Miranda van Uitert

Generalized Processor Sharing Queues

Stellingen

behorende bij het proefschrift

Generalized Processor Sharing Queues

van

Maria Johanna Gerarda van Uitert

I

In een buffer gevoed door veel aan/uit bronnen waarbij de aan- en/of uittijden een zwaarstaartige verdeling hebben, heeft een meting van het aantal bronnen dat aan is een grotere voorspellende waarde voor de evolutie van de bufferinhoud, dan wanneer de bronnen exponentieel verdeelde aan- en uittijden hebben (zie [1]).

[1] M. MANDJES, M.J.G. VAN UITERT. Transient analysis of traffic generated by bursty sources, and its application to measurement-based admission control. *Telecommunication Systems* 15 (2000), 295–321.

II

Beschouw een standaard Brownse beweging $\{B(t), t \geq 0\}$. Als $0 < c_2 < c_1$ dan geldt voor iedere $u \geq 0$ dat

$$\mathbb{P}\left(\sup_{t\geq 0}\{B(t) - c_2 t\} - \sup_{t\geq 0}\{B(t) - c_1 t\} > u\right) =$$

$$\frac{c_1 - 2c_2}{c_1 - c_2} e^{-2c_2 u} \mathbb{P}\left(\mathcal{N} < \frac{c_1 - 2c_2}{\sqrt{c_1 - c_2}}\sqrt{u}\right) + \frac{c_1}{c_1 - c_2}\mathbb{P}\left(\mathcal{N} > \frac{c_1}{\sqrt{c_1 - c_2}}\sqrt{u}\right),$$

waarbij $\mathcal{N}$ een stochast is met een standaard normale verdeling (zie [2]).

[2] K. DĘBICKI, M. MANDJES, M.J.G. VAN UITERT. Tandem queues and stationary independent increments. *In preparation* (2003).

III

Beschouw het tandem model zoals beschreven in Hoofdstuk 5, voor het specifieke geval van n bronnen die verkeer genereren volgens een *fractionele* Brownse beweging. We spreken van overflow wanneer bufferniveau nb wordt overschreden in de tweede wachtrij. Het meest waarschijnlijke pad naar overflow in het regime $n \to \infty$ bevat een interval tijdens welke op constante snelheid c_1 verkeer wordt gegenereerd.

IV

Beschouw een GPS systeem met 2 wachtrijen, bedieningssnelheid c en gewichten ϕ_1 en ϕ_2 (met $\phi_1 + \phi_2 = 1$). Wachtrij 1 wordt gevoed door A_1, een Gaussisch proces $\{A_1(t), t \geq 0, A_1(0) = 0\}$ met stationaire aangroeiingen, welke met kans 1 continue paden heeft. Voor de variantiefunctie $v(t) := \mathbb{V}\mathrm{ar}[A_1(t)]$ gelden de volgende aannames:

- $v(t) \in C[(0,\infty))$ is stijgend;
- $v(t)$ is regulier variërend in 0 met index $\beta \in (0,2]$ en $v(t)$ is regulier variërend in ∞ met index $\alpha \in (0,2)$.

Wachtrij 2 wordt gevoed door een aan/uit proces $A_2 := \{A_2(t), t \geq 0, A_2(0) = 0\}$, onafhankelijk van A_1, met regulier variërende aantijden. Tijdens een aanperiode wordt verkeer gegenereerd met een snelheid groter dan c. Gebruik makend van Stelling 6.1 in [3] kan worden aangetoond dat de volgende relatie geldt voor de verdeling van de stationaire werklast, Q_1, in wachtrij 1:

$$\lim_{u\to\infty} \frac{1}{u} \log \mathbb{P}(Q_1 > u) = \lim_{u\to\infty} \frac{1}{u} \log \mathbb{P}\left(Q_1^{\phi_1 c} > u\right),$$

met $Q_1^{\phi_1 c}$ de stationaire werklast in een wachtrij gevoed door A_1 met bedieningssnelheid $\phi_1 c$ (zie [4]).

[3] S.C. Borst, K. Dębicki, A.P. Zwart. Subexponential asymptotics of hybrid fluid and ruin models. *Submitted* (2003).

[4] K. Dębicki, M.J.G. van Uitert. Large buffer asymptotics for Generalized Processor Sharing queues with Gaussian inputs. *In preparation* (2003).

V

De discussie aan het eind van Hoofdstuk 5 wijst uit dat de Rough Full Link approximation in [5] een conservatieve heuristiek is en uitstekend voldoet voor praktische doeleinden.

[5] P. Mannersalo, I. Norros. Approximate formulae for Gaussian priority queues. *Proceedings ITC 17*, Salvador da Bahia, Brazil (2001), 991–1002.

VI

Beschouw een GPS systeem met bedieningssnelheid 1, dat klanten bedient van 2 klassen volgens gewichten ϕ_1 en ϕ_2 (met $\phi_1 + \phi_2 = 1$). Klanten binnen klasse 1 worden bediend volgens het Processor Sharing principe. Zij komen aan volgens een Poisson proces en hebben een bedieningsvraag met een *intermediate* regulier variërende verdeling. Neem voor de verkeersintensiteit van klasse 1, ρ_1, aan dat $\rho_1 < \phi_1$ en voor de totale verkeersintensiteit ρ dat $\rho < 1$. Voor een klant van klasse 1 die bij binnenkomst het systeem in evenwicht aantreft, geldt de volgende relatie voor de totale verblijftijