2 = \theta$ we arrive at

$$\begin{aligned} \varepsilon^{g,h+1}(\sigma) &= (h+1) \frac{e^{\frac{\theta C\tau_0}{h+1}}}{\theta\delta\tau_0} \left(\frac{Ce}{\delta} \right)^{\frac{h^2+h-2}{2(h+1)}} e^{-\frac{\theta}{h+1}\sigma} \\ &= \left(\frac{Ce}{\delta} \right)^{\frac{h^2+3h}{2(h+1)}} e^{-\frac{\theta}{h+1}\sigma}, \end{aligned} \tag{5.14}$$

which completes the induction. In the last equation we optimized $\tau_0 = (h+1)/(\theta C)$ using Lemma 5.4.

Next we compute bounds on the delay $W^H(t)$ with Theorem 4.6. Eq. (4.20) yields

$$\begin{aligned} d(\sigma) &= \inf \{d : \mathcal{G}_\delta^h(s) + \sigma \leq \mathcal{S}(s+d) \text{ for all } s \geq 0\} \\ &= \inf \{d : (r + \delta)s + \sigma \leq (C - r_c - \delta)(s+d) \text{ for all } s \geq 0\} \\ &= \frac{\sigma}{C - r_c - \delta}. \end{aligned}$$

By letting $t \rightarrow \infty$, it then follows that for any $d > 0$ we have the (steady-state) per-node delay bounds

$$Pr(W^h > d) \leq \varepsilon^{g,h+1} ((C - r_c - \delta)d).$$

Adding the per-node delay bounds we obtain for the (steady-state) end-to-end delay process W^{net}

$$Pr(W^{net} > d) \leq \inf_{d_1 + \dots + d_H = d} \left\{ \sum_{h=1}^H Pr(W^h > d_h) \right\}.$$

Evaluating the infimum with Lemma 5.3 we obtain

$$Pr(W^{net} > d) \leq K e^{-\frac{2\theta}{H(H+3)}(C-r_c-\delta)d}, \quad (5.15)$$

where

$$K = \frac{H(H+3)}{2} \left(\frac{Ce}{\delta} \right)^{\frac{(H+1)(H+5)}{3(H+3)}} \prod_{h=1}^H (h+1)^{-\frac{2(h+1)}{H(H+3)}}. \quad (5.16)$$

The optimal value of δ which minimizes the delay bound from Eq. (5.15) is obtained with Lemma 5.4, i.e.,

$$\delta = \min \left\{ \frac{H(H+1)(H+5)}{6\theta d}, \frac{C - (r + r_c)}{2} \right\}. \quad (5.17)$$

Setting the right-hand side of Eq. (5.15) equal to $1 - z$, the z -quantile $w^{net}(z)$ of the distribution of W^{net} is bounded by d where

$$d = \frac{H(H+3)}{2\theta(C - r_c - \delta)} \log \left(\frac{K}{1-z} \right). \quad (5.18)$$

It remains to choose δ . Since the optimal value of δ from Eq. (5.17) depends on d , the value of d in Eq. (5.18) can be determined for numerical purposes using an iterative method. We first let $\delta = \frac{C-(r+r_c)}{2}$ in Eq. (5.18) yielding a value d_0 . We then let $d = d_0$ in Eq. (5.17), and the resulting value of δ is finally used in Eq. (5.18).

The order of growth of $w^{net}(z)$ is now obtained by analyzing the constant K from Eq. (5.16). We can write

$$\begin{aligned} \log(K) &= \mathcal{O}(\log(H)) + \mathcal{O}(H) - \frac{1}{H(H+3)} \sum_{h=1}^H (h+1) \log(h+1) \\ &= \mathcal{O}(H) . \end{aligned}$$

Using the quadratic term from Eq. (5.18) we finally obtain that

$$w^{net}(z) = \mathcal{O}(H^3) .$$

We have thus shown that the method of using a statistical network service curve yields much tighter end-to-end delay bounds than the method of adding per-node results.

5.2 Network with Packetized Service

Rather than considering a fluid service model for packet networks, as in the previous section, we now consider a more realistic *packetized service model*. In this model packets become available for service at a node no sooner than fully processed at the next node upstream.

We first show how the packetized service model can be expressed in the terms of network calculus. Then we show that end-to-end delays in the network from Figure 2.1 grow as $\mathcal{O}(H \log H)$ when accounting for packetization. In other words, packetization does not change the $\mathcal{O}(H \log H)$ established in the previous section for networks with fluid service. Finally we establish the $\Omega(H \log H)$ lower bound on delays in a tandem network (no cross traffic), thus proving the $\Theta(H \log H)$ scaling behavior of delays.

Let us consider the network with cross traffic from Figure 2.1 where the capacity of each node is C . The arrivals of the through flow $A^1(t)$ are represented by a compound Poisson process, where packets arrive with rate λ and have exponentially distributed sizes X_i with mean $1/\mu$. Formally, we write

$$A^1(t) = \sum_{i=1}^{N(t)} X_i, \quad (5.19)$$

where $N(t)$ is a Poisson process with mean λt . Similarly, the arrivals of the cross flows $A_h(t)$, for $h = 1, \dots, H$, are represented by compound Poisson processes with arrival rates λ_c , and exponentially distributed packet sizes with mean $1/\mu$. Scheduling at the nodes is work-conserving and locally FIFO, such that our derivations hold even for the worst-case scenario when the packets of cross flows have preemptive priorities over the packets of the through flow. We denote the utilization factor by $\rho = (\lambda + \lambda_c)/(\mu C)$, and assume for stability that $\rho < 1$.

Consider the expression for the effective bandwidth of the through flow, i.e., $\alpha(\theta, t) = \frac{\lambda}{\mu - \theta}$ for some $0 < \theta < \mu$, and all $t \geq 0$ [52]. Then, following the construction of statistical envelope from effective bandwidth presented in [76], the through flow is bounded by an EBB statistical envelope

$$\mathcal{G}(t) = \frac{\lambda}{\mu - \theta} t \quad (5.20)$$

with error function $\varepsilon(\sigma) = e^{-\theta\sigma}$. This is obtained using the Chernoff bound as shown below for any $0 \leq s \leq t$, $\sigma \geq 0$, and some choice of $\theta > 0$

$$\begin{aligned} Pr \left(A(t) - A(s) > \frac{\lambda}{\mu - \theta} (t - s) + \sigma \right) &\leq E \left[e^{\theta(A(t) - A(s))} \right] e^{-\theta \frac{\lambda}{\mu - \theta} (t - s)} e^{-\theta\sigma} \\ &= e^{-\theta\sigma}. \end{aligned}$$

Similarly, at each node h , the cross flow $A_h(t)$ is bounded by an EBB statistical envelope

$$\mathcal{G}_h(t) = \frac{\lambda_c}{\mu - \theta_c} t \quad (5.21)$$

with error function $\varepsilon_h(\sigma) = e^{-\theta_c \sigma}$, for some $0 < \theta_c < \mu$.

To account for the effects of a packetized service model, we represent each node as a concatenation between a fluid server with rate C and a packetizer. The fluid server serves packets according to the fluid service model used in the previous section. The packetizer, denoted here by P^μ , has the role of a delay element that ensures that packets become available for service downstream after they have been fully processed upstream. We mention that packetizers have also been used in the context of the deterministic network calculus (see Parekh and Gallager [91], Le Boudec [15]). The packetizers introduced herein delay packets whose sizes are described by a distribution function, whereas the packetizers introduced in the deterministic network calculus delay packets whose sizes are always bounded.

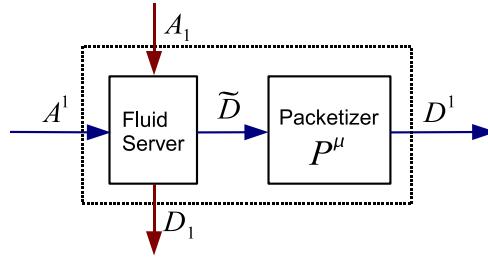


Figure 5.2: A statistical packetizer P^μ at the first node in a network with cross traffic.

In Figure 5.2 we illustrate the first node in the network from Figure 2.1. On one hand the fluid server is represented by a deterministic constant-rate service curve $\mathcal{R}(t) = Ct$, i.e., the output process $\tilde{D}(t)$ satisfies for all $t \geq 0$

$$\tilde{D}(t) = A^1 * \mathcal{R}(t). \quad (5.22)$$

Using the expression of the arrivals $A^1(t)$ from Eq. (5.19), we can also write

$$\tilde{D}(t) = \sum_{i=1}^{M(t)} X_i + X_f(t),$$

where $M(t)$ denotes the number of packets fully processed by time t , and $X_f(t)$ denotes the processed fraction of the packet (if any) currently in service at time t ; if the server is idle at

time t , then $X_f(t) = 0$. The process $\tilde{D}(t)$ is thus a virtual output process which represents the *fluid output* of the through flow at the fluid server.

On the other hand the packetizer P^μ takes the fluid output $\tilde{D}(t)$ as input, and produces the *packetized output* process $D^1(t)$ satisfying

$$D^1(t) = \sum_{i=1}^{M(t)} X_i .$$

This accounts for the fact that the second node can start processing a packet no sooner than the packet was completely processed by the fluid server at the first node. It then follows inductively that packetizers account for packetization in the entire network. Also, note that there is no loss in generality in assuming that the cross flows are not required to pass through packetizers (see Le Boudec [15]).

The next lemma shows that packetizers can be represented with statistical service curves.

Lemma 5.5 *Consider the network node represented in Figure 5.2, where packets arrive according to an exponential distribution with mean $1/\mu$. Then for any discretization parameter $\tau_0 > 0$ the function*

$$\mathcal{S}^\mu(t) = Ct$$

is a statistical envelope for the packetizer P^μ with error function

$$\varepsilon^\mu(\sigma) = e^{\mu C \tau_0} e^{-\mu \sigma} .$$

PROOF. Fix $\tau_0 > 0$ and $t \geq 0$, and let us denote \underline{t} as the beginning of the last busy period containing t at the fluid server. If $X_f(t)$ denotes the fraction already processed of the packet currently in service at the fluid server at time t , then

$$u \triangleq t - \frac{X_f(t)}{C} \tag{5.23}$$

is the time when the processing of the packet was started. It then follows that

$$D^1(t) - D^1(\underline{t}) = C(u - \underline{t}) . \quad (5.24)$$

We can write

$$\begin{aligned} Pr (D^1(t) < A^1 * [\mathcal{S}^\mu - \sigma]_+(t)) &\leq Pr (D^1(t) < D^1(\underline{t}) + [C(t - \underline{t}) - \sigma]_+) \\ &\leq Pr (C(t - u) > \sigma) \\ &\leq Pr (X_f(t) > \sigma) \\ &\leq e^{-\mu\sigma} . \end{aligned}$$

In the first line we used that \underline{t} is the beginning of a busy period. In the second line we applied Eq. (5.24). In the third line we applied the definition of u from Eq. (5.23), and in the last line we invoked the exponential distribution of packets' sizes.

We can now apply Lemma 4.4 and obtain that for all $t \geq 0$ and σ

$$Pr (D^1(t) < A^1 * [\mathcal{S}^\mu - \sigma]_+(t + \tau_0)) \leq e^{\mu C \tau_0} e^{-\mu\sigma} ,$$

which completes the proof. □

5.2.1 Upper bound

Here we compute end-to-end delay bounds in the network from Figure 2.1, and analyze their scaling properties. First we consider the simple case of a *tandem network*, i.e., there is no cross traffic. Then we consider the case of a network with positive cross traffic.

Using the deterministic service curve description of the fluid server from Eq. (5.22), and the statistical service curve description of the packetizer from Lemma 5.5, we represent the service offered by each node in the network from Figure 2.1 to the through flow by the EBB

statistical service curve

$$\begin{aligned}\mathcal{S}(t) &= \mathcal{R} * \mathcal{S}^\mu(t) \\ &= Ct,\end{aligned}$$

with error function

$$\varepsilon^s(\sigma) = e^{\mu C \tau_0} e^{-\mu \sigma}.$$

Having also the EBB statistical envelope for the through flow from Eq. (5.20), we next derive bounds on the end-to-end delay $W^{net}(t)$. We invoke Theorem 5.1 with the parameters

$$\begin{aligned}r_a &= \frac{\lambda}{\mu - \theta}, \quad M_a = 1, \quad \theta_a = \theta, \quad \text{and} \\ r_s &= C, \quad M_s = e^{\mu C \tau_0}, \quad \theta_s = \mu.\end{aligned}$$

Plugging these values into Eqs. (5.7), and choosing δ that satisfies the condition

$$0 < \delta \leq \frac{C - \frac{\lambda}{\mu - \theta}}{H},$$

we obtain the delay bound for $d \geq 0$

$$Pr(W^{net} > d) \leq \frac{\alpha}{\theta} (e^{\mu C \tau_0})^{\frac{H\theta}{\alpha}} \left(\frac{1}{\mu \delta \tau_0} \right)^{\frac{\alpha - \theta}{\alpha}} e^{-\frac{\theta \mu}{\alpha} (C - (H-1)\delta)d},$$

where

$$\alpha = H\theta + \mu.$$

Using Lemma 5.4 we minimize the violation probability with $\tau_0 = \frac{\alpha - \theta}{H\theta\mu C}$, yielding the delay bound

$$Pr(W^{net} > d) \leq K e^{-\frac{\theta \mu}{\alpha} (C - (H-1)\delta)d}, \quad (5.25)$$

where

$$K = \frac{\alpha}{\theta} \left(\frac{H\theta C e}{\delta(\alpha - \theta)} \right)^{\frac{\alpha - \theta}{\alpha}}.$$

Furthermore, setting the right-hand side of Eq. (5.25) equal to $1 - z$ for some $0 < z < 1$, the z -quantile $w^{net}(z)$ of the distribution of W^{net} is bounded by

$$d = \frac{\alpha}{\theta\mu(C - (H - 1)\delta)} \log \left(\frac{K}{1 - z} \right). \quad (5.26)$$

According to Lemma 5.4, the optimal value of δ is given by

$$\delta = \inf \left\{ \frac{\alpha - \theta}{(H - 1)\theta\mu d}, \frac{C - \frac{\lambda}{\mu - \theta}}{H} \right\}.$$

Since δ depends on d , the value of d in Eq. (5.26) can be determined for numerical purposes using the iterative method presented after Eq. (5.13).

The order of growth of $w^{net}(z)$ is given by Theorem 5.1, i.e.,

$$w^{net}(z) = \mathcal{O}(H \log H).$$

For the rest of the section we consider the case of the network from Figure 2.1 with positive cross traffic.

Having the statistical envelopes for the cross traffic from Eq. (5.21), we first apply Theorem 4.5 to derive a statistical leftover service curve for the through flow at each fluid server h for $h = 1, \dots, H$. The service curves are given for any $\delta > 0$ by the functions

$$T^h(t) = \left(C - \frac{\lambda_c}{\mu - \theta_c} - \delta \right) t, \quad (5.27)$$

with error functions

$$\varepsilon^{T,h}(\sigma) = \frac{e^{\theta_c C \tau_0}}{\theta_c \delta \tau_0} e^{-\theta_c \sigma}. \quad (5.28)$$

Next, having the description of the packetizers from Lemma 5.5, we invoke Theorem 4.7

to derive a statistical *network* service curve for the concatenation between a fluid server and the corresponding packetizer. The service curve is given at each node h by the function

$$\begin{aligned} \mathcal{S}^h(t) &= T^h * \mathcal{S}_{-\delta}^\mu(t) \\ &= \inf_{0 \leq s \leq t} \left\{ \left(C - \frac{\lambda_c}{\mu - \theta_c} - \delta \right) s + (C - \delta)(t - s) \right\} \\ &= \left(C - \frac{\lambda_c}{\mu - \theta_c} - \delta \right) t. \end{aligned}$$

The corresponding error function is

$$\begin{aligned} \varepsilon^h(\sigma) &= \tilde{\varepsilon}_{\delta\tau_0}^{T,h} * \varepsilon^\mu(\sigma) \\ &= \inf_{\sigma_1 + \sigma_2 = \sigma} \left\{ e^{\theta_c C \tau_0} \left(\frac{1}{\theta_c \delta \tau_0} \right)^2 e^{-\theta_c \sigma_1} + e^{\mu C \tau_0} e^{-\mu \sigma_2} \right\} \\ &= \frac{\theta_c + \mu}{\theta_c} \left(\left(\frac{e^{\theta_c C \tau_0}}{\theta_c \delta \tau_0} \right)^2 \frac{\theta_c}{\mu} \right)^{\frac{\mu}{\theta_c + \mu}} e^{-\frac{\theta_c \mu}{\theta_c + \mu} \sigma}. \end{aligned}$$

In the last line we applied Lemma 5.3.

Having the EBB statistical envelope for the through flow from Eq. (5.20), we now derive bounds on the end-to-end delay W^{net} by invoking Theorem 5.1 with the parameters

$$\begin{aligned} r_a &= \frac{\lambda}{\mu - \theta}, \quad M_a = 1, \quad \theta_a = \theta, \quad \text{and} \\ r_s &= C - \frac{\lambda_c}{\mu - \theta_c} - \delta, \quad M_s = \frac{\theta_c + \mu}{\theta_c} \left(\left(\frac{e^{\theta_c C \tau_0}}{\theta_c \delta \tau_0} \right)^2 \frac{\theta_c}{\mu} \right)^{\frac{\mu}{\theta_c + \mu}}, \quad \theta_s = \frac{\theta_c \mu}{\theta_c + \mu}. \end{aligned}$$

Plugging these values into Eq. (5.7) after first enforcing the condition

$$\delta \leq \frac{C - \frac{\lambda}{\mu - \theta} - \frac{\lambda_c}{\mu - \theta_c}}{H + 1},$$

we obtain the delay bound

$$Pr(W^{net} > d) \leq K e^{-\frac{\theta \theta_c \mu}{\gamma} (C - \frac{\lambda_c}{\mu - \theta_c} - H \delta) d}, \quad (5.29)$$

where

$$\begin{aligned} K &= \frac{\gamma}{\theta} \left(\frac{2HCe\theta\mu}{\delta\beta} \right)^{\frac{\beta}{\gamma}} (\theta_c + \mu)^{\frac{(H-1)\theta(\theta_c+\mu)}{\gamma}} \left(\frac{1}{\theta_c} \right)^{\frac{H\theta\theta_c}{\gamma}} \left(\frac{1}{\mu} \right)^{\frac{\beta-H\theta\mu}{\gamma}}, \\ \gamma &= H\theta(\theta_c + \mu) + \theta_c\mu, \text{ and} \\ \beta &= (3H - 1)\theta\mu + (H - 1)\theta\theta_c + \theta_c\mu. \end{aligned}$$

In the derivation we applied Lemma 5.4 that gives $\tau_0 = \frac{\beta}{2HC\theta\theta_c\mu}$.

Setting the right-hand side of Eq. (5.29) equal to $1 - z$ for some $0 < z < 1$, the z -quantile $w^{net}(z)$ of the distribution of W^{net} is bounded by

$$d = \frac{\gamma}{\theta\theta_c\mu(C - \frac{\lambda_c}{\mu-\theta_c} - H\delta)} \log \left(\frac{K}{1-z} \right). \quad (5.30)$$

According to Lemma 5.4, the optimal value of δ is given by

$$\delta = \inf \left\{ \frac{\beta}{H\theta\theta_c\mu d}, \frac{C - \frac{\lambda}{\mu-\theta} - \frac{\lambda_c}{\mu-\theta_c}}{H+1} \right\}.$$

Since δ depends on d , the value of d in Eq. (5.30) can be computed for numerical purposes with the iterative method presented after Eq. (5.13).

The order of growth of $w^{net}(z)$ is given by Theorem 5.1, i.e.,

$$w^{net}(z) = \mathcal{O}(H \log H).$$

5.2.2 Lower bound

Here we show that end-to-end delays of packets in a tandem network are bounded from below by $\Omega(H \log H)$.

First, let us make some important remarks about the differences between the delays of packets, and the virtual delays from Eq. (2.1) which we usually compute with the network calculus. Consider a single node where packets arrive as a Poisson process with rate λ and

have exponentially distributed service times with mean $1/(\mu C)$. Denote t_n as the departure time of packet n , and let $t_0 = 0$. Also, denote Y_n as the service time of packet n , and W'_n as the delay of packet n . Let $W(t)$ the virtual delay from Eq. (2.1) under either a fluid or packetized service model. Assuming a stability condition, let W' as the steady-state delay of packets, and W as the steady-state delay of $W(t)$.

We observe that while the packets' delays W'_n are always positive, the calculus delays satisfy $W(t) = 0$ whenever t does not belong to a busy period. In other words, $Pr(W(t) = 0) = 1 - \rho$, where $\rho = \lambda/(\mu C)$. Also, we have that for all $t \in [t_{n-1}, t_n)$

$$W(t) \leq W'_n \leq W(t) + Y_n ,$$

which yields that $E[W] \leq E[W']$, and $E[W'] - E[W] \leq 1/(\mu C)$. This means that for numerical purposes, the difference between the two delays is negligible. In a multi-node scenario, one can similarly show an $\mathcal{O}(H)$ difference between the end-to-end delays of packets and the virtual end-to-end delay computed with the calculus.

Now we give the main result of this section.

Theorem 5.6 ($\Omega(H \log H)$ END-TO-END DELAY BOUNDS) *Consider a flow traversing a tandem network with H nodes, each having capacity C . The arrivals of the flow consist of packets arriving with rate λ and having exponentially distributed packet sizes with mean $1/\mu$. All inter-arrival times and packet sizes are stationary and independent. Denote by W^{net} the steady-state end-to-end delay of the packets. Then for each $0 < z < 1$ there exist two constants γ_1 and γ_2 such that*

$$Pr\left(W^{net} < \gamma_1 H \log(\gamma_2 H)\right) \leq z . \quad (5.31)$$

From the definition of the z -quantile $w^{net}(z)$ (see Eq. (5.1)) we have that

$$w^{net}(z) \geq \gamma_1 H \log(\gamma_2 H) ,$$

implying the following result.

Corollary 5.7 *The z -quantile $w^{net}(z)$ of the distribution of the steady-state end-to-end delay of packets W^{net} satisfies*

$$w^{net}(z) = \Omega(H \log H) .$$

The lower bound remains valid in a network with cross traffic because the effect of adding cross traffic in the tandem network is an increase of the end-to-end delays. In conjunction with Corollary 5.2, and the observations we made at the beginning of this section about the differences between packets' and virtual delays, we obtain the scaling behavior of delays in networks with EBB envelopes and EBB service curves.

Corollary 5.8 *The z -quantile $w^{net}(z)$ of the distribution of the steady-state end-to-end delays of packets W^{net} satisfies*

$$w^{net}(z) = \Theta(H \log H) .$$

A critical assumption in the theorem leading to the $\Omega(H \log H)$ lower bounds for the end-to-end delays is that packets maintain their sizes at each of the traversed nodes. Analyzing such networks is difficult due to the correlations which arise between service at different nodes. To see where these correlations stem from, suppose that a very large packet arrives at the first node; then, the next packet is likely to experience very large queueing delays at each node along the network path, due to the high processing times of the large packet at the nodes.

Before proving the theorem let us review some related literature concerning the analysis of tandem networks with identical service times. A simple approach to analyze a packet-switched network, where the sizes of packets are implicitly identical, relies on Kleinrock's independence assumption [67]. The assumption states that the size of each packet is *independently* re-sampled at each of the traversed node. With this assumption, the network has a product-form (under proper assumptions on arrivals and service time distributions), and exact

results are readily available. Although not realistic, the independence assumption can lead to reasonably good approximations in certain scenarios (e.g. densely-connected networks or high utilizations) [10].

One of the first attempts for the analysis of a tandem network with identical service times was pursued by Boxma [17]. For a network with two nodes, Poisson arrivals and general service times distribution, Boxma derives the steady-state delay distribution of the delay at the second node. Also, he shows that the (positive) correlations of the delays at the two nodes are much higher than in the case of independent service times of packets. In a G/G/1 tandem network characterized by general arrivals and service time distributions, Calo shows in [24] that the delays of any packet at the network nodes, starting with the second, are nondecreasing. More precisely, Vinogradov shows that the delays grow as $\Theta(\log h)$, where h is the node, for Poisson input and exponentially distributed service times [109].

In [108] Vinogradov provides the expression for the steady-state end-to-end delay distribution, starting with the second node, for Poisson arrivals and general service times. This result is later extended to general arrivals by Le Gall [50]. Consider the case of a tandem network with $H + 1$ nodes, packets arriving with rate $\lambda < 1$, and exponentially distributed service times with mean one. If W^h denote the steady-state delay of packets at nodes $h = 1, \dots, H + 1$, then we have according to [108] for all $d \geq 0$

$$Pr(W^2 + \dots + W^{H+1} \leq d) = \frac{(1 - \lambda) \left(1 - e^{-\frac{d}{H}}\right) x\left(\frac{d}{H}\right)}{\lambda e^{-\frac{d}{H}}} e^{-H \int_{\frac{d}{H}}^{\infty} x(t) dt}, \quad (5.32)$$

where $x(u)$ is the solution in x of the equation

$$x = \lambda \left(1 - \frac{1}{x+1} e^{-(x+1)u}\right). \quad (5.33)$$

Vinogradov also shows in [109] that the expectation of the end-to-end delay for Poisson arrivals and exponentially distributed service times grows as $\Theta(H \log H)$ for $H \rightarrow \infty$; note that Corollary 5.8 reproduces this result.

Let us remark that the end-to-end delay distribution from Eq. (5.32) is not explicit, and numerical results are difficult to compute. Concretely, one needs to solve the transcendental equation (5.33), and further compute the integral of the solution in Eq. (5.32). In light of this apparent limitation of Eq. (5.32), the contribution of Theorem 5.6 is that it provides a simple expression for the end-to-end delay lower bound, that quickly lends to numerical results. Also, unlike the end-to-end delay expression from Eq. (5.32), Theorem 5.6 accounts for the first node in the tandem network.

Now we turn to the proof of Theorem 5.6, which will use a maximal inequality for supermartingales due to Doob. In order to state this inequality, let us first briefly introduce the concept of a supermartingale (we follow Grimmett and Stirzaker [52]).

Definition 5.9 (SUPERMARTINGALE) *Let a continuous-time stochastic process $\mathbf{X} = \{X(t) : t \geq 0\}$ such that $X(t)$ is integrable for all t , i.e., $E[|X(t)|] < \infty$. Let also a family $\mathbf{F} = \{\mathcal{F}_t : t \geq 0\}$ of sub- σ -algebras of \mathcal{F} satisfying two properties: (1) $\mathcal{F}_s \subseteq \mathcal{F}_t$ for all $s \leq t$, and (2) $X(t)$ is \mathcal{F}_t -measurable for all $t \geq 0$. (\mathbf{X}, \mathbf{F}) is a supermartingale if the inequality*

$$E[X(t) \mid \mathcal{F}_s] \leq X(s) \text{ a.s.},$$

holds for all $0 \leq s \leq t$.

The family of σ -algebras \mathcal{F} is appropriately called a *filtration* to reflect the inclusion property (1) from the definition. We usually work with σ -algebras \mathcal{F}_t of the form

$$\mathcal{F}_t = \sigma \{X(s) : s \leq t\},$$

i.e., the σ -algebras generated by the past history of the process. Automatically, the family of σ -algebras constructed in this way satisfies the property (2) in the definition. To simplify notation, we usually say that $X(t)$ is a supermartingale when the corresponding filtration is implied.

The quantity $E[X(t) \parallel \mathcal{F}_s]$ is the expectation of $X(t)$ conditioned on the σ -algebra \mathcal{F}_s . This type of conditional expectation is a random variable itself and satisfies the following useful properties.

Lemma 5.10 (BASIC PROPERTIES OF CONDITIONAL EXPECTATION) *Let X, Y two random variables, and let \mathcal{G} a σ -algebra satisfying $\mathcal{G} \subseteq \mathcal{F}$. If X is statistically independent of \mathcal{G} then*

$$E[X \parallel \mathcal{G}] = E[X] \text{ a.s..}$$

If X is \mathcal{G} -measurable then

$$E[XY \parallel \mathcal{G}] = X E[Y \parallel \mathcal{G}] \text{ a.s..}$$

Now we present Doob's maximal inequality.

Lemma 5.11 (DOOB'S MAXIMAL INEQUALITY - [52]) *Let (X, \mathbf{F}) a nonnegative supermartingale. Then we have*

$$Pr \left(\sup_{0 \leq s \leq t} X(s) > x \right) \leq \frac{E[X(0)]}{x}, \quad (5.34)$$

for all $t \geq 0$ and $x > 0$.

PROOF OF THEOREM 5.6. Fix $0 < z < 1$. For any $j \geq 1$ let the following variables.

- X_j : inter-arrival time between the $(j-1)^{th}$ and j^{th} packets
- Y_j : the service time of packet j
- $T_{H,j}$: the departure time of j^{th} packet from node H
- $W_{H,j}$: total delay experienced by the j^{th} packet across the network

For $n \geq 1$ we can bound $T_{H,n}$ as

$$T_{H,n} \geq \sup_{0 \leq j \leq n} \left\{ \sum_{i=0}^j X_i + HY_j + \sum_{i=j+1}^n Y_i \right\},$$

where we denoted for convenience $X_0 = 0$ and $Y_0 = 0$. The equation states that for any $j \leq n$ the following is true: $T_{H,n}$ is greater than the arrival time of the j^{th} packet at the first node, plus the minimum time required by the j^{th} packet to traverse the network, and plus the total service times of packets $j+1, \dots, n$ at node H .

Since the n^{th} packet arrives at the first node at time $\sum_{i=1}^n X_i$, it immediately follows that its end-to-end delay $W_{H,n}$ is bounded below by

$$W_{H,n} \geq \sup_{0 \leq j \leq n} \left\{ HY_j + \sum_{i=j+1}^n (Y_i - X_i) \right\}.$$

In terms of probabilities, we have that for all $d \geq 0$ and $K > 0$

$$\begin{aligned} & Pr(W_{H,n} \geq d) \\ & \geq Pr \left(\sup_{0 \leq j \leq n} \left\{ HY_j + \sum_{i=j+1}^n (Y_i - X_i) \right\} \geq d \right) \\ & \geq Pr \left(\sup_{0 \leq j \leq n} \left\{ HY_j - K(n-j) + \sum_{i=j+1}^n (Y_i - X_i) + K(n-j) \right\} \geq d \right) \\ & \geq Pr \left(\sup_{0 \leq j \leq n} \{ HY_j - K(n-j) \} \geq d \right) \\ & \quad - Pr \left(\inf_{0 \leq j \leq n} \left\{ \sum_{i=j+1}^n (Y_i - X_i) + K(n-j) \right\} < 0 \right). \end{aligned} \quad (5.35)$$

In the last line we used the implication

$$\left\{ \sup_j a_j \geq 0 \right\} \Rightarrow \left\{ \sup_j \{a_j + b_j\} \geq 0 \right\} \cup \left\{ \inf_j b_j < 0 \right\},$$

for some numbers a_j, b_j . Also, K is a free parameter which is to be chosen below such that

the the two probabilities in Eq. (5.35) are properly balanced.

We now evaluate the second term of the right-hand side of Eq. (5.35). Let $0 < \theta < \lambda$ and choose K such that

$$\frac{\lambda}{\lambda - \theta} \frac{\mu C}{\mu C + \theta} e^{-\theta K} = 1 - z. \quad (5.36)$$

Using the following notations for all $1 \leq j \leq n$

$$\begin{aligned} U_j &= -(Y_{n-j+1} - X_{n-j+1}), \\ S_j &= \sum_{i=1}^j U_i, \text{ and} \\ T_j &= e^{\theta(S_j - Kj)}, \end{aligned}$$

the last probability in Eq. (5.35) can be rewritten as

$$P \left(\sup_{1 \leq j \leq n} T_j > 1 \right). \quad (5.37)$$

This will be estimated next.

Let the filtration of σ -algebras $\mathcal{F}_j = \sigma(U_1, U_2, \dots, U_j)$. We have for any $j \geq 1$ that

$$\begin{aligned} E [T_{j+1} \mid \mathcal{F}_j] &= E [T_j e^{\theta(U_{j+1} - K)} \mid \mathcal{F}_j] \\ &= T_j E [e^{\theta(U_{j+1} - K)}] \\ &= T_j \frac{\lambda}{\lambda - \theta} \frac{\mu C}{\mu C + \theta} e^{-\theta K} \\ &\leq T_j. \end{aligned}$$

In the second and third lines we applied the statistical independence of X_i and Y_i (see the properties of conditional expectation from Lemma 5.10). In the third line we used the moment generating functions $E [e^{\theta X_j}] = \frac{\lambda}{\lambda - \theta}$ and $E [e^{-\theta Y_j}] = \frac{\mu C}{\mu C + \theta}$. In the last line we used the definition of K from Eq. (5.36).

Since $T_j \geq 0$ for all j it then follows that T_j is a (discrete-time) supermartingale. Invoking

Doob's maximal inequality for supermartingales (Lemma 5.11) we have that

$$\begin{aligned}
 P\left(\sup_{1 \leq j \leq n} T_j > 1\right) &\leq E[T_1] \\
 &\leq \frac{\lambda}{\lambda - \theta} \frac{\mu C}{\mu C + \theta} e^{-\theta K} \\
 &\leq 1 - z.
 \end{aligned} \tag{5.38}$$

We now turn to the first term in Eq. (5.35). Choosing a value d such that

$$P\left(\sup_{0 \leq j \leq n} (HY_j - K(n - j)) \geq d\right) = 2(1 - z),$$

Eqs. (5.35) and (5.38) yield

$$Pr(W_{H,n} \geq d) \geq 1 - z. \tag{5.39}$$

Moreover, Eq. (5.38) implies that

$$1 - 2(1 - z) = \prod_{j=0}^n \left(1 - e^{-\frac{\mu C}{H}(d + Kj)}\right),$$

such that we can write

$$\begin{aligned}
 \log(1 - 2(1 - z)) &= \sum_{j=0}^n \log\left(1 - e^{-\frac{\mu C}{H}(d + Kj)}\right) \leq -\sum_{j=0}^n e^{-\frac{\mu C}{H}(d + Kj)} \\
 &\leq -\frac{e^{-\frac{\mu C}{H}d} \left(1 - e^{-\frac{\mu C}{H}(n+1)K}\right)}{1 - e^{-\frac{\mu C}{H}K}} \\
 &\leq -\frac{H}{\mu CK} e^{-\frac{\mu C}{H}d}.
 \end{aligned} \tag{5.40}$$

In the second line we applied the inequality $\log(1 - x) \leq -x$ for all $0 \leq x < 1$. In the third line we evaluated the series, and in the last line we let $n \rightarrow \infty$ and applied the inequality $1 - e^{-x} \leq x$ for all $x \geq 0$.

From Eq. (5.40) we can choose

$$\begin{aligned}\gamma_1 &= \frac{1}{\mu C}, \text{ and} \\ \gamma_2 &= -\frac{1}{\mu C K \log(1 - 2(1 - z))}.\end{aligned}$$

Since $W^{net} = \lim_{n \rightarrow \infty} W_{H,n}$, it follows from Eq. (5.39) that the proof is complete. \square

5.3 Numerical Examples

In this section we present two sets of numerical examples. In the first we illustrate the delay bounds derived in Section 5.1 in the case of a network with cross traffic, Markov-modulated On-Off arrivals, and a fluid service model. In the second we compare the upper and lower bounds on end-to-end delays derived in Sections 5.2.1 and 5.2.2 for tandem networks and a packetized service model.

5.3.1 Markov-modulated On-Off processes

Here we present numerical examples which illustrate that the method of using a statistical network service curve yields much smaller end-to-end delay bounds than the method of adding per-node delay bounds. Delay bounds with the two methods have been computed in Section 5.1.

We assume that the through traffic and each cross traffic at nodes $h = 1, \dots, H$ consist of aggregates of statistically independent Markov-modulated On-Off processes, which can be defined as follows. Let a two-state homogenous and continuous-time Markov chain $X(t)$ with the transition matrix

$$Q = \begin{pmatrix} -\mu & \mu \\ \lambda & -\lambda \end{pmatrix}.$$

The two states of $X(t)$ are denoted ‘On’ and ‘Off’, and μ and λ represent the transition rates

from the ‘On’ state to the ‘Off’ state, and vice-versa, respectively. In the steady-state, the average spending time of the process $X(t)$ in the ‘On’ state is $\frac{1}{\mu}$, and the average spending time in the ‘Off’ state is $\frac{1}{\lambda}$ [101]. Therefore, the value

$$T = \frac{1}{\mu} + \frac{1}{\lambda} \quad (5.41)$$

is the average time for the Markov process $X(t)$ to change states twice. Also, T reflects the burstiness of the process with small values of T indicating a low burstiness and vice-versa.




Figure 5.3: A Markov-modulated On-Off traffic model.

A continuous-time arrival process $A(t)$ is a Markov-modulated On-Off process driven by a Markov process $X(t)$ if the arrival rate of the process is either P or zero, depending whether $X(t)$ is in the ‘On’ and ‘Off’ states, respectively. Figure 5.3 illustrates a Markov-modulated On-Off process.

In our network scenario, we assume that there are N through flows and N_c cross flows at each node. Consider the expression of the effective bandwidth for a single Markov-modulated On-Off process $A(t)$ (see Chang [26], Kelly [64])

$$\alpha_A(\theta, t) \leq \frac{1}{2\theta} \left(P\theta - \mu - \lambda + \sqrt{(P\theta - \mu + \lambda)^2 + 4\lambda\mu} \right). \quad (5.42)$$

Then, using the statistical independence of the flows and the Chernoff bound (see also Li *et al.* [76]), we obtain that the function

$$\mathcal{G}(t) = \frac{N}{2\theta} \left(P\theta - \mu - \lambda + \left((P\theta - \mu + \lambda)^2 + 4\lambda\mu \right)^{\frac{1}{2}} \right) t$$

is an EBB statistical envelope for the through traffic with error function $\varepsilon(\sigma) = e^{-\theta\sigma}$ for any

$\theta > 0$. Analogous expressions for the cross traffic envelopes are obtained with N replaced by N_c . These descriptions further permit the calculation of the delay bounds given in Eqs. (5.13) and (5.18) from Section 5.1.

Burstiness	T (ms)	P (Mbps)	r (Mbps)	μ (ms^{-1})	λ (ms^{-1})
low	10	1.5	0.15	1.0	0.11
high	100	1.5	0.15	0.1	0.01

Table 5.1: Parameters of a Markov-modulated On-Off Markov process.

Let us consider the following numerical settings. Time is measured in milliseconds and we plot bounds on the z -quantiles $w^{net}(z)$ with $z = 1 - 10^{-9}$. The capacity of each node is $C = 100$ Mbps. The peak rate of a flow is $P = 1.5$ Mbps and the average rate is $r = 0.15$ Mbps. We consider both the case of Markov processes with high burstiness ($T = 10$ ms) and small burstiness ($T = 100$ ms). Note that the values of P , r , and T determine the values for the transition rates μ and λ (see Table 5.1). Lastly, the only unknown variable, i.e., θ , is subject to numerical optimizations.

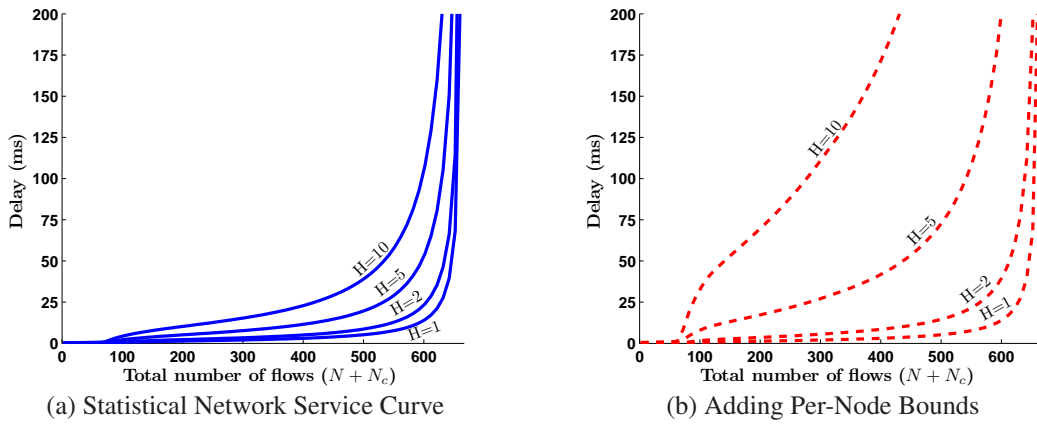


Figure 5.4: End-to-end delay $w^{net}(z)$ for Markov-modulated On-Off processes as a function of the number of flows $N + N_c$ ($H = 1, 2, 5, 10$, $T = 10$ ms (low burstiness), $N = N_c$, $z = 1 - 10^{-9}$)

In Figures 5.4.(a) and (b) we show the bounds on the z -quantiles $w^{net}(z)$ obtained with the method of using a statistical network service curve (Eq. (5.13)), and with the method of

adding per-node bounds (Eq. (5.18)), as a function of the total number of flows $N + N_c$. We consider several number of nodes ($H = 1, 2, 5$, and 10), equal share of through and cross flows at each node ($N = N_c$), and low burstiness for a flow ($T = 10$ ms). We observe that as the number of nodes increases, the differences between the bounds obtained with the two methods become more accentuate. Also, the bounds obtained with Eq. (5.18) are much more sensitive than the bounds from Eq. (5.13) when the number of nodes H increases.

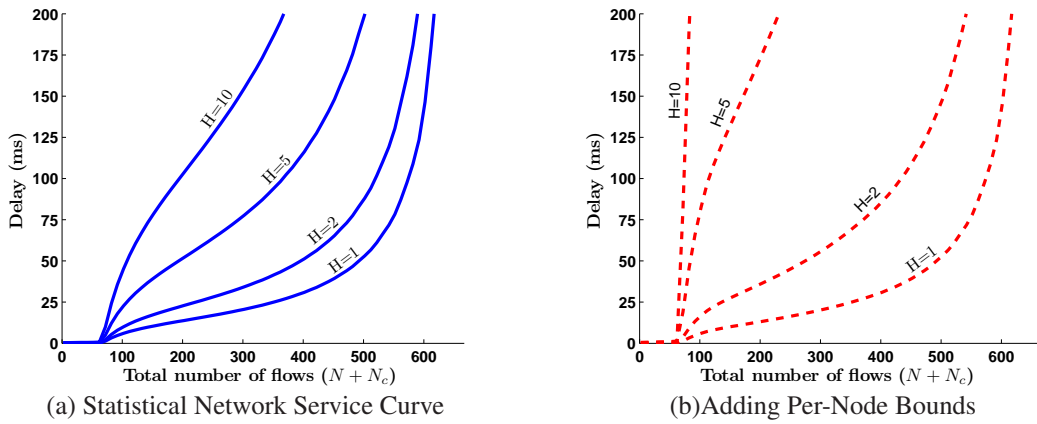


Figure 5.5: End-to-end delay $w^{net}(z)$ for Markov-modulated On-Off processes as a function of the number of flows $N + N_c$ ($H = 1, 2, 5, 10$, $T = 100$ ms (high burstiness), $N = N_c$, $z = 1 - 10^{-9}$)

In Figures 5.5.(a) and (b) we consider similar numerical settings as in Figures 5.4.(a) and (b), but for flows with higher burstiness ($T = 100$ ms). Because the flows are more bursty, the resulted delay bounds are more pessimistic than the corresponding bounds in Figures 5.4.(a) and (b).

In Figure 5.6 we show the bounds on the z -quantiles $w^{net}(z)$ obtained with the two methods as a function of the total number of flows $N + N_c$. Different from Figure 5.4 is that we now consider different percentages p of the through flows out of the total number of flows ($p = 10\%$ and 90%). Figure 5.6.(a) shows that for a small number of nodes ($H = 2$) the differences between the bounds obtained with the two methods are relatively small depending on the load of through traffic. However, as Figure 5.6.(b) shows, the differences between the bounds become more pronounced when the number of nodes increases ($H = 10$).

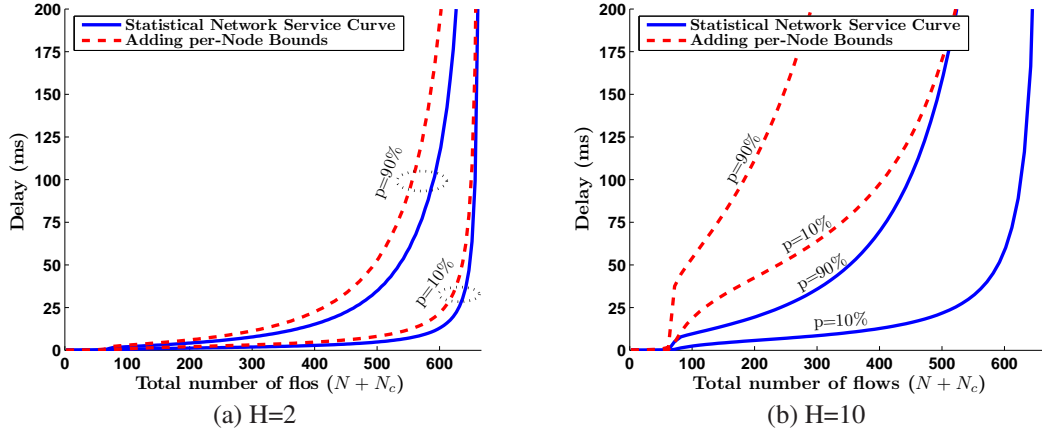


Figure 5.6: End-to-end delay $w^{net}(z)$ as a function of the number of flows $N + N_c$ ($N = 10\%$, 90% out of the total number of flows $N + N_c$, $H = 2, 10$, $T = 10\text{ms}$, $z = 1 - 10^{-9}$)

In Figure 5.7 we show the bounds on the z -quantiles $w^{net}(z)$ obtained with the two methods as a function of the total number of nodes H . We now consider two utilization factors ($\rho = 0.5$ and $\rho = 0.9$) and set $T = 10$ ms, and $N = N_c$. The figure best illustrates the $\mathcal{O}(H \log H)$ and $\mathcal{O}(H^3)$ scaling behaviors of the end-to-end delay bounds obtained with the two methods.




Figure 5.7: End-to-end delay $w^{net}(z)$ as a function of the number of nodes H in the network (utilization factor $\rho = 0.5, 0.9$, $T = 10$ ms, $N = N_c$, $z = 1 - 10^{-9}$)

5.3.2 Comparison between upper and lower bound

Here we give numerical examples to illustrate a comparison between the upper bound from Eq. (5.26), and the corresponding lower bound from Theorem 5.6. We recall that both bounds were derived in tandem networks with Poisson arrivals, exponentially distributed packet sizes, identical service times at the nodes, and a packetized service model.

We also plot the exact results obtained by using Kleinrock's independence assumption stating that the size of each packet is independently regenerated at each traversed node [67]. Additionally, we show simulation results on $w^{net}(z)$ obtained as follows. We start with an empty network and record the largest 100 delays among the first 10^8 packets that complete service. The smallest of the recorded delays is the simulated value of $w^{net}(z)$.

We consider the following numerical settings. Time is measured in milliseconds and we plot bounds on the z -quantiles $w^{net}(z)$ with $z = 1 - 10^{-6}$. The capacity of each node is $C = 100$ Mbps and the average size of packets is $1/\mu = 400$ Bytes [87]. We note that for an utilization factor ρ , the arrival rate of packets is $\lambda = \rho\mu C$. Lastly, the unknown variable θ (in the expression for the upper bound from Eq. (5.26)) is numerically optimized.

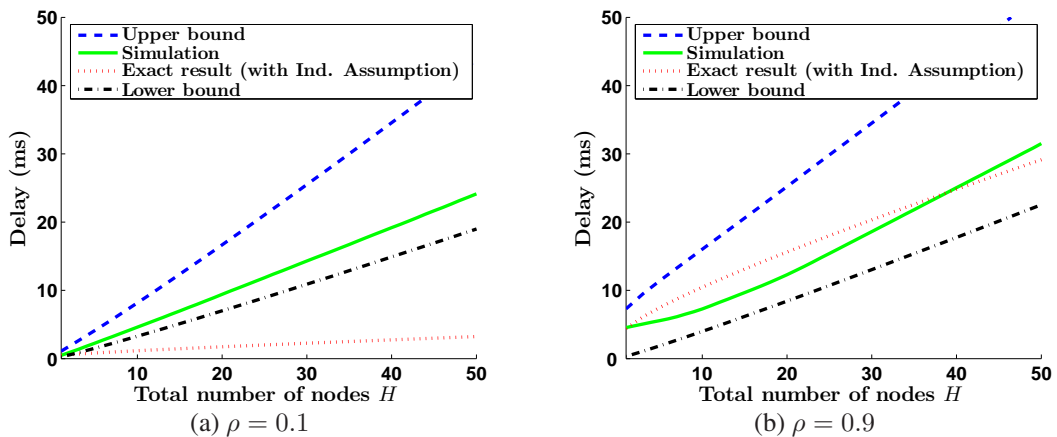


Figure 5.8: End-to-end delay $w^{net}(z)$ as a function of the number of nodes H (utilization factor ($\rho = 0.1$ and $\rho = 0.9$), $C = 100$ Mbps, average packet size $\mu^{-1} = 400$ Bytes, $z = 1 - 10^{-6}$)

In Figure 5.8 we show the end-to-end delay bounds as a function of the total number of

nodes H in the network. We consider both low utilization factor ($\rho = 0.1$) in Figure 5.8.(a), and high utilization factor ($\rho = 0.9$) in Figure 5.8.(b). The figures show that the utilization factor has little impact on simulation results, a fact that is captured by both the upper and lower bounds. The formula for the exact result is too optimistic at low utilizations, and becomes more accurate at high utilizations. The last observation has also been pointed out in the literature (see [10]), in order to explain in which scenarios Kleinrock's independence assumption is justified.

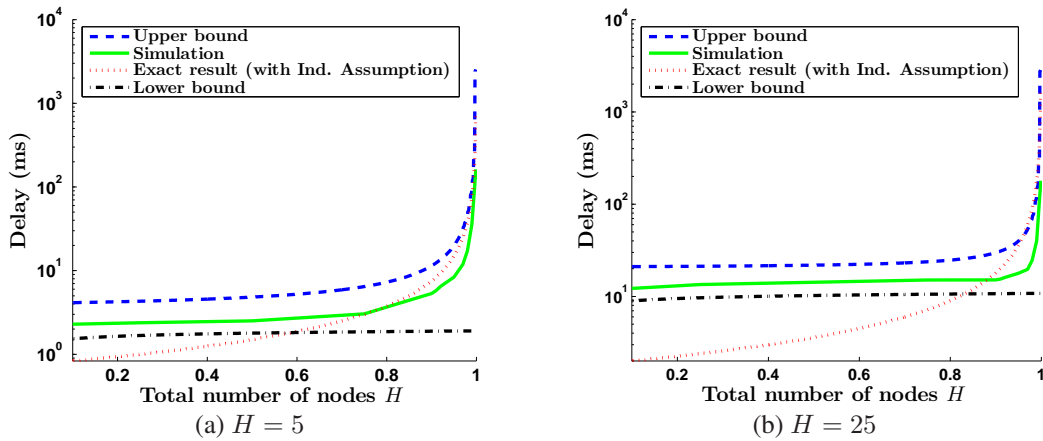


Figure 5.9: End-to-end delay $w^{net}(z)$ as a function of the number of the utilization factor ρ (number of nodes ($H = 5$ and $H = 25$), $C = 100$ Mbps, average packet size $\mu^{-1} = 400$ Bytes, $z = 1 - 10^{-6}$)

In Figure 5.9 we show (on a logarithmic scale) the end-to-end delay bounds by varying the utilization factor ρ (the last shown value is $\rho = 0.999$). We consider a small number of nodes ($H = 5$) in Figure 5.9.(a), and a high number of nodes ($H = 25$) in Figure 5.9.(b). The figures shows that both the lower and upper bounds are quite accurate at utilization factors less than 0.8. At high utilizations the lower bound does not capture the exponential increase of the delay, whereas the upper bound is too pessimistic.

For small number of nodes, the formula for the exact results predicts overly optimistic delays at low utilizations, becomes accurate at high utilization, but predicts overly pessimistic delays at very high utilizations (e.g. $\rho = 0.999$). For high number of nodes, the formula for

the exact results is inaccurate at most utilizations; the formula even predicts more pessimistic delays than the upper bounds around an utilization factor of $\rho = 0.97$.

Chapter 6

Accounting for Statistical Independence

The network calculus formulation presented in Chapter 4 is generally suitable to analyze network scenarios where arrivals and service at the nodes may be statistically correlated. Consequently, the calculus may yield pessimistic end-to-end performance bounds in scenarios with independent arrivals or service. To account for statistical independence across multiple nodes, and yet capture possible correlations among arrivals or service, we now extend the network calculus formulation with moment generating functions proposed by Chang [29] and Fidler [48].

The key extension of the proposed network calculus formulation is a new statistical service curve model. Similar to the service models used in [29, 48], the proposed service model can account for statistical independence of service across multiple nodes. Unlike the service models from [29, 48], the proposed service model permits the analysis of network scenarios where service at the nodes may be statistically correlated.

One scenario where the proposed network calculus formulation is particularly useful is a network with statistically independent cross traffic and a packetized service model. If packets maintain their sizes along the traversed nodes, then the service at the nodes is statistically correlated. These correlations can be captured within the proposed service curve model, and yet the statistical independence of the cross traffic can be accounted for. Another scenario is a network with correlated cross traffic, packetized service model, but where the size of each

packet is independently re-sampled at each traversed node. Although both arrivals and service are clearly correlated, the proposed network calculus can exploit the statistical independence of service arising from re-sampling the packets sizes. In the next chapter we will derive explicit delay bounds for these two network scenarios.

As in [29, 48], the methods in this chapter are restricted to the class of arrival processes having bounded moment generating functions. In particular, the class of heavy-tailed arrival processes is not covered. Besides the general case of arrivals with bounded MGFs, we also consider the special case of stationary arrivals with independent increments, for which improved performance bounds can be derived.

The rest of this chapter is organized as follows. In Sections 6.1-6.4 we consider independent arrivals or service: we give the arrivals and service descriptions, and then present results on single-node performance bounds and the construction of statistical network service curves. In Section 6.5 we specialize to the case of arrival processes with stationary and independent increments.

6.1 Statistical Envelope

As in the network calculus formulation proposed by Chang [29] and Fidler [48], we characterize the arrivals by bounds on their moment generating functions (MGF). Unlike in [29, 48], we now work in a continuous time setting.

Definition 6.1 (MGF ENVELOPE FOR ARRIVALS) *An arrival process $A(t)$ is bounded by an MGF envelope, with rate r and prefactor M , for some choice of $\theta > 0$, if for all $0 \leq s \leq t$*

$$E [e^{\theta A(s,t)}] \leq M e^{\theta r(t-s)}. \quad (6.1)$$

In general, both the rate r and the prefactor M depend on θ .

If the arrival process is stationary, the MGF envelope model is equivalent to the linear envelope model from Eq. (3.4); note that the rate r in Eq. (6.1) equals to the effective bandwidth $\alpha_A(\theta, t)$ in Eq. (3.4).

Next we relate the MGF envelope model from Definition 6.1 with the statistical envelope model from Definition 4.1. This result will later enable relating results obtained in this chapter with results obtained using the network calculus formulations from Chapter 4.

Lemma 6.2 *Suppose that an arrival process $A(t)$ is bounded by an MGF envelope with rate r and prefactor M for some choice of $\theta > 0$. Then the function*

$$\mathcal{G}(t) = rt \tag{6.2}$$

is a statistical envelope for $A(t)$ with error function $\varepsilon(\sigma) = Me^{-\theta\sigma}$, in the sense of Definition 4.1 ($\mathcal{G}(t)$ is an EBB envelope).

The proof of the lemma follows directly from the Chernoff bound.

6.2 Statistical Service Curve

We use a statistical service curve model that combines the service curve models from Definitions 3.5 and 4.3. In this way the new service model can capture correlations between the service at different nodes, and yet it can account for statistical independence of service where available.

Definition 6.3 (STATISTICAL SERVICE CURVE) *A doubly-indexed random process $S(s, t)$ is a statistical service curve with error function $\varepsilon(\sigma)$ for an arrival process $A(t)$ if the corresponding departure process $D(t)$ satisfies for all $t \geq 0$ and σ*

$$Pr\left(D(t) < A * [S - \sigma]_+(t + \tau_0)\right) \leq \varepsilon(\sigma), \tag{6.3}$$

where τ_0 is a discretization parameter.

The random process $S(s, t)$ is decreasing in s , increasing in t , and satisfies $S(s, t) = S(s, u) + S(u, t)$ for all $0 \leq s \leq u \leq t$. The error function $\varepsilon(\sigma)$ is assumed to be nonnegative and nonincreasing.

This statistical service curve is different from the one used in the calculus from Section 4 in that $S(s, t)$ is now defined as a random process. If $S(s, t)$ is non-random and depends only on $t - s$ then Definition 6.3 reduces to Definition 4.3, and we further impose the condition on $\varepsilon(\sigma)$ from Eq. (4.8). If $S(s, t)$ is random and $\varepsilon(\sigma) = 0$ for all σ , then Definition 6.3 reduces to Definition 3.5 with $\tau_0 = 0$.

Since service curves are defined here as random processes, it is convenient to use an MGF envelope model to bound them (see Fidler [48]). Unlike the MGF envelope model from Definition 6.1 bounding the arrivals from above, we now bound the service curves from below.

Definition 6.4 (MGF BOUND FOR SERVICE CURVES) *A statistical service curve $S(s, t)$ has an MGF (exponential) bound, with rate r and prefactor M , for some choice of $\theta > 0$, if for all $0 \leq s \leq t$*

$$E [e^{-\theta S(s,t)}] \leq M e^{-\theta r(t-s)}. \quad (6.4)$$

In general, both the rate r and the prefactor M depend on θ .

The next result is an adaptation of Lemma 4.2 that was used for the derivation of single-node results in the calculus from Chapter 4. Unlike Lemma 4.2 where the arrivals are described with statistical envelopes, the arrivals herein are described with MGF envelopes.

Lemma 6.5 (SAMPLE-PATH BOUNDS WITH BOOLE'S INEQUALITY) *Suppose that an arrival process $A(t)$ is bounded for some choice of $\theta > 0$ by an MGF envelope with rate r_a and prefactor M_a . For some discretization parameter τ_0 , let a statistical service curve $S(s, t)$ statistically independent of $A(t)$. For the same θ , $S(s, t)$ has an MGF bound with rate r_s and prefactor $M_s = M'_s \binom{\lfloor \frac{t-s}{\tau_0} \rfloor + H - 1}{H-1}$ for some integer $H > 0$, where M'_s does not depend on $t - s$. Denote $M = M_a M'_s$, $r = r_s - r_a$, and assume for stability that $r > 0$. Then we have for all*

$t \geq 0$ and σ

$$Pr \left(\sup_{0 \leq s \leq t} \{A(t) - A(s) - S(s, t + \tau_0)\} > \sigma \right) \leq M \left(\frac{1}{\theta r \tau_0} \right)^H e^{-\theta \sigma}. \quad (6.5)$$

In applications, H corresponds to the number of nodes. The case when M_s does not depend on $t - s$ corresponds to $H = 1$. The complementary case correspond to $H > 1$; the dependency is caused by the binomial factor which generally stems from evaluating multi-node convolutions (for further technical details see Theorem 6.8).

PROOF. Fix $t \geq 0$ and σ . Using the discretization technique used in the proof of Lemma 4.2, we can write

$$\begin{aligned} & Pr \left(\sup_{0 \leq s \leq t} \{A(t) - A(s) - [S(s, t + \tau_0) - \sigma]_+\} > 0 \right) \\ & \leq Pr \left(\sup_{j \geq 0} \{A(t) - A([t - (j + 1)\tau_0]_+) - S(t - j\tau_0, t + \tau_0)\} > \sigma \right) \\ & \leq M \sum_{j \geq 1} \binom{j + H - 1}{H - 1} e^{-\theta r j \tau_0} e^{-\theta \sigma} \\ & \leq M \left(1 - \left(\frac{1}{1 - e^{-\theta r \tau_0}} \right)^H \right) e^{-\theta \sigma} \\ & \leq M \left(\frac{1}{\theta r \tau_0} \right)^H e^{-\theta \sigma}. \end{aligned}$$

In the third line we applied Boole's inequality. In the fourth line we used $\sum_{j \geq 0} \binom{j + H - 1}{H - 1} a^j = \left(\frac{1}{1 - a} \right)^H$ for all $0 < a < 1$ (see [48]). Last we used that $\left(\frac{1}{1 - e^{-x}} \right)^H - 1 \leq \left(\frac{1}{x} \right)^H$ for all $x > 0$. The proof is thus complete. \square

For the rest of this section we show how to construct statistical leftover service curves that satisfy Definition 6.4. The presented result will be used in the next chapter for the derivation of end-to-end performance bounds in networks with cross traffic where arrivals are described with MGF envelopes as in Definition 6.1.

Theorem 6.6 (STATISTICAL LEFTOVER SERVICE CURVE) *Consider a node with capacity C serving two arrival processes $A(t)$ and $A_c(t)$, whose corresponding departure processes are $D(t)$ and $D_c(t)$, respectively. Assume that $A_c(t)$ is bounded by an MGF envelope with rate r_c and prefactor 1 for some choice of $\theta > 0$, and that $r_c < C$. Fix τ_0 . Then we have the following two constructions.*

1. *The process*

$$S(s, t) = [C(t - s - \tau_0) - A_c(s, t - \tau_0)]_+ \quad (6.6)$$

is a statistical service curve for $A(t)$ with error function zero. It has an MGF bound with rate $C - r_c$ and prefactor $e^{\theta(C-r_c)\tau_0}$.

2. *For any choice of $\delta > 0$ the process*

$$S(s, t) = [C - r_c - \delta]_+(t - s) \quad (6.7)$$

is a statistical service curve for $A(t)$ with error function $\varepsilon(\sigma) = \frac{e^{\theta C \tau_0}}{\theta \delta \tau_0} e^{-\theta \sigma}$.

On one hand, the construction from Eq. (6.6) extends the corresponding construction from Theorem 3.6 to a continuous time setting. It is generally useful when $A(t)$ and $A_c(t)$ are statistically independent. On the other hand, the construction from Eq. (6.7) extends the corresponding construction from Theorem 4.5 to arrivals described with MGF envelopes; in this case $S(s, t)$ is a non-random function. The second construction is useful when $A(t)$ and $A_c(t)$ may be statistically correlated.

PROOF. Fix $t \geq 0$ and $\delta > 0$. For the first case we invoke Theorem 3.6 and obtain that the function $T(s, t) = [C(t - s) - A_c(s, t)]_+$ satisfies

$$D(t) \geq A * T(t) \text{ a.s. } . \quad (6.8)$$

Next, as in Lemma 4.4, it can be shown that the function $S(s, t) = T(s, t - \tau_0)$ is a statistical leftover service curve in the sense of Definition 6.3. The corresponding MGF bound for the

service curve $S(s, t)$ follows from expanding $E [e^{-\theta S(s, t)}]$.

For the second case we first apply Lemma 6.2 and obtain that the function $\mathcal{G}(t) = r_c t$ is a statistical envelope for $A_c(t)$ with error function $\varepsilon(\sigma) = e^{-\theta\sigma}$. Then the claim from Eq. (6.7) follows directly by invoking Theorem 4.5, which completes the proof. \square

6.3 Single-Node Performance Bounds

The next result gives single-node performance bounds for a flow having the arrivals described with MGF envelopes as in Definition 6.1, and service described with the service curve model from Definition 6.3.

Theorem 6.7 (PROBABILISTIC PERFORMANCE BOUNDS) *Consider a flow at a node with arrivals and departures denoted by the processes $A(t)$ and $D(t)$, respectively. For some discretization parameter τ_0 , assume that the service available to the flow is given by a statistical service curve $S(s, t)$, that is statistically independent from $A(t)$, with error function $\varepsilon^s(\sigma)$. For some choice of $\theta > 0$ assume that $A(t)$ is bounded by an MGF envelope with rate r_a and prefactor M_a , and also $S(s, t)$ has an MGF bound with rate r_s and prefactor $M_s = M'_s \left(\frac{\lfloor \frac{t-s}{\tau_0} \rfloor + H - 1}{H - 1} \right)$ for some integer $H > 0$, where M'_s does not depend on time parameters. Denote $M = M_a M'_s$, $r = r_s - r_a$, and assume for stability that $r > 0$. Also, let us define the error function*

$$\varepsilon(\sigma) = \inf_{\sigma^a + \sigma^s = \sigma} \left\{ M \left(\frac{1}{\theta r \tau_0} \right)^H e^{-\theta \sigma^a} + \varepsilon^s(\sigma^s) \right\}. \quad (6.9)$$

Then we have the following probabilistic bounds.

1. **OUTPUT RATE ENVELOPE:** *If $H = 1$ and $\varepsilon^s(\sigma) = 0$ for all σ , then the output process $D(t)$ is bounded by an MGF envelope with rate r_a and prefactor $M \left(\frac{1}{\theta r \tau_0} \right)^H$.*

2. BACKLOG BOUND: A bound on the backlog process $B(t)$ is given for all $t, \sigma \geq 0$ by

$$\Pr\left(B(t) > \sigma\right) \leq \varepsilon(\sigma). \quad (6.10)$$

3. DELAY BOUND: A bound on the delay process $W(t)$ is given for all $t, \sigma \geq 0$ by

$$\Pr\left(W(t) > \frac{\sigma}{r_s}\right) \leq \varepsilon(\sigma). \quad (6.11)$$

Let us make some observations about the theorem. First, in the case when the service curve is given by the function $S(s, t) = r_s(t - s)$, the backlog and delay bounds in the theorem are exactly the bounds obtained with Theorem 4.6. This can be shown by applying Lemma 6.2 to express the arrival process with a statistical envelope in the sense of Definition 4.1. Second, in the case when $S(s, t) = r_s(t - s)$ and $\varepsilon^s = 0$, then the backlog bound was derived by Chang in [26] in a discrete-time setting. If $\varepsilon^s(\sigma) = 0$ similar bounds were derived by Fidler in [48] in a discrete-time setting. Last, the results in the theorem depend on the discretization parameter τ_0 ; we will later apply these results and show how to optimize τ_0 using convex optimizations.

PROOF. Fix $\tau_0 > 0$. To prove the output rate envelope let us choose $0 \leq s \leq t$, such that we can write

$$\begin{aligned} E \left[e^{\theta(D(t) - D(s))} \right] &\leq E \left[e^{\theta(A(t) - A * S(s + \tau_0))} \right] \\ &\leq E \left[e^{\theta \sup_{j \geq 0} \{A(t) - A([s - (j+1)]_+) - S((j+1)\tau_0)\}} \right] \\ &\leq M_a M'_s e^{\theta r_a(t-s)} \sum_{j \geq 1} \binom{j + H - 1}{H - 1} e^{-\theta(r_s - r_a)j\tau_0} \\ &\leq M \left(\frac{1}{\theta r \tau_0} \right)^H e^{\theta r_a(t-s)}. \end{aligned}$$

In the first line we used that $D(t) \leq A(t)$ and that $D(s) \geq A * S(s + \tau_0)$ *a.s.* In the second line we used the discretization technique used in the proof of Lemma 4.2. In the third line we

used Boole's inequality and the rest follows as in the proof of Theorem 6.5.

To prove the backlog bound, let us choose $\sigma \geq 0$ and σ^a, σ^s such that $\sigma^a + \sigma^s = \sigma$. Assume that for a particular sample-path the following inequalities

$$A(t) - A(s) \leq S(s, t + \tau_0) + \sigma^a, \quad (6.12)$$

hold for all $0 \leq s \leq t$. Also, assume that

$$D(t) \geq A * [S - \sigma^s]_+(t + \tau_0). \quad (6.13)$$

Then we can express the backlog process as follows.

$$\begin{aligned} B(t) &\leq A(t) - D(t) \\ &\leq \sup_{0 \leq s \leq t + \tau_0} \{A(t) - A(s) - [S(s, t + \tau_0) - \sigma^s]_+\} \\ &\leq \sigma. \end{aligned}$$

In the second line we used Eq. (6.13), and in the third line we used Eq. (6.12).

Since we started by assuming Eqs. (6.12) and (6.13), we arrive at

$$\begin{aligned} Pr(B(t) > \sigma) &\leq P(\text{Eqs. (6.12) or (6.13) fail}) \\ &\leq M \left(\frac{1}{\theta r \tau_0} \right)^H e^{-\theta \sigma^a} + \varepsilon^s(\sigma^s) \end{aligned}$$

In the last equation we applied Lemma 6.5 and the definition of the statistical service curve.

Minimizing over $\sigma^a + \sigma^s = \sigma$ we obtain Eq. (6.10).

Finally, to prove the delay bound, let us choose $t, \sigma \geq 0$ and σ^a, σ^s such that $\sigma^a + \sigma^s = \sigma$.

Denote $d = \frac{\sigma}{r_s}$, and assume that for a particular sample-path the inequalities

$$A(t - d) - A(s) \leq [S(s, t - d + \tau_0) + S(t - d + \tau_0, t + \tau_0) - \sigma^s]_+ \quad (6.14)$$

hold for all $0 \leq s \leq t - d$. Also, assume that Eq. (6.13) holds. From Eq. (6.14) we successively obtain

$$\begin{aligned}
& \sup_{0 \leq s \leq t-d} \{A(t-d) - A(s) - [S(s, t + \tau_0) - \sigma^s]_+\} \leq 0 \\
& \Rightarrow \sup_{0 \leq s \leq t+\tau_0} \{A(t-d) - A(s) - [S(s, t + \tau_0) - \sigma^s]_+\} \leq 0 \\
& \Rightarrow A(t-d) \leq D(t) \\
& \Rightarrow W(t) \leq d.
\end{aligned}$$

In the second line we extended the range of the supremum using the positivity constraints. In the third line we applied Eq. (6.13), and then we used the definition of the delay process.

As in the proof for the backlog bound, we arrive at

$$\begin{aligned}
Pr(W(t) > d) & \leq P(\text{Eqs. (6.14) or (6.13) fail}) \\
& \leq M \left(\frac{1}{\theta r \tau_0} \right)^H e^{-\theta r_s d} e^{\theta \sigma^s} + \varepsilon^s(\sigma^s) \\
& \leq M \left(\frac{1}{\theta r \tau_0} \right)^H e^{-\theta \sigma^a} + \varepsilon^s(\sigma^s).
\end{aligned}$$

In the second line we applied Lemma 6.5 with $\sigma = S(t - d + \tau_0, t + \tau_0) - \sigma^s$. The proof is completed by minimizing over $\sigma^a + \sigma^s = \sigma$. \square

6.4 Statistical Network Service Curve

Let us consider now a flow traversing H nodes in series. For some discretization parameter $\tau_0 > 0$ and all $h = 1, \dots, H$, assume that $S^h(s, t)$ is a statistical service curve for the flow at node h with error function $\varepsilon^h(\sigma)$, in the sense of Definition 6.3. Next we provide the construction for the corresponding statistical network service curve for the flow. We distinguish two cases.

If all the error functions $\varepsilon^h(\sigma) = 0$ for all σ , then we have the same statistical network

service as in Eq. (3.47), i.e.,

$$S^{net}(s, t) = S^1 * S^2 * \dots * S^H(s, t). \quad (6.15)$$

Otherwise, if $\varepsilon^h(\sigma) \geq 0$ then the corresponding statistical network service curve is given for any choice of $\delta > 0$, as in Eq. (4.25), by

$$S^{net}(s, t) = S^1 * S_{-\delta}^2 * \dots * S_{-(H-1)\delta}^H(s, t), \quad (6.16)$$

with the error function from Eq. (4.26) (the proof for Eq. (6.16) follows the lines of the proof for Theorem 4.7, but for service curves defined with doubly-indexed random processes). In applications, we require the statistical independence of the service curves $S^h(s, t)$ in both Eqs. (6.15) and (6.16).

Having the expressions from Eqs. (6.15) and (6.16), we next derive the MGF bounds for the statistical network service curves. To keep the notation simple, we only consider the special case of identical service curves.

Theorem 6.8 (MGF BOUND FOR STATISTICAL NETWORK SERVICE CURVE). *Consider the multi-node scenario from the beginning of this section. For some choice of $\theta > 0$, assume that the service curves $S^h(s, t)$ are independent, and each has an MGF bound with rate r_s and prefactor M_s that does not depend on time parameters. Then the flow's statistical network service curve has the MGF bound*

$$E \left[e^{-\theta S^{net}(s, t)} \right] \leq M^{net} e^{-\theta r_s (t-s)}, \quad (6.17)$$

where the prefactor M^{net} depends on the construction of the network service curve.

1. If the statistical network service curve is given by Eq. (6.15) then

$$M^{net} = M_s^H \binom{\lfloor \frac{t-s}{\tau_0} \rfloor + H - 1}{H - 1} e^{(H-1)\theta r_s \tau_0}. \quad (6.18)$$

2. If the statistical network service curve is given by Eq. (6.16) then

$$M^{net} = M_s^H \binom{\lfloor \frac{t-s}{\tau_0} \rfloor + H - 1}{H - 1} e^{(H-1)\theta(r_s + \delta + \delta \lfloor \frac{t-s}{\tau_0} \rfloor)\tau_0}. \quad (6.19)$$

We remark that the construction from Eq. (6.17), with network service curve as in Eq. (6.15), extends a result of Fidler [48] to the continuous time setting.

PROOF. Fix $\delta, \tau_0 > 0$ and $0 \leq s \leq t$. In the first case we can expand the MGF of $S^{net}(s, t)$ by applying Boole's inequality and the discretization technique used in the proof of Lemma 4.2.

$$\begin{aligned} E \left[e^{-\theta S^{net}(s,t)} \right] &\leq E \left[\sup_{s \leq x_1 \leq \dots \leq x_{H-1} \leq t} e^{-\theta(S^1(s, x_1) + \dots + S^H(x_{H-1}, t))} \right] \\ &\leq \sum_{0 \leq j_1 \leq \dots \leq j_{H-1} \leq \lfloor \frac{t-s}{\tau_0} \rfloor} E \left[e^{-\theta(S^1(s, [t-(j_1+1)]_+ \tau_0) + \dots + S^H((t-j_{H-1})\tau_0, t))} \right] \\ &\leq M^H e^{(H-1)\theta r_s \tau_0} e^{-\theta r_s (t-s)} \sum_{0 \leq j_1 \leq \dots \leq j_{H-1} \leq \lfloor \frac{t-s}{\tau_0} \rfloor} 1 \\ &\leq M^H \binom{\lfloor \frac{t-s}{\tau_0} \rfloor + H - 1}{H - 1} e^{(H-1)\theta r_s \tau_0} e^{-\theta r_s (t-s)}. \end{aligned}$$

In the third line we expanded the MGF of $S^h(s, t)$ by using statistical independence, and then collected terms. In the fourth line the binomial coefficient is the number of combinations with repetitions.

For the second case we proceed similarly as in the first case.

$$\begin{aligned} E \left[e^{-\theta S^{net}(s,t)} \right] &\leq E \left[\sup_{s \leq x_1 \leq \dots \leq x_{H-1} \leq t} e^{-\theta(S^1(s, x_1) + \dots + S_{-(H-1)\delta}^H(x_{H-1}, t))} \right] \\ &\leq E \left[\sup_{s \leq x_1 \leq \dots \leq x_{H-1} \leq t} e^{-\theta(S^1(s, x_1) + \dots + S^H(x_{H-1}, t))} e^{\theta \delta ((H-1)t - (x_1 + \dots + x_{H-1}))} \right] \\ &\leq \sum_{0 \leq j_1 \leq \dots \leq j_{H-1} \leq \lfloor \frac{t-s}{\tau_0} \rfloor} E \left[e^{-\theta(S^1(s, (t-(j_1+1))\tau_0) + \dots + S^H((t-j_{H-1})\tau_0, t))} \right] e^{\theta \delta \sum_{h=1}^{H-1} (j_h+1)\tau_0} \\ &\leq M^H \binom{\lfloor \frac{t-s}{\tau_0} \rfloor + H - 1}{H - 1} e^{(H-1)\theta(r_s + \delta)\tau_0} e^{-\theta r_s (t-s)} e^{(H-1)\theta \delta \lfloor \frac{t-s}{\tau_0} \rfloor \tau_0}. \end{aligned}$$

In the last line we bounded each j_h by $\lfloor \frac{t-s}{\tau_0} \rfloor$. The proof is thus complete. \square

6.5 The Special Case of Stationary Processes with Independent Increments

We now consider the case of arrival and service processes having stationary and independent increments. Examples of such processes include Poisson or compound Poisson processes. We will show that by accounting for these properties we can improve the single-node results obtained so far. In this sense we present a tighter construction of a statistical leftover service curve, and also improved single-node performance bounds.

The stationary and independent increments properties of a process $A(t)$ are formally expressed as follows.

1. *Stationary Increments:*

$$Pr\left(A(s, t) \leq x\right) = Pr\left(A(s+u, t+u) \leq x\right) \quad (6.20)$$

for all $s \leq t$ and $u, x \geq 0$.

2. *Independent Increments:* $A(u, v)$ and $A(s, t)$ are statistically independent for all $u \leq v \leq s \leq t$.

As we have shown in Section 3.4, the evaluation of backlog and delay bounds in the stochastic network calculus reduces to the evaluation of sample-path bounds. The technique which we used so far in the network calculus formulations from Chapters 4 and 6 was based on Boole's inequality (see also Section 3.4.2). This technique does not exploit the independent increments properties of arrivals, where available, and is more appropriate for processes with correlated increments. To account for the independent increments properties, we propose to evaluate sample-path bounds using Doob's maximal inequality for supermartingales (see Lemma 5.11).

The technique of using maximal inequalities to evaluating sample-path bounds is applied in a classical note by Kingman [66] to derive exponential backlog bounds in GI/GI/1 queues. Since Kingman's note, several works use related supermartingales techniques to derive exponential bounds. For example, Ross [99] improves Kingman's bounds. Extensions of Kingman's bounds to the case of Markov-modulated arrivals are carried out by Buffet and Duffield [20], Artiges and Nain [3], or Liu *et. al.* [83]. An extension to the multi-node case for stochastic linear systems under the $(max, +)$ algebra is carried out by Chang [27].

The application of Doob's maximal inequality for continuous-time processes does not require a discretization parameter τ_0 . For this reason, we use the definition of a statistical service curve from Definition 6.3 with $\tau_0 = 0$.

We now present a critical result for the evaluation of sample-path bounds. Unlike Lemma 6.5 that uses Boole's inequality to evaluate sample-path bounds, the next result uses Doob's maximal inequality.

Lemma 6.9 (SAMPLE-PATH BOUNDS WITH DOOB'S MAXIMAL INEQUALITY) *Suppose that an arrival process $A(t)$ is bounded for some choice of $\theta > 0$ by an MGF envelope with rate r_a and prefactor $M_a \leq 1$. Let also a statistical service curve $S(s, t)$ that has an MGF bound with rate r_s and prefactor $M_s \leq 1$, for the same θ . Assume that the inequality*

$$r_a \leq r_s \tag{6.21}$$

holds, and that $A(t)$ and $S(s, t)$ are independent processes having stationary and independent increments. Then we have for all $t, \sigma \geq 0$

$$Pr(A(t) > A * S(t) + \sigma) \leq e^{-\theta\sigma} . \tag{6.22}$$

PROOF. Fix $t, \sigma \geq 0$. First, using the stationary and independent increments properties,

we can rewrite the probability in Eq. (6.22) as follows

$$Pr \left(\sup_{0 \leq s \leq t} \{A(s) - S(s)\} > \sigma \right) = Pr \left(\sup_{0 \leq s \leq t} e^{\theta(A(s) - S(s))} > e^{\theta\sigma} \right). \quad (6.23)$$

Let us now construct the process

$$Y(s) = e^{\theta(A(s) - S(s))}, \quad (6.24)$$

and the corresponding filtration $\mathbf{F} = \{\mathcal{F}_s : s \geq 0\}$ with $\mathcal{F}_s = \sigma\{Y(u) : 0 \leq u \leq s\}$, for all $s \geq 0$. For $u, s \geq 0$ we can write

$$\begin{aligned} E[Y(s+u) \mid \mathcal{F}_s] &= E[Y(s)e^{\theta(A(s,s+u) - S(s,s+u))} \mid \mathcal{F}_s] \\ &= Y(s)E[e^{\theta(A(s,s+u) - S(s,s+u))} \mid \mathcal{F}_s] \\ &\leq Y(s)M_a M_s e^{-\theta(r_s - r_a)u} \\ &\leq Y(s), \end{aligned}$$

which shows that $Y(s)$ is a nonnegative supermartingale; note that the integrability of $Y(s)$ is guaranteed by the a-priori bounds on the MGF's of $A(t)$ and $S(s, t)$. In the second and third lines we used the properties of conditional expectation from Lemma 5.10, and also the stationary and independent increments properties. In the last line we used the condition from Eq. (6.21).

Therefore, we can apply Doob's maximal inequality from Lemma 5.11 to evaluate Eq. (6.23)

$$\begin{aligned} Pr \left(\sup_{0 \leq s \leq t} e^{\theta(A(s) - S(s))} > e^{\theta\sigma} \right) &\leq E[Y(0)] e^{-\theta\sigma} \\ &\leq e^{-\theta\sigma}, \end{aligned}$$

which completes the proof. □

In the following we address the construction of a statistical leftover service curve. This improves Theorem 6.6 for the special case of arrival processes with stationary and independent increments. We only consider the case of arrivals which may be statistically correlated. Note that the construction of leftover service curves in the case of statistically independent arrivals does not involve the evaluation of sample-path bounds (see Eq. (6.6) in Theorem 6.6).

Theorem 6.10 (STATISTICAL LEFTOVER SERVICE CURVE FOR ARRIVALS WITH STATIONARY AND INDEPENDENT INCREMENTS) *Consider a node with capacity C serving two arrival processes $A(t)$ and $A_c(t)$, whose corresponding departure processes are $D(t)$ and $D_c(t)$, respectively. Assume that $A_c(t)$ has stationary and independent increments and is also bounded by an MGF envelope with rate r_c and prefactor $M \leq 1$ for some choice of $\theta > 0$, and that $r_c \leq C$. Then the function*

$$S(s, t) = (C - r_c)(t - s) \quad (6.25)$$

is a statistical service curve for $A(t)$ with error function $\varepsilon(\sigma) = e^{-\theta\sigma}$.

The rate of the leftover service curve from Eq. (6.25) is larger than the rate of the leftover service curve from Eq. (6.7). Also, unlike the error function from Theorem 6.6, the error function in Eq. (6.25) is not obtained as a series.

PROOF. Fix $t \geq 0$ and σ . Assume that for a particular sample-path the following inequality

$$A_c(t) \leq \inf_{0 \leq s \leq t} \{A_c(s) + r_c(t - s) + \sigma\} . \quad (6.26)$$

holds. We recall from the proof of Theorem 2.3 that

$$D(t) \geq \inf_{0 \leq s \leq t} \{A(s) + [C(t - s) - (A_c(t) - A_c(s))]_+\} . \quad (6.27)$$

Inserting Eq. (6.26) into Eq. (6.27) yields

$$\begin{aligned} D(t) &\geq \inf_{0 \leq s \leq t} \{A(s) + [C(t-s) - r_c(t-s) - \sigma]_+\} \\ &\geq A * [S - \sigma]_+(t) . \end{aligned}$$

Since we started by assuming Eq. (6.26) we arrive at

$$\begin{aligned} Pr(D(t) < A * [S - \sigma]_+(t)) &\leq Pr(\text{Eq. (6.26) fails}) \\ &\leq e^{-\theta\sigma} . \end{aligned}$$

Here we applied Lemma 6.9 for the arrival process $A_c(t)$ and the function $\mathcal{T}(s, t) = r_c(t-s)$ which is to replace the process $S(s, t)$ from Lemma 6.9. The proof is thus complete. \square

For the rest of this section we address the derivation of single-node performance bounds. In this sense, we extend Theorem 6.7 to the special case of processes with stationary and independent increments.

Theorem 6.11 (PROBABILISTIC PERFORMANCE BOUNDS FOR PROCESSES WITH STATIONARY AND INDEPENDENT INCREMENTS) *Consider a flow at a node with arrivals and departures given by the processes $A(t)$ and $D(t)$, respectively. The service available to the flow is given by a statistical service curve $S(s, t)$ with error function $\varepsilon^s(\sigma)$. For some choice of $\theta > 0$ assume that $A(t)$ is bounded by an MGF envelope with rate r_a and prefactor $M_a \leq 1$, and also $S(s, t)$ has an MGF bound with rate r_s and prefactor $M_s \leq 1$. Assume that the inequality*

$$r_a \leq r_s$$

holds, and that $A(t)$ and $S(s, t)$ are independent processes having stationary and independent increments. Also, let us define the error function

$$\varepsilon(\sigma) = \inf_{\sigma^a + \sigma^s = \sigma} \{e^{-\theta\sigma^a} + \varepsilon^s(\sigma^s)\} .$$

Then we have the following probabilistic bounds:

1. BACKLOG BOUND: A bound on the backlog process $B(t)$ is given for all $t, \sigma \geq 0$ by

$$\Pr(B(t) > \sigma) \leq \varepsilon(\sigma). \quad (6.28)$$

2. DELAY BOUND: A bound on the delay process $W(t)$ is given for all $t, \sigma \geq 0$ by

$$\Pr\left(W(t) > \frac{\sigma}{r_s}\right) \leq \varepsilon(\sigma). \quad (6.29)$$

PROOF. Let $t, \sigma \geq 0$ and σ^a, σ^s such that $\sigma^a + \sigma^s = \sigma$. Following the proof for the backlog bound in Theorem 6.7 we arrive at

$$\begin{aligned} \Pr(B(t) > \sigma) &\leq \Pr\left(\sup_{0 \leq s \leq t} \{A(t) - A(s) - S(s, t)\} > \sigma^a\right) \\ &\quad + \Pr(D(t) < A * [S - \sigma^s]_+(t)) \\ &\leq e^{-\theta\sigma^a} + \varepsilon^s(\sigma^s). \end{aligned}$$

In the last equation we applied Lemma 6.9 and the definition of the statistical service curve.

Let us denote $d = \frac{\sigma}{r_s}$. Following the proof for the delay bound in Theorem 6.7 we arrive at

$$\begin{aligned} \Pr(W(t) > d) &\leq \Pr\left(\sup_{0 \leq s \leq t-d} \{A(t-d) - A(s) - S(s, t-d)\} > S(t-d, t) - \sigma^s\right) \\ &\quad + \Pr(D(t) < A * [S - \sigma^s]_+(t)) \\ &\leq e^{-\theta r_s d} e^{\theta\sigma^s} + \varepsilon^s(\sigma^s) \\ &\leq e^{-\theta\sigma^a} + \varepsilon^s(\sigma^s). \end{aligned}$$

In the second line we applied Lemma 6.9 and the definition of the statistical service curve, and finally we replaced d with $\frac{\sigma}{r_s}$. The proof is completed by minimizing over $\sigma^a + \sigma^s = \sigma$.

□

Chapter 7

Comparison of Delay Bounds with Exact Results

Unlike other theories for network analysis which express performance metrics such as backlog or delay in terms of *exact results*, the stochastic network calculus yields performance *bounds*. For this reason, a potential concern in using the calculus for network analysis is whether the obtained backlog or delay bounds are sufficiently tight.

The purpose of this chapter is to shed light on the accuracy of performance bounds obtained with the calculus. Towards this goal, we derive network calculus bounds in network scenarios where exact results are available, and then show numerical comparisons between the obtained bounds and the exact results.

There are several factors which may influence the accuracy of network calculus performance bounds. First, traffic is usually expressed in terms of statistical envelopes which set bounds on the arrivals. Second, the common technique for evaluating sample-path bounds relies on Boole's inequality (see Section 3.4.2) which may be loose since it allows for correlations within arrival processes. Third, the analysis of a network with cross traffic relies on a worst-case service view for the through traffic by constructing leftover service curves.

Statistical envelopes provide tight characterizations for a wide class of arrival processes due to accounting for statistical multiplexing (see Boorstyn *et. al.* [12], or Li *et. al.* [76]). For this reason, our comparison study will mainly focus on the impact of using Boole's inequality to the estimation of sample-path bounds, and the impact of using leftover service curves to

derive performance bounds.

The next two sections treat the single-node and the multi-node cases, respectively. The third section then provides numerical comparisons for each case.

7.1 The Single-Node Case

In this section we derive network calculus delay bounds for three queueing models where exact results are available: M/M/1, M/D/1 and M/M/1 queues with two priorities. We find that the bounds obtained are tight at most utilization factors. This provides evidence that the method of using Boole's inequality for estimating sample-path bounds is quite accurate. Also, we show improved bounds by accounting for the independent increments property of arrivals.

For the case of the M/M/1 queue with priorities we focus on the expected delay for the lower priority packets. We derive bounds by first constructing a statistical leftover service curve, and then deriving bounds on the expected delay. The obtained bounds are very accurate when compared to the exact results, unless the amount of lower priority flows is negligible when compared to the amount of high priority packets. This indicates that leftover service curves provide an accurate characterization for the service received by the lowest priority flows at a SP scheduler.

7.1.1 The M/M/1 queue

We consider the following M/M/1 queueing model. Let a node with capacity C where packets arrive according to an exponential distribution with mean inter-arrival distance $1/\lambda$. The size of each packet i is denoted by X_i and is exponentially distributed with mean $1/\mu$. We assume that the inter-arrival times and the packet sizes are statistically independent. We denote the utilization factor by $\rho = \lambda/(\mu C)$, and assume for stability that $\rho < 1$.

The exact distribution of the steady-state delay of packets W' in the M/M/1 queue is given by [68]

$$P(W' > d) = e^{-\mu C(1-\rho)d} . \quad (7.1)$$

Next we derive two network calculus (virtual) delay bounds for the M/M/1 queue according to Eq. (2.1). First we invoke Theorem 6.7 that evaluates sample-path bounds with Boole's inequality, and second we invoke Theorem 6.11 that evaluates sample-path bounds with Doob's inequality by accounting for the independent increments properties of arrivals.

To fit the queueing model with network calculus, we model the arrival process as in Subsection 5.2, i.e.,

$$A(t) = \sum_{i=1}^{N(t)} X_i . \quad (7.2)$$

where $N(t)$ is a Poisson process with mean λt .

Since the effective bandwidth of $A(t)$ is $\alpha_A(\theta, t) = \frac{\lambda}{\mu - \theta}$ for some choice of θ with $0 < \theta < \mu$, we obtain that $A(t)$ is bounded by an MGF envelope (see Definition 6.1) with rate and prefactor given by

$$r_a = \frac{\lambda}{\mu - \theta} , \quad M_a = 1 . \quad (7.3)$$

On the other hand, the service at a single queue can be modelled for any choice of the discretization parameter τ_0 with the statistical service curve

$$S(s, t) = C(t - s - \tau_0) ,$$

and error function $\varepsilon(\sigma) = 0$ for $\sigma \geq 0$, and $\varepsilon(\sigma) = 1$ for $\sigma < 0$, according to Definition 6.3.

We have that $S(s, t)$ has an MGF bound with rate and prefactor given by

$$r_s = C , \quad M_s = e^{\theta C \tau_0} , \quad (7.4)$$

for the same $\theta > 0$ as above.

Choosing

$$\theta < \mu(1 - \rho), \quad (7.5)$$

we see that $r_a < r_s$, as necessary for Theorem 6.7. We can now plug the values from Eqs. (7.3) and (7.4) into Theorem 6.7 and obtain the delay bound for all $\sigma \geq 0$

$$\begin{aligned} Pr\left(W(t) > \frac{\sigma}{r_s}\right) &\leq \inf_{\sigma^a + \sigma^s = \sigma} \left\{ \frac{M_a M_s}{\theta (r_s - r_a) \tau_0} e^{-\theta \sigma^a} + \varepsilon(\sigma^s) \right\} \\ &\leq \frac{e^{\theta C \tau_0}}{\theta \left(C - \frac{\lambda}{\mu - \theta}\right) \tau_0} e^{-\theta \sigma}. \end{aligned}$$

Optimizing $\tau_0 = 1/(\theta C)$ with Lemma 5.4, replacing σ with dC , and letting $t \rightarrow \infty$ yields the steady-state delay bound for all $d \geq 0$

$$Pr(W > d) \leq \frac{1}{1 - \frac{\rho}{1 - \frac{\theta}{\mu}}} e^{-\theta C d}. \quad (7.6)$$

This bound can be further optimized over θ , subject to the constraint from Eq. (7.5).

Let us now derive a delay bound using Theorem 6.11 that accounts for the independent increments property. Since there is no need for the discretization parameter τ_0 , the statistical service curve is now given by

$$S(s, t) = C(t - s),$$

with error function $\varepsilon(\sigma) = 0$ for $\sigma \geq 0$, and $\varepsilon(\sigma) = 1$ for $\sigma < 0$, satisfying Definition 6.3 with $\tau_0 = 0$. We have that $S(t)$ has an MGF bound with rate and prefactor given by

$$r_s = C, \quad M_s = 1. \quad (7.7)$$

The processes $A(t)$ and $S(s, t)$ have stationary and independent increments, and are also statistically independent. As before, by imposing the condition on θ from Eq. (7.5), we can plug the values from Eqs. (7.3) and (7.7) into Theorem 6.11 and obtain the delay bound for

all $d \geq 0$

$$\begin{aligned} Pr(W(t) > d) &\leq \inf_{\sigma} \{e^{-\theta r_s d} e^{\theta \sigma} + \varepsilon(\sigma)\} \\ &\leq e^{-\theta C d}. \end{aligned}$$

In this case we can optimize the value of θ by letting $\theta \rightarrow \mu(1 - \rho)$, yielding the steady-state delay bound

$$P(W > d) \leq e^{-\mu C(1-\rho)d}. \quad (7.8)$$

Although Eq. (7.8) looks exactly as Eq. (7.1), we would like to point out that W' and W are different steady-state delays (for more details see the related discussion from the beginning of Section 5.2.2).

The next result gives the asymptotic improvement of the delay bound from Eq. (7.8) over the delay bound from Eq. (7.6), when the utilization factor ρ approaches one.

Theorem 7.1 For $0 < z < 1$, let $d_1(z)$ and $d_2(z)$ be the bounds on the z -quantiles from Eq. (7.6) and Eq. (7.8), respectively. Then for $\rho \rightarrow 1$ we have

$$d_1(z) - d_2(z) = \Omega\left(\frac{\log(1 - \rho)^{-1}}{1 - \rho}\right).$$

The result shows that the benefits of using the independent increments property are substantial at very high utilizations.

PROOF. From Eq. (7.6) we have that

$$d_1(z) = -\frac{1}{\theta C} \log\left(1 - \frac{\rho}{1 - \frac{\theta}{\mu}}\right) (1 - z),$$

and from Eq. (7.8) we have that

$$d_2(z) = \frac{1}{\theta C} \log \frac{1}{1 - z},$$

Using that $\theta < \mu$ we get

$$d_1(z) - d_2(z) \geq \frac{1}{\theta C} \log \left(\frac{1}{1 - \rho} \right) .$$

Finally, using the constraint on θ from Eq. (7.5), the main claim follows. \square

7.1.2 The M/M/1 queue with priorities

We now consider an M/M/1 network node serving a through and a cross flow. The packets of the flows arrive according to Poisson processes with rates λ and λ_c , respectively, and have exponentially distributed sizes with mean $1/\mu$. The packets of the cross flow have *preemptive* priority over the packets of the through flow. The capacity of the server is C , and we assume for stability that the utilization factor $\rho = (\lambda + \lambda_c)/(\mu C)$ is less than one.

With queueing theory we have the exact result on the expectation $E[W']$ of the steady-state delay of the low priority flow's packets (see Gross and Harris [53])

$$E[W'] = \frac{\mu C - \lambda_c(1 - \rho)}{\mu C (\mu C - \lambda_c)(1 - \rho)} . \quad (7.9)$$

Next we derive a bound with network calculus on $E[W]$, where W is the corresponding steady-state virtual delay according to Eq. (2.1).

The through and the cross flows are represented by the processes $A(t) = \sum_{i=1}^{N(t)} X_i$ and $A_c(t) = \sum_{i=1}^{N_c(t)} Y_i$, respectively, where $N(t)$ and $N_c(t)$ are Poisson processes with rates λ and λ_c , respectively; also X_i and Y_i are independent and exponentially distributed with mean $1/\mu$. As before we obtain that $A(t)$ and $A_c(t)$ are bounded by MGF envelopes with rates and prefactors given by

$$\begin{aligned} r_a &= \frac{\lambda}{\mu - \theta}, \quad M_a = 1, \\ r_c &= \frac{\lambda_c}{\mu - \theta}, \quad M_c = 1, \end{aligned} \quad (7.10)$$

for some choice of θ with $0 < \theta < \mu$.

Having the description of the cross flow from above we can apply Theorem 6.10 and obtain that the function

$$S(s, t) = C(t - s) - A_c(s, t)$$

is a statistical leftover service curve for the through flow satisfying Definition 6.3 with $\tau_0 = 0$.

We then obtain that $S(s, t)$ has an MGF bound with rate and prefactor given by

$$r_s = C - \frac{\lambda_c}{\mu - \theta}, \quad M_s = 1. \quad (7.11)$$

Choosing

$$\theta < \mu(1 - \rho), \quad (7.12)$$

we see that $r_a < r_s$, as required for Theorem 6.11. Since $A(t)$ and $S(s, t)$ have stationary and independent increments, we can plug the values from Eqs. (7.10) and (7.11) into Theorem 6.11 and obtain the steady-state delay bound for all $d \geq 0$

$$P(W > d) \leq e^{-\frac{\lambda(1-\rho)}{\rho}d}, \quad (7.13)$$

after letting $\theta \rightarrow \mu(1 - \rho)$ (see Eq. (7.12)), and $t \rightarrow \infty$.

Next, using the formula for the expectation

$$E[W] = \int_0^\infty Pr(W > x) dx,$$

we obtain the upper bound

$$E[W] \leq \frac{\rho}{\lambda(1 - \rho)}. \quad (7.14)$$

Let us compare asymptotically the bound from Eq. (7.14) with the exact result from Eq. (7.9). By letting the percentage of through traffic approach zero (i.e. $\lambda \rightarrow 0$), we observe that the exact result is bounded, whereas the upper bound converges to infinity. This indicates

that for very small amount of through traffic, the upper bound is overly pessimistic (see also the corresponding numerical results for further observations).

7.1.3 The M/D/1 queue

The M/D/1 queueing model is similar to the M/M/1 queueing model, with the exception that the size X_i of each packet i is now constant, i.e., $X_i = 1/\mu$ for all i .

The distribution of the steady-state delay W' of packets is given by [57]

$$P(W' > d) = 1 - (1 - \rho)e^{\lambda d} \sum_{k=0}^T \frac{(k\rho - \lambda d)^k}{k!} e^{-(k-1)\rho}, \quad (7.15)$$

where $T = \lfloor d\mu \rfloor$ denotes the largest integer less than or equal to $d\mu$. It is known that this formula poses numerical complications when ρ approaches one, due to the appearance of large alternating, very nearly cancelling terms (note that the factor $k\rho - \lambda d$ is negative). There are several numerical algorithms to evaluate Eq. (7.15), of which we choose one due to Iversen and Staalhagen [57].

We next derive two network calculus virtual delay bounds (according to Eq. (2.1)) with Theorems 6.7 and 6.11.

The network calculus models for fitting the M/D/1 queue are similar to those for the M/M/1 queue with one difference. Since the size of each packet is now a constant, we have that the effective bandwidth of the arrival process $A(t)$ from Eq. (7.2) is given for all $\theta > 0$ by

$$\begin{aligned} \alpha_A(\theta, t) &= \frac{1}{\theta t} \log (E [e^{\theta A(t)}]) \\ &= \frac{1}{\theta t} \log \left(\sum_{n \geq 0} E [e^{\theta \sum_{i=1}^n X_i}] e^{-\lambda t} \frac{(\lambda t)^n}{n!} \right) \\ &= \frac{\lambda}{\theta} \left(e^{\frac{\theta}{\mu}} - 1 \right) \end{aligned}$$

From here it follows that $A(t)$ is now bounded by an MGF envelope with rate and prefactor

given by

$$r_a = \frac{\lambda}{\theta} \left(e^{\frac{\theta}{\mu}} - 1 \right), \quad M_a = 1. \quad (7.16)$$

Choosing

$$\rho \frac{\mu}{\theta} \left(e^{\frac{\theta}{\mu}} - 1 \right) < 1, \quad (7.17)$$

we see that $r_a < r_s$, as required for Theorem 6.7. We can now plug the values from Eqs. (7.16) and (7.4) into Theorem 6.7, and after optimizing $\tau_0 = 1/(\mu C)$, we obtain the steady-state delay bound for all $d \geq 0$

$$Pr(W > d) \leq \frac{e}{1 - \rho \frac{\mu}{\theta} \left(e^{\frac{\theta}{\mu}} - 1 \right)} e^{-\theta C d}. \quad (7.18)$$

The arrival process $A(t)$ is a compound Poisson processes, such that we can plug the values from Eqs. (7.16) and (7.7) into Theorem 6.11, and obtain the steady-state delay bound for all $d \geq 0$

$$P(W > d) \leq e^{-\theta C d}, \quad (7.19)$$

This bound is similar to the bound from Eq. (7.8) obtained for a M/M/1 queue; the difference is that the parameter θ is now subject to the restriction from Eq. (7.17), rather than Eq. (7.5).

To quantify the benefits of accounting for the independent increments property of arrivals, we next give a similar result to the one from Theorem 7.1, but for the M/D/1 queue.

Theorem 7.2 For $0 < z < 1$, let us denote $d_1(z)$ and $d_2(z)$ as the bounds on the z -quantiles of the delay bounds from Eq. (7.18) and Eq. (7.19), respectively. Then for $\rho \rightarrow 1$ we have

$$d_1(z) - d_2(z) = \Omega \left(\frac{\log(1 - \rho)^{-1}}{1 - \rho} \right).$$

We remark that the asymptotic benefits of accounting for the independent increments property are similar for the M/M/1 and M/D/1 queues.

PROOF. Following the the proof of Theorem 7.1, we arrive at

$$\begin{aligned} d_1(z) - d_2(z) &\geq \frac{1}{\theta C} \log \left(\frac{1}{1 - \rho \frac{\mu}{\theta} (e^{\frac{\theta}{\mu}} - 1)} \right) \\ &\geq \frac{1}{\theta C} \log \left(\frac{1}{1 - \rho} \right). \end{aligned}$$

In the last line we used the inequality $e^x \geq x + 1$ for all $x \geq 0$.

Furthermore, from the constraint on θ from Eq. (7.17) we obtain that $\theta = \mathcal{O}(1 - \rho)$, which completes the proof. \square

7.2 The Multi-Node Case

In this section we continue the discussion on the accuracy of network calculus bounds by analyzing multi-node networks. Besides deriving bounds in scenarios with exact solutions, we also derive bounds by relaxing the assumptions of the statistical independence of arrivals or service, that are necessary for the existence of exact solutions. The latter bounds provide insight into the role of statistical independence in network calculus.

We consider the network with cross traffic from Figure 2.1 where we make the following assumptions. The through flow and each of the cross flows consist of packets arriving according to Poisson processes with rates λ and λ_c , respectively. The size of each packet is exponentially distributed with mean $1/\mu$, and the size of each of the through flow's packets is independently regenerated at each traversed node [67]. We assume statistical independence among the Poisson processes and the sizes of packets. The network is stable, i.e., the utilization factor $\rho = (\lambda + \lambda_c)/(\mu C)$ is less than one. Under these assumptions, the network is an M/M/1 network and exact results are available [67].

If the scheduling at the nodes is FIFO, then the steady-state end-to-end delay of packets

$W^{net'}$ of the through flow has a Gamma distribution $\Gamma(\mu C(1 - \rho), H)$ [68]:

$$P(W^{net'} > d) = \left(\sum_{k=0}^{H-1} \frac{(\mu C(1 - \rho)d)^k}{k!} \right) e^{-\mu C(1 - \rho)d} . \quad (7.20)$$

In the following we use the network calculus formulation from Chapter 6 to derive bounds on the virtual end-to-end delay according to Eq. (2.1) (see also the discussion from the beginning of Section 5.2.2 concerning packets' delays vs. virtual delays). We consider three scenarios.

1. *Scenario with independent arrivals and independent service times* (Section 7.2.1): We consider that both arrivals and service times are statistically independent; exact solutions are available in this scenario (e.g. Eq. (7.20)). Unlike the exact results which hold for FIFO scheduling, the network calculus bounds are obtained by assuming that the through flow has a lower priority than the cross flows at the nodes. The obtained bounds are thus expected to be more pessimistic than the exact results.
2. *Scenario with correlated arrivals and independent service times* (Section 7.2.2): We consider the case when the cross flows may be correlated, but we account for the statistical independence due to regenerating the packets' sizes of the through flow at the traversed nodes. By comparing the bounds obtained in this scenario with the bounds from Section 7.2.1 we address the impact of accounting for the statistical independence of arrivals in network calculus.
3. *Scenario with independent arrivals and correlated service times* (Section 7.2.3): We consider the case when the packets of the through flow maintain their sizes at the traversed nodes (i.e. inducing statistical correlations of the service), but the arrivals are statistically independent. By comparing the bounds obtained in this scenario with the bounds from Section 7.2.1, we address the impact of accounting for the statistical independence of service in network calculus.

For numerical comparisons, we also consider the scenario with both correlated arrivals and service times. This is the most pessimistic scenario and the corresponding bounds are available in Section 5.2.

We point out that by allowing for correlated arrivals, our results are not restricted to some particular joint distributions on the arrivals, but they rather hold for any arbitrary correlation structures; the same remark holds in the case of correlated service.

The bounds for the above scenarios are derived using a packetized service model. For the first two scenarios we also derive bounds using a fluid service model. These bounds permit a discussion on whether the assumption of using a fluid service model, which yields simpler bounds than the packetized service model, is justified in network calculus.

To fit the queueing network with network calculus we represent the arrivals by compound Poisson processes, as done in Section 5.2. Accordingly, we obtain that the through arrival flow $A(t)$ is bounded by an MGF envelope with rate and prefactor given by

$$r_a = \frac{\lambda}{\mu - \theta}, \quad M_a = 1, \quad (7.21)$$

and each of the cross arrival flows $A_h(t)$ are bounded by an MGF envelope with rate and prefactor given by

$$r_c = \frac{\lambda_c}{\mu - \theta}, \quad M_c = 1,$$

for some choice of θ with $0 < \theta < \mu$ (see also Section 7.1.1).

To account for the packetized model of service, we recall from Section 5.2 that each node in the network from Figure 2.1 can be represented as the concatenation between a fluid flow server and a packetizer P^μ . The next result gives a statistical service curve representation for a packetizer, that is useful when the service times are statistically independent.

Lemma 7.3 *Consider the network node represented in Figure 5.2. Then the function*

$$S^\mu(s, t) = [C(t - s) - X_f(t)]_+,$$

is a statistical service curve for the packetizer P^μ with error function $\varepsilon^\mu = 0$, in the sense of Definition 6.3 with $\tau_0 = 0$, where $X_f(t)$ denotes the time already spent in service by the packet (if any) at the fluid server at time t .

We note that the theorem complements Lemma 5.5 which considers the case of correlated service times. Also, we remark our preference for a service curve construction which dispenses with a discretization parameter τ_0 . The reason is that when such service curves are convolved, we can obtain simpler expressions for the rate envelope of the corresponding network service curves (with Theorem 6.8). Then we can express the derived network service curve according to Definition 6.3, that further accounts for a discretization parameter, by invoking Lemma 4.4. Last we can derive performance bounds with Theorem 6.7.

PROOF. We closely follow the proof of Lemma 5.5. Fix $t \geq 0$, and let us denote \underline{t} as the beginning of the last busy period before t at the fluid server. If $X_f(t)$ denotes the fraction already processed of the packet currently in service at the fluid server at time t , then

$$u = t - \frac{X_f(t)}{C}$$

is the starting processing time of the packet currently serviced. It then follows that

$$\begin{aligned} D^1(t) &= D^1(\underline{t}) + C(u - \underline{t}) \\ &= A^1(\underline{t}) + C\left(t - \frac{X_f(t)}{C} - \underline{t}\right) \\ &= A^1(\underline{t}) + S^\mu(\underline{t}, t) \\ &\geq A^1 * S^\mu(t), \end{aligned}$$

which completes the proof. □

7.2.1 Independent arrivals and service

Here we derive network calculus delay bounds by accounting for the statistical independence of both arrivals and service times. In the last part we derive further bounds by using a fluid service model, rather than a packetized service model.

Let us consider the representation of each of the network nodes as in Figure 5.2 (i.e. as the concatenation between a fluid server and a packetizer P^μ). By enforcing the condition that $\theta < \mu - \lambda_c/C$, we can invoke Theorem 6.6 (by dispensing with the discretization parameter τ_0) and obtain that the function

$$T^h(s, t) = [C(t - s) - A_h(s, t)]_+ \quad (7.22)$$

is a statistical leftover service curve at the h^{th} fluid server. Then, by using the service curve representation of each packetizer as in Lemma 7.3, we further obtain with Eq. (6.15) that each node in the network from Figure 2.1 can be described with the statistical (network) service curve

$$\begin{aligned} S^h(s, t) &= T^h * S^{\mu, h}(s, t) \\ &= \inf_{s \leq u \leq t} \left([C(u - s) - A_h(s, u)]_+ + [C(t - u) - X_f^h(t)]_+ \right) \\ &\geq [C(t - s) - A_h(s, t) - X_f^h(t)]_+ , \end{aligned}$$

where $X_f^h(t)$ denotes the fraction already processed of the packet currently in service (if any) at node h at time t . Moreover, each service curve $S^h(s, t)$ has an MGF bound with rate and prefactor given by

$$r_s = C - r_c, \quad M_s = \frac{\mu}{\mu - \theta},$$

where we used that $E \left[e^{\theta X_f^h(t)} \right] = \frac{\mu}{\mu - \theta}$.

Next we can construct the statistical network service curve for the through flow along the H nodes. At this point we make the transition to a service curve representation with the dis-

cretization parameter τ_0 (complying with Definition 6.3). Using Eq. (6.15) and Lemma 4.4, we obtain the statistical network service curve

$$S^{net}(s, t) = S^1 * \dots * S^H(s, t - \tau_0),$$

that has (according to Eq. (6.18) from Theorem 6.8) an MGF bound with rate and prefactor given by

$$r^{net} = r_s, \quad M^{net} = \left(\frac{\mu}{\mu - \theta} e^{2\theta r_s \tau_0} \right)^H \binom{\lfloor \frac{t-s}{\tau_0} \rfloor + H - 1}{H - 1}.$$

We remark that the contribution of using Lemma 4.4 to the prefactor M^{net} from Eq. (6.18) in Theorem 6.8 is $e^{H\theta r_s \tau_0}$.

Finally, having the through flow's MGF envelope description from Eq. (7.21) and the network service curve just derived, we can invoke Theorem 6.7 and derive delay bounds. First, let us denote

$$r = r_s - r_a$$

and enforce the stability condition that

$$r > 0 \Leftrightarrow \theta < \mu(1 - \rho).$$

Then Eq. (6.11) from Theorem 6.7 gives the delay bound for all $\sigma \geq 0$

$$P\left(W^{net}(t) > \frac{\sigma}{r^{net}}\right) \leq \left(\frac{e^{2\theta r_s \tau_0}}{\theta r \tau_0} \frac{\mu}{\mu - \theta}\right)^H e^{-\theta \sigma}.$$

Optimizing the discretization parameter $\tau_0 = \frac{1}{2\theta r_s}$, as in Lemma 5.4, replacing σ with $d \cdot r^{net}$, and letting $t \rightarrow \infty$ we obtain the steady-state delay bound for all $d \geq 0$

$$P(W^{net} > d) \leq \left(e \frac{2r_s}{r} \frac{\mu}{\mu - \theta}\right)^H e^{-\theta(C - \frac{\lambda_c}{\mu - \theta})d}. \quad (7.23)$$

Rather than considering a packetized service model, we now assume that the network

treats traffic in a fluid manner. With this assumption we can view each node in the network from Figure 2.1 as a fluid server, and consequently derive the leftover service curves

$$S^h(s, t) = [C(t - s) - A_h(s, t)]_+,$$

i.e., the expression for $T^h(s, t)$ from Eq. (7.22).

To derive end-to-end delay bounds we can proceed as before, with the difference that the prefactor of the MGF bound of $S^h(s, t)$ is now $M_s = 1$, rather than $M_s = \frac{\mu}{\mu - \theta}$.

The steady-state delay bound assuming the fluid service model thus becomes

$$P(W^{net} > d) \leq \left(e \frac{2r_s}{r}\right)^H e^{-\theta(C - \frac{\lambda_c}{\mu - \theta})d}. \quad (7.24)$$

7.2.2 Correlated arrivals, independent service

Here we derive network calculus bounds for the M/M/1 network by dispensing with the statistical independence of cross arrivals, but accounting for the statistical independence of service arising from independently re-sampling the sizes of the through packets at each traversed node. In the last part we derive further bounds by also assuming a fluid service model.

As in the previous section, let a positive number θ_c such that $\theta_c < \mu - \lambda_c/C$, and denote

$$r_s(\theta_c) = C - \frac{\lambda_c}{\mu - \theta_c}.$$

Since each of the cross arrival processes $A_h(t)$ has stationary and independent increments, we can apply Theorem 6.10 and obtain that the function

$$\mathcal{T}^h(s, t) = r_s(\theta_c)(t - s)$$

is a statistical (leftover) service curve for the through flow at the h^{th} fluid server with error function $\varepsilon^s(\sigma) = e^{-\theta_c \sigma}$. Since the error function corresponding to the service curve $S^{\mu, h}$

from Lemma 7.3 is zero, we further obtain that for some $\tau_0 > 0$ the function

$$\begin{aligned} S^h(s, t) &= \mathcal{T}^h * S^{\mu, h}(s, t - \tau_0) \\ &\geq [r_s(\theta_c)(t - s) - X_f^h(t) - r_s(\theta_c)\tau_0]_+ \end{aligned}$$

is a statistical service curve for the through flow at the h^{th} node in the network from Figure 2.1 with error function $\varepsilon^s(\sigma)$ (as before, $X_f^h(t)$ denotes the fraction already processed of the packet currently in service (if any) at node h at time t). We remark that this service curve complies with Definition 6.3, by accounting for the discretization parameter τ_0 .

Next we construct the statistical network service curve $S^{net}(s, t)$ as in Eq. (6.16) for the through flow along the H nodes. The corresponding error function is given as in Eq. (4.26) from Theorem 4.7, i.e.,

$$\varepsilon^{net}(\sigma) = \underbrace{\tilde{\varepsilon}_{\delta\tau_0}^s * \dots * \tilde{\varepsilon}_{\delta\tau_0}^s}_{H-1 \text{ times}} * \varepsilon^s(\sigma),$$

where

$$\tilde{\varepsilon}_{\delta\tau_0}^s(\sigma) = \frac{1}{\delta\tau_0} \int_{\sigma}^{\infty} e^{-\theta_c u} du = \frac{1}{\theta_c \delta\tau_0} e^{-\theta_c \sigma},$$

for some $\delta > 0$. Using Lemma 5.3 we can optimize the expression of the error function as

$$\varepsilon^{net}(\sigma) = H \left(\frac{1}{\theta_c \delta\tau_0} \right)^{\frac{H-1}{H}} e^{-\frac{\theta_c}{H} \sigma}.$$

On the other hand, we have from Theorem 6.8 (more exactly Eq. (6.19)) that $S^{net}(s, t)$ has an MGF bound with rate and prefactor given by

$$\begin{aligned} r^{net} &= r_s(\theta_c) - (H - 1)\delta, \\ M^{net} &= \left(\frac{\mu}{\mu - \theta} \right)^H e^{(2H-1)\theta r_s(\theta_c)\tau_0} e^{(H-1)\theta\delta\tau_0} \left(\left\lfloor \frac{t-s}{\tau_0} \right\rfloor + H - 1 \right). \end{aligned} \quad (7.25)$$

Finally, having the through flow's rate envelope description from Eq. (7.21) and the network service curve just derived, we can invoke Theorem 6.7 and derive delay bounds. First,

let us denote

$$r = r^{net} - r_a$$

and enforce the stability condition that

$$r > 0 \Leftrightarrow \delta < \frac{1}{H-1} \left(C - \frac{\lambda}{\mu - \theta} - \frac{\lambda_c}{\mu - \theta_c} \right).$$

Then Eq. (6.11) from Theorem 6.7 gives the delay bound for all $\sigma \geq 0$

$$P \left(W^{net}(t) > \frac{\sigma}{r^{net}} \right) \leq \inf_{\sigma^a + \sigma^s = \sigma} \left\{ M' \left(\frac{1}{\theta r \tau_0} \right)^H e^{-\theta \sigma^a} + \varepsilon^{net}(\sigma^s) \right\},$$

where $M' = \left(\frac{\mu}{\mu - \theta} \right)^H e^{H(2\theta r_s(\theta_c) + \delta)\tau_0}$ (obtained by slightly relaxing the term before the binomial factor in Eq. (7.25)).

We can optimize this expression using Lemma 5.3. Then, replacing σ with $d \cdot r^{net}$, and letting $t \rightarrow \infty$ we obtain the steady-state delay bound for all $d \geq 0$

$$Pr(W^{net} > d) \leq K e^{-\frac{\theta \theta_c}{\alpha} \left(C - \frac{\lambda_c}{\mu - \theta_c} - (H-1)\delta \right) d}, \quad (7.26)$$

where

$$\begin{aligned} K &= \frac{\alpha}{\theta_c} \left(\frac{\mu}{\mu - \theta} \right)^{\frac{H\theta_c}{\alpha}} \left(\frac{He\theta_c(2r_s(\theta_c) + \delta)}{\beta r} \right)^{\frac{\beta}{\alpha}} \left(\frac{r}{\delta} \right)^{\frac{(H-1)\theta}{\alpha}} \left(\frac{\theta_c}{\theta} \right)^{\frac{\theta}{\alpha}} \\ \alpha &= H\theta + \theta_c \\ \beta &= (H-1)\theta + H\theta_c. \end{aligned}$$

Let us assume now that the network treats traffic in a fluid manner. As shown at the end of Section 7.2.1, the derivation of the corresponding bounds proceeds as before with the difference that the term $\frac{\mu}{\mu - \theta}$ is to be replaced by 1. Consequently, the steady-state delay

bound takes the form

$$Pr(W^{net} > d) \leq K e^{-\frac{\theta \theta_c}{\alpha} (C - \frac{\lambda_c}{\mu - \theta_c} - (H-1)\delta)d}, \quad (7.27)$$

where

$$K = \frac{\alpha}{\theta_c} \left(\frac{H e \theta_c (2r_s(\theta_c) + \delta)}{\beta r} \right)^{\frac{\beta}{\alpha}} \left(\frac{r}{\delta} \right)^{\frac{(H-1)\theta}{\alpha}} \left(\frac{\theta_c}{\theta} \right)^{\frac{\theta}{\alpha}},$$

and α, β are as above.

7.2.3 Independent arrivals, correlated service

Here we derive network calculus bounds for the M/M/1 network by accounting for the statistical independence of cross arrivals, but assuming identical service times for packets at the traversed nodes.

Following the steps from the previous sections, we first apply Theorem 6.6 and Lemma 5.5 and obtain that the function

$$S^h(s, t) = [C(t - s) - A_h(s, t)]_+$$

is a statistical service curve (in the sense of Definition 6.3) for the through flow at the h^{th} node with error function $\varepsilon^h(\sigma) = e^{\mu C \tau_0} e^{-\mu \sigma}$, for some $\tau_0 > 0$. Let us observe that the service curve has an MGF bound with rate and prefactor given by

$$r_s = C - r_c, \quad M_s = 1,$$

for some positive θ with $\theta < \mu - \frac{\lambda_c}{C}$.

Next we construct the statistical network service curve $S^{net}(s, t)$ as in Eq. (6.16) for the through flow along the H nodes. The corresponding error function is given as in Eq. (4.26)

from Theorem 4.7, and can be written after optimizations with Lemma 5.3 as

$$\varepsilon^{net}(\sigma) = H e^{\mu C \tau_0} \left(\frac{1}{\mu \delta \tau_0} \right)^{\frac{H-1}{H}} e^{-\frac{\mu}{H} \sigma} .$$

for some $\delta > 0$.

On the other hand, we have from Theorem 6.8 (more exactly Eq. (6.19)) that $S^{net}(s, t)$ has an MGF bound with rate and prefactor given by

$$r^{net} = r_s - (H - 1)\delta, \quad M^{net} = e^{(H-1)\theta(r_s+\delta)\tau_0} \left(\left\lfloor \frac{t-s}{\tau_0} \right\rfloor + H - 1 \right) / (H - 1) .$$

Finally, having the through flow's rate envelope description from Eq. (7.21) and the network service curve just derived, we can invoke Theorem 6.7 and derive delay bounds. First, let us denote

$$r = r^{net} - r_a$$

and enforce the stability condition that

$$r > 0 \Leftrightarrow \delta < \frac{1}{H - 1} \left(C - \frac{\lambda + \lambda_c}{\mu - \theta} \right) .$$

Then Eq. (6.11) from Theorem 6.7 yields the following delay bound for all $\sigma \geq 0$

$$P \left(W^{net}(t) > \frac{\sigma}{r^{net}} \right) \leq \inf_{\sigma^a + \sigma^s = \sigma} \left\{ \left(\frac{e^{\theta(r_s+\delta)\tau_0}}{\theta r \tau_0} \right)^H e^{-\theta \sigma^a} + \varepsilon^{net}(\sigma^s) \right\} .$$

We can optimize this expression using Lemma 5.3. Then, replacing σ with $d \cdot r^{net}$, and letting $t \rightarrow \infty$ we obtain the steady-state delay bound for all $d \geq 0$

$$Pr(W^{net} > d) \leq K e^{-\frac{\theta \mu}{\alpha} (C - \frac{\lambda_c}{\mu - \theta} - (H-1)\delta)d}, \quad (7.28)$$

where

$$\begin{aligned}
 K &= \frac{\alpha}{\mu} \left(\frac{He\mu(C + r_s + \delta)}{\beta r} \right)^{\frac{\beta}{\alpha}} \left(\frac{r}{\delta} \right)^{\frac{(H-1)\theta}{\alpha}} \left(\frac{\mu}{\theta} \right)^{\frac{\theta}{\alpha}} \\
 \alpha &= H\theta + \mu \\
 \beta &= (H-1)\theta + H\mu.
 \end{aligned}$$

7.3 Numerical Examples

We consider a similar numerical setting as in Section 5.3.2: time is measured in milliseconds and we plot bounds on the z -quantiles $w^{net}(z)$ with $z = 1 - 10^{-9}$; also, the capacity of each node is $C = 100$ Mbps and the average size of packets is $1/\mu = 400$ Bytes.

First we recall from Section 5.2.2 that the difference between exact packets' delays and the virtual delays computed with the calculus is negligible for numerical purposes.

7.3.1 The single-node case

Here we present numerical comparisons between exact results on packets delays and network calculus virtual delay bounds for the M/M/1, M/D/1, and M/M/1 queues with priorities.

In Figure 7.1.(a) we show the delay bounds from Eq. (7.6) obtained with Theorem 6.7 (based on Boole's inequality) and the exact results from Eq. (7.1) for the M/M/1 queue. The bounds from Eq. (7.8) obtained with Theorem 6.11 (based on Doob's maximal inequality) look exactly as the exact results, but they are bounds on the steady-state virtual delay. In Figure 7.1.(b) we show the delay bounds from Eqs. (7.18) and (7.19) obtained with Theorem 6.7 and Theorem 6.11, respectively, and the exact results from Eq. (7.15) for the M/D/1 queue. We show the bounds as a function of the utilization factor ρ ; the bounds closely match the exact results at utilization factors $\rho < 0.9$ and are not depicted.

For both the M/M/1 and M/D/1 queues, the bounds obtained with Theorem 6.11 improve the bounds obtained with Theorem 6.7 because they take advantage of the independent in-

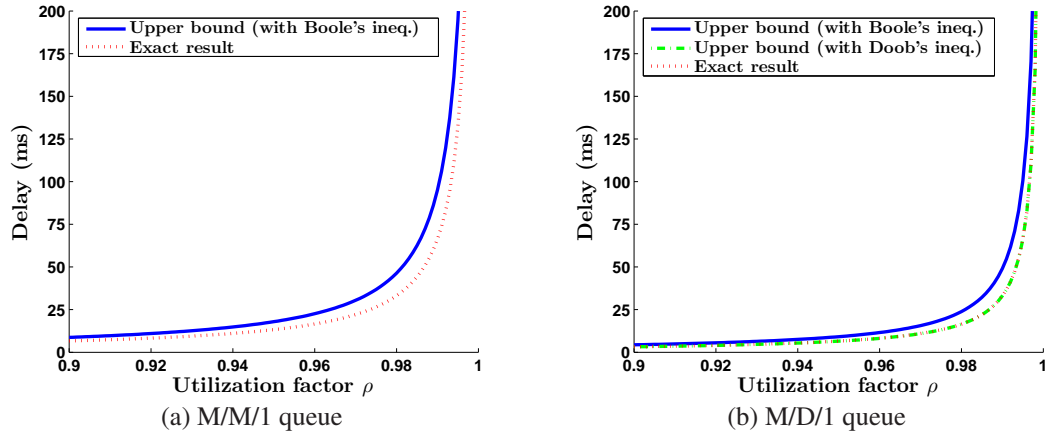


Figure 7.1: Delay $w(z)$ in M/M/1 and M/D/1 queues as a function of the utilization factor ρ ($C = 100$ Mbps, average packet size $\mu^{-1} = 400$ Bytes, $z = 1 - 10^{-9}$)

crements properties of arrival processes. The improvement becomes visible at very high utilization. This indicates that the use of Boole's inequality can lead to conservative bounds, but only at a very high utilizations.

In Figure 7.2 we show the bounds from Eq. (7.14) and the exact results from Eq. (7.9) for the expected delay of the lower priority flow in the M/M/1 queue. We show the bounds as a function of the utilization factor ρ and consider two percentages of lower priority traffic: (a) low ($p = 0.1$) and (b) very low ($p = 0.01$). We observe that in the former case the upper bounds closely match the exact results; the match becomes more pronounced when increasing the percentage of lower priority traffic. However, for very low percentages of through traffic, the upper bounds can become extremely conservative when compared to the exact results. The gap between the two vanishes when $\rho \rightarrow 1$, as pointed out at the end of Section 7.1.2.

7.3.2 The multi-node case

Here we present numerical examples illustrating three aspects of the network calculus bounds derived in Sections 7.2.1-7.2.3: (1) how conservative are the bounds when compared to exact results, (2) how do they behave when relaxing the independence assumptions of arrivals or

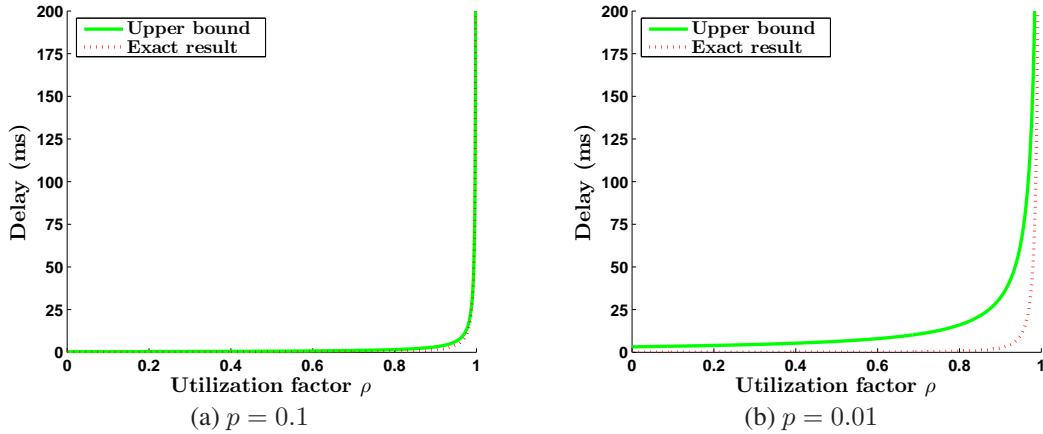


Figure 7.2: Expected delay for the lowest priority flow in an M/M/1 queue as a function of the utilization factor ρ ($C = 100$ Mbps, percentage of through traffic ($p = 0.1$ and $p = 0.01$), average packet size $\mu^{-1} = 400$ Bytes)

service, and (3) how do they behave when neglecting packetization.

In Figure 7.3 we show the bounds obtained using the statistical assumptions on both arrivals and service (Eq. (7.23)), and the exact results, as a function of the number of nodes. We consider both a moderate utilization factor ($\rho = 0.75$) in Figure 7.3.(a), and a high utilization factor ($\rho = 0.9$) in Figure 7.3.(b). There are three different plots for the upper bounds, each reflecting a different percentage p of through traffic, out of the total traffic mix: $p = 0.1$, $p = 0.5$, and $p = 0.9$, respectively (note that the exact results do not depend of p). We observe that for moderate utilization factors and high percentages of through traffic the calculus bounds are reasonably accurate. However, when the amount of cross traffic is increased, the calculus bounds become pessimistic. This indicates that the representation of the service for the through flow with leftover service curves may be too conservative (this observation is also supported by the corresponding discussion from Section 7.1.2). At higher utilizations, the network calculus bounds become pessimistic even for high amounts of through traffic; also, the bounds for higher amounts of cross traffic degrade more rapidly than in the case of lower utilizations.

Similar conclusions can be drawn from Figure 7.4, which shows the bounds as a function of the utilization factor ρ for two cases: (a) small number of nodes ($H = 5$), and (b) large

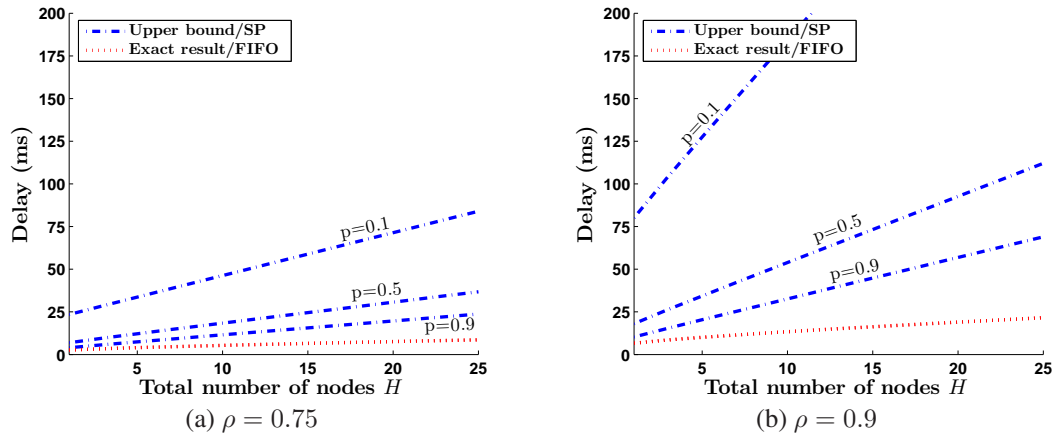


Figure 7.3: Comparison of network calculus bounds with exact results in an M/M/1 network. End-to-end delay $w^{net}(z)$ as a function of the number of nodes H ($C = 100$ Mbps, utilization factor ($\rho = 0.75$ and $\rho = 0.9$), percentage of through traffic ($p = 0.1$, $p = 0.5$, and $p = 0.9$), average packet size $\mu^{-1} = 400$ Bytes, $z = 1 - 10^{-9}$)

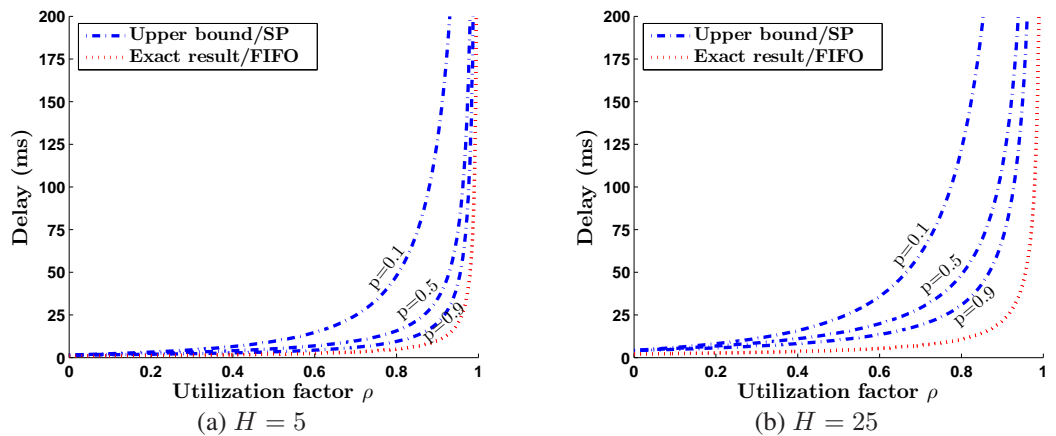


Figure 7.4: Comparison of network calculus bounds with exact results in an M/M/1 network. End-to-end delay $w^{net}(z)$ as a function of the utilization factor ρ (number of nodes ($H = 5$ and $H = 25$), $C = 100$ Mbps, percentage of through traffic ($p = 0.1$, $p = 0.5$, and $p = 0.9$), average packet size $\mu^{-1} = 400$ Bytes, $z = 1 - 10^{-9}$)

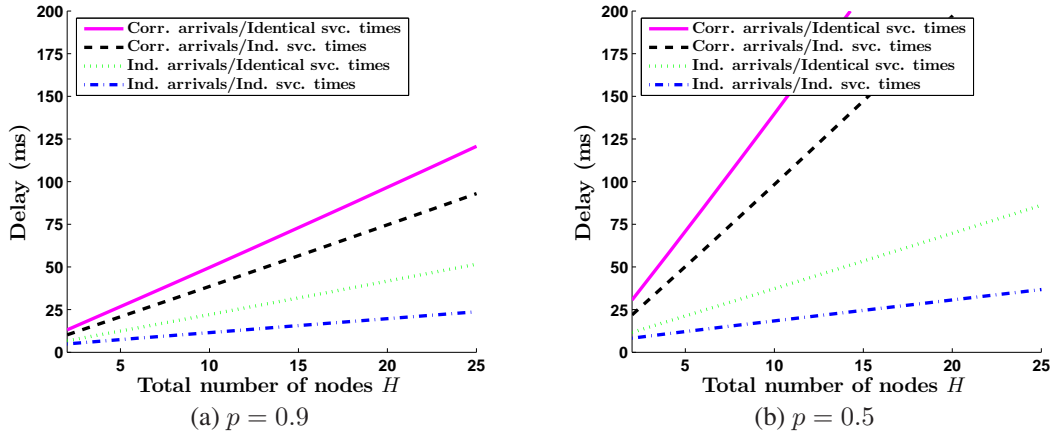


Figure 7.5: The impact of relaxing the statistical independence assumptions of arrivals and service in an M/M/1 network. End-to-end delay $w^{net}(z)$ as a function of the number of nodes H ($C = 100$ Mbps, utilization factor $\rho = 0.75$, percentage of through traffic ($p = 0.9$ and $p = 0.5$), average packet size $\mu^{-1} = 400$ Bytes, $z = 1 - 10^{-9}$)

number of nodes ($H = 25$). The calculus bounds are quite accurate for small number of nodes, and a low percentage of through traffic. However, increasing the number of nodes or the amount of cross traffic leads to much more pessimistic bounds.

In Figure 7.5 we illustrate the behavior of calculus bounds by relaxing the statistical independence assumptions of arrivals and service. We plot the end-to-end delay bounds as a function of the number of nodes H and consider two cases: (a) large amount of through traffic ($p = 0.9$), and (b) medium amount of through traffic ($p = 0.5$). The plots correspond to Eqs. (7.23), (7.28), (7.26), and (5.29), respectively, in an increasing order of the bounds. Both figures show that dispensing with the independence of service has a similar effect on the bounds, in both cases of independent or correlated arrivals. Dispensing with the independence assumption of arrivals has a much more noticeable effect in Figure 7.5.(b), due to the increase in the amount of cross traffic. The bounds obtained for correlated arrivals but independent service appear to be more pessimistic than the bounds obtained for independent arrivals but correlated service, i.e., correlations within arrivals have a more noticeable effect than correlations within service.

Similar conclusions can be drawn from Figure 7.6, that is similar to Figure 7.5, but it

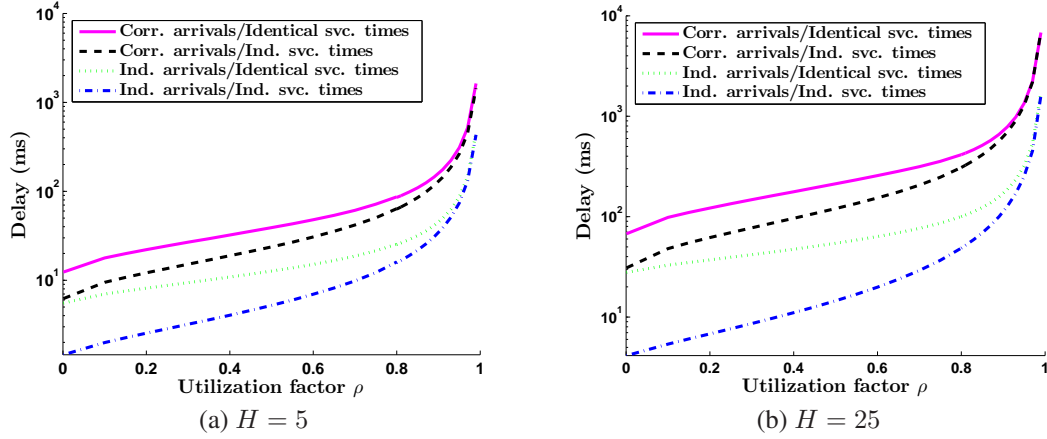


Figure 7.6: The impact of relaxing the statistical independence assumptions of arrivals and service in an M/M/1 network. End-to-end delay $w^{net}(z)$ as a function of the utilization factor ρ (number of nodes ($H = 5$ and $H = 25$), $C = 100$ Mbps, percentage of through traffic ($p = 0.5$), average packet size $\mu^{-1} = 400$ Bytes, $z = 1 - 10^{-9}$)




Figure 7.7: The impact of using a fluid service model in an M/M/1 network. End-to-end delay $w^{net}(z)$ as a function of the utilization factor ρ (number of nodes ($H = 5$ and $H = 25$), $C = 100$ Mbps, percentage of through traffic ($p = 0.5$), average packet size $\mu^{-1} = 400$ Bytes, $z = 1 - 10^{-9}$)

shows the bounds (on a logarithmic scale) as a function of the utilization factor ρ for two cases: (a) small number of nodes ($H = 5$), and (b) large number of nodes ($H = 25$). Also, in both cases (a) and (b) we let an equal share of through and cross traffic ($p = 0.5$). Remarkably, the figures show that for both independent and correlated arrivals, the assumption of identical service times for packets is justified at high utilizations. This observation is first pointed out in the context of queueing networks theory (for independent arrivals only) by means of simulations, whereas we reach it by means of analytical and numerical results.

Last, Figure 7.7 illustrates the effects of dispensing with the packetized service model at the nodes. We consider a small number of nodes ($H = 5$) in Figure 7.7.(a) and a higher number of nodes ($H = 25$) in Figure 7.7.(b). In both figures we consider both correlated and independent arrivals, and plot the bounds as a function of the utilization factor. For the case of correlated arrivals, the bounds obtained using the two service models (packetized and fluid) closely match in all situations. A similar behavior is observed for independent arrivals, with the difference that the fluid model predicts more optimistic bounds than the packetized model, but only at very low utilization factors. Based on the figures, we can conclude that using a fluid flow service model is generally justified at high data rates.

Chapter 8

Conclusions and Future Work

8.1 Conclusions

We have made theoretical contributions in the stochastic network calculus, and have provided new analytical insights into the scaling behavior of network delays.

1. A stochastic network calculus formulation with a new statistical network service curve: We have formulated a stochastic network calculus that is generally suitable to analyze network scenarios where arrivals and service at the nodes may be statistically correlated. We have provided mathematical models to characterize arrivals and service, and have also presented analytical results for the derivation of single-node and multi-node bounds on network performance metrics such as backlog or delay.

The main technical contribution of this calculus formulation, and of this thesis in general, is the construction of a new statistical network service curve. The network service curve provides a service description for a flow in a network, as though the flow had traversed a single-node only. The network service curve allows the computation of multi-node performance bounds using single-node results. The proposed statistical network service curve can be applied to a wide class of arrivals, and lends itself to explicit numerical results. We point out that the formulation of such network service curves was a long-standing research problem.

2. *The existence of $\Theta(H \log H)$ scaling of network performance metrics:* We have applied our statistical network service curve in a network scenario with EBB arrivals and service, and have shown that it yields end-to-end performance bounds which grow as $\mathcal{O}(H \log H)$. This scaling behavior is much smaller than the corresponding $\mathcal{O}(H^3)$ scaling behavior of end-to-end bounds obtained by using the alternative method of adding per-node bounds proposed in the early 1990s.

For a particular network scenario with EBB arrivals and service, we have also shown the $\Omega(H \log H)$ scaling of network performance bounds. The resulting $\Theta(H \log H)$ result demonstrates a different scaling behavior of network performance bounds than is currently predicted with existing analytical tools. For example, queueing networks theory predicts a linear scaling behavior, i.e., $\Theta(H)$, by relying on simplifying assumptions of the statistical independence of arrivals and service at the nodes.

3. *A stochastic network calculus formulation accounting for statistical independence:* We have formulated a stochastic network calculus that can account for the properties of statistical independence of arrivals or service at the nodes, where available. Unlike other formulations of the calculus which require the independence of both arrivals and service, our formulation also allows for correlations between either arrivals or service at different nodes. A scenario where the proposed network calculus formulation is particularly useful is a network with statistically independent cross traffic, but statistically correlated service at the nodes. These correlations are generally inherent in packet networks since each packet maintains its size constant at each of the traversed nodes.

We have specialized the second calculus formulation to the class of arrival processes characterized by stationary and independent increments properties. Using these additional properties, we have applied supermartingales based techniques in network calculus. The benefit of using these techniques is that we could improve existing bounds in the network calculus, especially at high utilizations of the nodes.

4. *Relationship with existing theories:* We have investigated the accuracy of stochastic

network calculus bounds by comparing them with exact results available in product-form queueing networks. In the single-node case, we have shown for M/M/1, M/D/1, and M/M/1 queues with priorities, that the network calculus bounds are quite accurate. In the multi-node case, we have shown that the network calculus bounds can become accurate in M/M/1 networks where the amount of cross traffic is low. We have quantified the impact of accounting for statistical independence in the multi-node analysis with the calculus, and have also shown that fluid service models provide good approximations for packetized service models at high utilizations.

8.2 Future Work

Here we outline some open problems in the stochastic network calculus, and possible applications of the calculus to networking problems.

A problem of interest in the stochastic network calculus theory concerns the existence of ‘convolution-form’ networks, i.e., the class of networks in which the service given to flows can be expressed in terms of $(\min, +)$ convolution formulas. Networks with cross traffic, fixed routing, and arrivals described with statistical envelopes satisfying certain integrability conditions have been shown in this thesis to have a convolution-form. It is an open problem whether networks with arbitrary topologies, possibly containing cycles, probabilistic routing, and more general classes of arrivals can also be described with convolution formulas. A potential insight into this direction may be offered by the results of Wischik [114] who showed that under certain asymptotic regimes, the per-flow effective bandwidths are preserved at the output of a link, i.e., the characteristics of traffic are preserved within networks.

The aim of seeking a general result on the existence of convolution-form networks is to enable the performance analysis of networks in a simplified manner, as it is currently the case for product-form queueing networks. Convolution-form networks may exist for a much larger class of arrival processes than the class of Poisson processes, which represents a restriction for product-form networks.

Another theoretical problem concerns the formulation of a stochastic network calculus which can enable the derivation of backlog and delay bounds for the class of heavy-tailed arrivals processes. Such processes present interest since they have been shown to model traffic in networks (see Park and Willinger [92]). We have pointed out that the calculus formulation from Chapter 6 does not apply to heavy-tailed processes, because they have unbounded moment generating functions. Also, the calculus formulation from Chapter 4 may not apply to heavy-tailed processes in some scenarios, such as networks with heavy-tailed cross traffic; the reason is that the construction of leftover service curves, followed by the derivation of network service curve, would result in unbounded error functions.

In Chapter 6 we have shown that the stochastic network calculus can exploit the properties of independent increments of arrival processes in the single-node case. It is currently open whether similar results hold in the multi-node case, which may significantly improve the delay bounds obtained in Chapter 7 for M/M/1 networks.

A possible application of the stochastic network calculus is to the computation of per-user throughput capacity and delay bounds in multi-hop wireless networks. There are currently many results available in this directions concerning access mechanisms such as ALOHA (see Silvester and Kleinrock [102]) or the distributed coordination function (DCF) characteristic to 802.11 wireless networks (see Bianchi [11]). The potential benefit of using the calculus is that it can yield bounds on both the capacity and delay distributions. Also, the calculus can compute delay bounds for a wide class of arrival traffic; existing results concerning average delays are available only for specific arrivals, such as Bernoulli or infinite arrival models.

The key challenge in applying the calculus to the analysis of wireless networks is capturing with statistical service curve the noise induced by collisions and the different power levels used for transmissions. Possible insights into this problem may be offered by related works of Wu and Negi [115], Jiang and Emstad [61], and Fidler [49]. Having available the service curves within cells, i.e., areas where nodes can hear each other, end-to-end results on capacity or delay can be then obtained with the convolution theorem.

Glossary of Notation

Notation	Description	Page
D-BIND	Deterministic Bounding INterval-length Dependent envelope model	19
EBB	Exponentially Bounded Burstiness envelope model	36
EDF	Earliest Deadline First scheduling algorithm	43
FBM	Fractional Brownian Motion	6
FIFO	First In First Out scheduling algorithm	23
GPS	Generalized Processor Sharing scheduling algorithm	51
gSBB	generalized Stochastically Bounded Burstiness envelope model	40
LRD	Long range dependence	5
LTI	Linear Time-Invariant system	20
MGF	Moment Generating Function	45
MPEG	Moving Picture Experts Group video compression standard	40
SBB	Stochastically Bounded Burstiness envelope model	37
SCED	Service Curve-based Earliest Deadline First scheduling algorithm	68
SP	Static Priority scheduling algorithm	23
$(\min, +)$	$(\min, +)$ algebra	20

Notation	Description	Page
---^a	Superscript referring to arrivals	78
---_a	Subscript referring to arrivals	91
---^s	Superscript referring to service	78
---_s	Subscript referring to service	91
---_c	Subscript referring to cross traffic	24
---_h	Subscript referring to cross traffic at node h	29
---^h	Superscript referring to through traffic at node h	29
---^{net}	Superscript referring to through traffic across the network	29
$A(t), A(s, t)$	Arrival process	16
$D(t), D(s, t)$	Departure process	16
$B(t)$	Backlog process	16
$W(t)$	Delay process	17
$w(z)$	z -quantile	90
$\mathcal{G}(t)$	Envelope function (deterministic or statistical)	17
$G(t), G(s, t)$	Statistical envelope as random process	44
$\mathcal{G}(t, \varepsilon)$	Effective envelope function	39
$\mathcal{G}(t, \beta, \varepsilon)$	Global effective envelope function	41
$\mathcal{S}(t)$	Service curve function (deterministic or statistical)	21
$\mathcal{R}(t)$	Constant-rate service curve function (deterministic or statistical)	22
$S(t), S(s, t)$	Service curve as random process	52
ε	Violation probability value	39
z	Violation probability value (usually $z = 1 - \varepsilon$)	90

Notation	Description	Page
$\varepsilon(\sigma)$	Error function for statistical envelope or service curve	78
$\tilde{\varepsilon}_a(\sigma)$	The integral of an error function: $\frac{1}{a} \int_{\sigma}^{\infty} \varepsilon(u) du$	73
u, s, t	Time indexes	17
\underline{t}	The start of the busy period containing t	49
C	Capacity at the node	15
H	Number of nodes in the network	15
h	Node index	29
r	Traffic rate	18
P	Traffic peak rate	33
δ	Relaxation rate (e.g. $\mathcal{S}_{-\delta}(t) = \mathcal{S}(t) - \delta t$)	77
σ	Traffic burst	18
$d, d(\sigma)$	Delays	98
θ	Exponential decay rate in the violation probability; space parameter in effective bandwidth	36
M	Prefactor in exponentially decaying violation probability functions	36
K, α, β, γ	Constants	108
λ, μ	Transition rates between the states of a Markov-modulated On-Off process	118
λ	Arrival rates of packets (through flows)	101
μ	Service rate of packets	101
$X_f(t)$	Processed fraction of a packet	103
$\alpha_A(\theta, t)$	Effective bandwidth of an arrival flow $A(t)$	36
$B(t, p)$	Binomial random variable	44

Notation	Description	Page
$E[X]$	Expectation of a random variable	36
$E[X \parallel \mathcal{F}]$	Conditional expectation of X with respect to the σ -algebra \mathcal{F}	113
$\Gamma(x, H)$	Gamma distribution	153
I	Indicator function	44
$Var[X]$	Variance of a random variable X	39
$f \approx g$	Approximative pointwise equality of two functions	34
$f \triangleq g$	Equality by definition	16
$*$	$(min, +)$ convolution operator	21, 52
$*_t$	$(min, +)$ modified convolution operator	48
\oslash	$(min, +)$ deconvolution operator	25
$f(t) = \mathcal{O}(g(t))$	Landau Big-Oh asymptotic notation	28
$f(t) = \Omega(g(t))$	Landau Big-Omega asymptotic notation	28
$f(t) = \Theta(g(t))$	Landau Big-Theta asymptotic notation	28
e	Euler's constant: $e = 2.718281 \dots$	96
$[x]_+$	Positive part $\max\{x, 0\}$ of a number x	22
$\log(x)$	Natural logarithm of x	36
\square	Halmos symbol to end a proof	75

Index

- Busy period, 23
 - of a flow, 23
- Convolution, 21
 - of doubly-indexed functions, 52
- Deterministic envelope, 17
 - IntServ (example), 19
 - leaky-bucket (example), 18
- Deterministic service curve, 21
 - constant-rate (example), 22
 - latency-rate (example), 23
 - leftover, 24
- EBB, *see* Statistical envelope
- Effective bandwidth, 36
- Effective envelope, *see* Statistical envelope
- Effective service curve, *see* Statistical service curve
- Envelope, *see* Deterministic envelope, Statistical envelope
- Expectation, 36
- Filtration, 113
- Function
 - indicator, 44
 - sub-additive, 18
- Global effective envelope, *see* Statistical envelope
- gSBB, *see* Statistical envelope
- Inequality
 - Boole, 57
 - Chernoff, 57
 - Doob (maximal), 114
- Landau notation, 29
- Leaky-bucket, *see* Deterministic envelope
- Network
 - tandem, 105
 - with cross traffic, 15
- Network service curve, 27
 - Deterministic, 27
 - Statistical, 69
- Packetizer, 103
- Pay-bursts-only-once, 32
- Process, 16
 - Arrivals, 16

- Backlog, 16
- Delay, 17
- Departures, 16
- SBB, *see* Statistical envelope
- Scheduling, 23
 - earliest eadline first, 43
 - FIFO, 24
 - generalized processor sharing, 52
 - locally-FIFO, 24
 - static priority, 23
 - workconserving, 24
- Statistical envelope, 36
 - as non-random function, 36
 - EBB, 36
 - effective envelope, 39
 - global effective envelope, 41
 - gSBB, 40
 - SBB, 37
 - as random process, 43
 - with *a.s.* ordering, 45
 - with stochastic ordering, 44
- Statistical service curve, 47
 - as non-random function, 47
 - adaptive, 48
 - EBB, 47
 - effective service curve, 47
 - leftover, 50
 - network, 82
 - as random process, 52
 - leftover, 53
 - network, 69
 - Traffic source (examples)
 - Bernoulli, 33
 - Compound Poisson, 102
 - FBM, 73
 - Markov-modulated On-Off, 119
 - Variance, 39
 - z-quantile, 90

Bibliography

- [1] R. Agrawal, R. L. Cruz, C. Okino, and R. Rajan. Performance bounds for flow control protocols. *IEEE/ACM Transactions on Networking*, 7(3):310–323, June 1999.
- [2] D. Anick, D. Mitra, and M.M. Sondhi. Stochastic theory of a data-handling system with multiple sources. *Bell Systems Technical Journal*, 61(8):1871–1894, October 1982.
- [3] D. Artiges and P. Nain. Upper and lower bounds for the multiplexing of multiclass Markovian on/off sources. *Performance Evaluation*, 27-28(4):673–698, October 1996.
- [4] S. Ayyorgun and R. Cruz. A service-curve model with loss and a multiplexing problem. In *Proceedings of the 24th IEEE International Conference on Distributed Computing System (ICDCS)*, pages 756–765, March 2004.
- [5] S. Ayyorgun and W. Feng. A probabilistic definition of burstiness characterization. Technical Report LA-UR 03-3668, Los Alamos National Laboratory, May 2003.
- [6] S. Ayyorgun and W. Feng. A systematic approach for providing end-to-end probabilistic QoS guarantees. In *Proceedings of the International Conference On Computer Communications and Networks (ICCCN)*, pages 115–122, October 2004.
- [7] F. L. Baccelli, G. Cohen, G. J. Olsder, and J.-P. Quadrat. *Synchronization and Linearity: An Algebra for Discrete Event Systems*. John Wiley and Sons, 1992.

- [8] F. Baskett, K. M. Chandy, R. R. Muntz, and F. G. Palacios. Open, closed and mixed networks of queues with different classes of customers. *Journal of the ACM*, 22(2):248–260, April 1975.
- [9] J. Beran, R. Sherman, M. S. Taqqu, and W. Willinger. Long-range-dependence in variable-bit-rate video traffic. *IEEE Transactions on Communications*, 43(2/3/4):1566–1579, February/March/April 1995.
- [10] D. Bertsekas and R. Gallager. *Data Networks*. Prentice Hall, 1992.
- [11] G. Bianchi. Performance analysis of the IEEE 802.11 distributed coordination function. *IEEE Journal on Selected Areas in Communications*, 18(3):535–547, March 2000.
- [12] R. Boorstyn, A. Burchard, J. Liebeherr, and C. Ottamakorn. Effective envelopes: Statistical bounds on multiplexed traffic in packet networks. In *Proceedings of IEEE Infocom 2000*, pages 1223–1232, Tel Aviv, March 2000.
- [13] R. Boorstyn, A. Burchard, J. Liebeherr, and C. Ottamakorn. Statistical service assurances for traffic scheduling algorithms. *IEEE Journal on Selected Areas in Communications Special Issue on Internet QoS*, 18(12):2651–2664, December 2000.
- [14] J.-Y. Le Boudec. Application of network calculus to guaranteed service networks. *IEEE/ACM Transactions on Information Theory*, 44(3):1087–1097, May 1998.
- [15] J.-Y. Le Boudec. Some properties of variable length packet shapers. *IEEE/ACM Transactions on Networking*, 10(3):329–337, June 2002.
- [16] J.-Y. Le Boudec and P. Thiran. *Network Calculus*. Springer Verlag, Lecture Notes in Computer Science, LNCS 2050, 2001.
- [17] O. Boxma. On a tandem queueing model with identical service times at both counters. part 1,2. *Advances in Applied Probability*, 11(3):616–659, 1979.

- [18] R. Braden, D. Clark, and S. Shenker. Integrated services in the Internet architecture: an overview. IETF RFC 1633, July 1994.
- [19] E. Brockmeyer, H.L. Hallstrom, and A. Jensen. The life and works of A. K. Erlang. *Transactions of the Danish Academy of Technical Sciences*, (2), 1948.
- [20] E. Buffet and N. G. Duffield. Exponential upper bounds via martingales for multiplexers with Markovian arrivals. *Journal of Applied Probability*, 31(4):1049–1060, December 1994.
- [21] A. Burchard, J. Liebeherr, and F. Ciucu. On $\Theta(H \log H)$ scaling of network delays. In *Proceedings of IEEE Infocom*, May 2007.
- [22] A. Burchard, J. Liebeherr, and S. D. Patek. A calculus for end-to-end statistical service guarantees. Technical Report CS-2001-19, University of Virginia, Computer Science Department, May 2002.
- [23] A. Burchard, J. Liebeherr, and S. D. Patek. A min-plus calculus for end-to-end statistical service guarantees. *IEEE Transactions on Information Theory*, 52(9):4105 – 4114, September 2006.
- [24] S. B. Calo. Delay properties of message channels. In *Proc. IEEE ICC, Boston, Mass.*, pages 43.5.1–43.5.4, 1979.
- [25] C.-S. Chang. Stability, queue length and delay, Part II: Stochastic queueing networks. In *Proceedings of the 31st IEEE Conference on Decision and Control*, pages 1005–1010, December 1992.
- [26] C.-S. Chang. Stability, queue length, and delay of deterministic and stochastic queueing networks. *IEEE Transactions on Automatic Control*, 39(5):913–931, May 1994.
- [27] C.-S. Chang. On the exponentiality of stochastic linear systems under the max-plus algebra. *IEEE Transactions on Automatic Control*, 41(8):1182–1188, August 1996.

- [28] C.-S. Chang. On deterministic traffic regulation and service guarantees: a systematic approach by filtering. *IEEE/ACM Transactions on Information Theory*, 44(3):1097–1110, May 1998.
- [29] C.-S. Chang. *Performance Guarantees in Communication Networks*. Springer Verlag, 2000.
- [30] C.-S. Chang, Y. Chiu, and W. T. Song. On the performance of multiplexing independent regulated inputs. In *Proceedings of ACM Sigmetrics*, pages 184–193, June 2001.
- [31] J. Choe and N. B. Shroff. A central-limit-theorem-based approach for analyzing queue behavior in high-speed network. *IEEE/ACM Transactions on Networking*, 6(5):659–671, October 1998.
- [32] G. Choudhury, D. Lucantoni, and W. Whitt. Squeezing the most out of ATM. *IEEE Transactions on Communications*, 44(2):203–217, February 1996.
- [33] F. Ciucu, A. Burchard, and J. Liebeherr. Scaling properties of statistical end-to-end bounds in the network calculus. *IEEE Transactions on Information Theory*, 52(6):2300–2312, June 2006.
- [34] C. Courcoubetis, V. A. Siris, and G. D. Stamoulis. Application and evaluation of large deviation techniques for traffic engineering in broadband networks. In *Proceedings of ACM Sigmetrics*, pages 212–221, June 1998.
- [35] Mark E. Crovella and Azer Bestavros. Self-similarity in world wide web traffic: Evidence and possible causes. In *SIGMETRICS Conf. Measurement & Modeling of Comput. Syst.*, pages 160–169, May 1996.
- [36] R. Cruz. A calculus for network delay, parts I and II. *IEEE Transactions on Information Theory*, 37(1):114–141, January 1991.

- [37] R. L. Cruz. A calculus for network delay, Part I : Network elements in isolation. *IEEE Transactions of Information Theory*, 37(1):114–121, January 1991.
- [38] R. L. Cruz. A calculus for network delay, Part II : Network analysis. *IEEE Transactions on Information Theory*, 37(1):132–141, January 1991.
- [39] R. L. Cruz. Quality of service guarantees in virtual circuit switched networks. *IEEE Journal on Selected Areas in Communications*, 13(6):1048–1056, August 1995.
- [40] R. L. Cruz. Quality of service management in integrated services networks. In *1st Semi-Annual Research Review, CWC, University of California at San Diego*, June 1996.
- [41] R. L. Cruz and C. Okino. Service gurantees for window flow control. In *Proceedings of the 34th Allerton Conference on Communications, Control and Computating*, October 1996.
- [42] J. N. Daigle and J. D. Langford. Models for analysis of packet voice communications systems. *IEEE Journal on Selected Areas in Communications*, 4(6):847–855, September 1986.
- [43] N. G. Duffield and N. O’Connell. Large deviations and overflow probabilities for the general single-server queue, with applications. *Mathematical Proceedings of the Cambridge Philosophical Society*, 118(2):363–374, September 1995.
- [44] A. Elwalid and D. Mitra. Design of generalized processor sharing schedulers which statistically multiplex heterogeneous QoS classes. In *Proceedings of IEEE Infocom*, pages 1220–1230, March 1999.
- [45] A. Elwalid, D. Mitra, and R. Wentworth. A new approach for allocating buffers and bandwidth to heterogeneous, regulated traffic in an ATM node. *IEEE Journal on Selected Areas in Communications*, 13(6):1115–1127, August 1995.

- [46] A.K. Erlang. The theory of probabilities and telephone conversations. *Nyt Tidsskrift for Matematik B*, 20:33–39, 1909 (in Danish).
- [47] A.K. Erlang. Solution of some problems in the theory of probabilities of significance in automatic telephone exchanges. *Elektroteknikerne*, 13, 1917 (in Danish).
- [48] M. Fidler. An end-to-end probabilistic network calculus with moment generating functions. In *IEEE 14th International Workshop on Quality of Service (IWQoS)*, pages 261–270, June 2006.
- [49] M. Fidler. A network calculus approach to probabilistic quality of service analysis of fading channels. In *Globecom*, November 2006.
- [50] P. Le Gall. The overall sojourn time in tandem queues with identical successive service times and renewal input. *Stochastic Processes and Their Applications*, 52(1):165–178, August 1994.
- [51] R. J. Gibbens and P. J. Hunt. Effective bandwidths for the multi-type UAS channel. *Queueing Systems*, 9(1-2):17–28, September 1991.
- [52] G. Grimmett and D. Stirzaker. *Probability and Random Processes*. Oxford University Press, 2001.
- [53] D. Gross and C.M. Harris. *Fundamentals of Queueing Theory (2nd Edition)*. John Wiley and Sons, 1985.
- [54] R. Guérin, H. Ahmadi, and M. Naghshineh. Equivalent capacity and its application to bandwidth allocation in high-speed networks. *IEEE Journal on Selected Areas in Communications*, 9(7):968–981, September 1991.
- [55] H. Heffes and D. M. Lucantoni. A Markov modulated characterization of packetized voice and data traffic and related statistical multiplexer performance. *IEEE Journal on Selected Areas in Communications*, 4(6):856–867, September 1986.

- [56] J.Y. Hui. Resource allocation for broadband networks. *IEEE Journal of Selected Areas in Communications*, 6(9):1598–1608, December 1988.
- [57] V.B. Iversen and L. Staalhagen. Waiting time distribution in M/D/1 queueing systems. *Electronics Letters*, 35(25):2184–2185, December 1999.
- [58] J.R. Jackson. Networks of waiting lines. *Operations Research*, 5:518–521, August 1957.
- [59] D. Jagerman, B. Melamed, and W. Willinger. Stochastic modeling of traffic processes. In J. Dshalalow, editor, *Frontiers in Queueing: Models, Methods and Problems*. CRC Press, 1996.
- [60] Y. Jiang. A basic stochastic network calculus. In *ACM Sigcomm*, pages 123–134, September 2006.
- [61] Y. Jiang and P. J. Emstad. Analysis of stochastic service guarantees in communication networks: A server model. In *Proceedings of the International Workshop on Quality of Service (IWQoS)*, pages 233–245, June 2005.
- [62] Y. Jiang and P. J. Emstad. Analysis of stochastic service guarantees in communication networks: A traffic model. In *19th International Teletraffic Congress (ITC)*, August 2005.
- [63] F. P. Kelly. Networks of queues with customers of different types. *Journal of Applied Probability*, 3(12):542–554, September 1975.
- [64] F. P. Kelly. Notes on effective bandwidths. In *Stochastic Networks: Theory and Applications*. (Editors: F.P. Kelly, S. Zachary and I.B. Ziedins) Royal Statistical Society Lecture Notes Series, 4, pages 141–168. Oxford University Press, 1996.

- [65] G. Kesidis and T. Konstantopoulos. Worst-case performance of a buffer with independent shaped arrival processes. *IEEE Communications Letters*, 4(1):26–28, January 2000.
- [66] J. F. C. Kingman. A martingale inequality in the theory of queues. *Cambridge Philosophical Society*, 59:359–361, 1964.
- [67] L. Kleinrock. *Communication Nets : Stochastic Message Flow and Delay*. McGraw-Hill, Inc., 1964.
- [68] L. Kleinrock. *Queueing Systems*, volume 1. John Wiley and Sons, 1975.
- [69] E. W. Knightly. H-BIND: A new approach to providing statistical performance guarantees to VBR traffic. In *Proceedings of IEEE Infocom*, pages 1091–1099, March 1996.
- [70] E. W. Knightly. Second moment resource allocation in multi-service networks. In *Proceedings of ACM Sigmetrics*, pages 181–191, June 1997.
- [71] E. W. Knightly. Enforceable quality of service guarantees for bursty traffic streams. In *Proceedings of IEEE Infocom*, pages 635–642, March 1998.
- [72] E. W. Knightly and H. Zhang. D-BIND: an accurate traffic model for providing QoS guarantees to VBR traffic. *IEEE/ACM Transactions on Networking*, 5(2):219–231, April 1997.
- [73] J. Kurose. On computing per-session performance bounds in high-speed multi-hop computer networks. In *Proceedings of ACM Sigmetrics*, pages 128–139, June 1992.
- [74] E. A. Lee and P. Varaiya. *Structure and Interpretation of Signals and Systems*. Addison Wesley, 2002.
- [75] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson. On the self-similar nature of ethernet traffic. *IEEE/ACM Transactions on Networking*, 2(1):1–15, February 1994.

- [76] C. Li, A. Burchard, and J. Liebeherr. A network calculus with effective bandwidth. Technical Report CS-2003-20, University of Virginia, November 2003.
- [77] J. Liebeherr. A note on statistical multiplexing and scheduling in video networks at high data rates. In *The Internet as a Large-Scale Complex System* (Eds: K. Park, W. Willinger), pages 179–201. Oxford University Press, 2005.
- [78] J. Liebeherr, A. Burchard, and S. Patek. A stochastic networks calculus: theory and tools for the analysis of future networks. *Proposal to the National Science Foundation*, April 2004.
- [79] J. Liebeherr, M. Fidler, and S. Valaee. A min-plus system interpretation of bandwidth estimation. In *Proceedings of IEEE Infocom*, May 2007.
- [80] J. Liebeherr, S. Patek, and A. Burchard. Statistical per-flow service bounds in a network with aggregate provisioning. In *Proceedings of IEEE Infocom*, March 2003.
- [81] J. Liebeherr, D. Wrege, and D. Ferrari. Exact admission control for networks with bounded delay services. *IEEE/ACM Transactions on Networking*, 4(6):885–901, December 1996.
- [82] Y. Liu, C.K. Tham, and Y. Jiang. A stochastic network calculus. Technical report, ECE-CCN-0301, Dept. of Electrical and Computer Engineering, National University of Singapore, November 2004.
- [83] Z. Liu, P. Nain, and D. Towsley. Exponential bounds with applications to call admission. *Journal of the ACM*, 44(3):366–394, May 1997.
- [84] R. M. Loynes. The stability of a queue with non-independent inter-arrival and service times. *Proceedings of the Cambridge Philosophical Society*, 58:497–520, 1962.

- [85] B. Maglaris, D. Anastassiou, P. Sen, G. Karlsson, and J.D. Robbins. Performance models of statistical multiplexing in packet video communications. *IEEE Transactions on Communications*, 36(7):834, July 1988.
- [86] L. Massoulié and A. Simonian. Large buffer asymptotics for the queue with fractional Brownian input. *Journal of Applied Probability*, 36(3):894–906, September 1999.
- [87] S. McCreary and K. Claffy. Trends in wide area IP traffic patterns. In *Proceedings of 13th ITC Specialist Seminar on Internet Traffic Measurement and Modeling*, September 2000.
- [88] T. Mikosch, S. Resnick, H. Rootzn, and A. Stegeman. Is network traffic approximated by stable Lévy motion or fractional Brownian motion? *Annals of Applied Probability*, 12(1):23–68, February 2002.
- [89] I. Norros. On the use of fractional Brownian motion in the theory of connectionless networks. *IEEE Journal on Selected Areas in Communications*, 13(6):953–962, August 1995.
- [90] A. Parekh and R. Gallager. A generalized processor sharing approach to flow control: The single-node case. *IEEE/ACM Transactions on Networking*, 1(3):344–357, June 1993.
- [91] A. K. Parekh and R. G. Gallager. A generalized processor sharing approach to flow control in integrated services networks: The multiple node case. *IEEE/ACM Transactions on Networking*, 2(2):137–150, April 1994.
- [92] K. Park and W. Willinger, editors. *Self-Similar Network Traffic and Performance Evaluation*. John Wiley & Sons, 2000.
- [93] V. Paxson and S. Floyd. Wide-area traffic: The failure of Poisson modelling. *IEEE/ACM Transactions on Networking*, 3(3):226–244, June 1995.

- [94] F. Lo Presti, Z.-L. Zhang, J. Kurose, and D. Towsley. Source time scale and optimal buffer/bandwidth tradeoff for heterogeneous regulated traffic in a network node. *IEEE/ACM Transactions on Networking*, 7(4):490–501, August 1999.
- [95] J. Qiu, K. Cetinkaya, C. Li, and E. W. Knightly. Inter-class resource sharing using statistical service envelopes. 1999. (extended version of [96]; available online at <http://www.ece.rice.edu/networks/papers/QCLK00.ps>).
- [96] J. Qiu and E. W. Knightly. Inter-class resource sharing using statistical service envelopes. In *Proceedings of IEEE Infocom*, pages 36–42, March 1999.
- [97] M. Reisslein, K. W. Ross, and S. Rajagopal. A framework for guaranteeing statistical QoS. *IEEE/ACM Transactions of Networking*, 10(1):27–42, February 2002.
- [98] L. G. Roberts. The evolution of packet switching. *Proceedings of the IEEE*, 66(11):1307–1313, November 1978.
- [99] S. M. Ross. Bounds on the delay distribution in GI/G/1 queues. *Journal of Applied Probability*, 11(2):417–421, June 1974.
- [100] H. Sariowan, R. L. Cruz, and G. Polyzos. Scheduling for quality of service guarantees via service curves. In *International Conference on Computer Communications and Networks (ICCCN)*, pages 512–520, September 1995.
- [101] M. Schwartz. *Broadband Integrated Networks*. Prentice Hall, 1996.
- [102] J.A. Silvester and L. Kleinrock. On the capacity of multihop slotted ALOHA networks with regular structure. *IEEE Transactions on Communications*, 31(8):974–982, August 1983.
- [103] V. Sivaraman and F. M. Chiussi. Providing end-to-end statistical delay guarantees with earliest deadline first scheduling and per-hop traffic shaping. In *Proceedings of IEEE Infocom*, pages 603–612, March 2000.

- [104] V. Sivaraman, F. M. Chiussi, and M. Gerla. End-to-end statistical delay service under GPS and EDF scheduling: A comparison study. In *Proceedings of IEEE Infocom*, pages 1113–1122, April 2001.
- [105] D. Starobinski and M. Sidi. Stochastically bounded burstiness for communication networks. *IEEE Transactions on Information Theory*, 46(1):206–212, January 2000.
- [106] D. Stiliadis and A. Varma. Latency-rate servers: a general model for analysis of traffic scheduling algorithms. *IEEE/ACM Transactions on Networking*, 6(5):611–624, October 1998.
- [107] D. Stoyan. *Comparison Methods for Queues and Other Stochastic Models*. John Wiley & Sons Inc., 1983.
- [108] O. P. Vinogradov. A multiphase system with identical service. *Soviet Journal of Computer and Systems Sciences.*, 24(2):28–31, March 1986.
- [109] O. P. Vinogradov. A multiphase system with many servers and identical service times. *Stochastic Processes and Their Applications, MIEM*, 24:42–45, 1989 (in Russian).
- [110] M. Vojnovic and J.-Y. Le Boudec. Bounds for independent regulated inputs multiplexed in a service curve network element. In *Proceedings of IEEE Globecom*, pages 1857–1861, November 2001.
- [111] M. Vojnovic and J.-Y. Le Boudec. Stochastic analysis of some expedited forwarding networks. In *Proceedings of IEEE Infocom*, pages 1004–1013, June 2002.
- [112] M. Vojnovic and J.-Y. Le Boudec. Bounds for independent regulated inputs multiplexed in a service curve network element. *IEEE Transactions on Communications*, 51(5):735–740, May 2003.
- [113] W. Whitt. Tail probabilities with statistical multiplexing and effective bandwidths in multi-class queues. *Telecommunications Systems*, 2(1):71–107, December 1993.

- [114] D. Wischik. The output of a switch, or, effective bandwidths for networks. *Queueing Systems*, 32(4):383–396, 1999.
- [115] D. Wu and R. Negi. Effective capacity: a wireless link model for support of quality of service. *IEEE Transactions on Wireless Communications*, 2(4):630–643, July 2003.
- [116] O. Yaron and M. Sidi. Performance and stability of communication networks via robust exponential bounds. *IEEE/ACM Transactions on Networking*, 1(3):372–385, June 1993.
- [117] O. Yaron and M. Sidi. Generalized processor sharing networks with exponentially bounded burstiness arrivals. In *Proceedings of IEEE Incocom*, pages 628–634, June 1994.
- [118] Q. Yin, Y. Jiang, S. Jiang, and P. Y. Kong. Analysis on generalized stochastically bounded bursty traffic for communication networks. In *Proceedings of IEEE Local Computer Networks (LCN)*, pages 141–149, November 2002.
- [119] X. Yu, I.L.-J. Thng, Y. Jiang, and C. Qiao. Queueing processes in GPS and PGPS with LRD traffic inputs. *IEEE/ACM Transactions on Networking*, 13(3):676–689, June 2005.
- [120] H. Zhang and E. W. Knightly. Providing end-to-end statistical performance guarantees with bounding interval dependent stochastic models. In *Proceedings of ACM Sigmetrics*, pages 211–220, May 1994.
- [121] Z.-L. Zhang, D. Towsley, and J. Kurose. Statistical analysis of generalized processor sharing scheduling discipline. *IEEE Journal on Selected Area in Communications*, 13(6):1071–1080, August 1995.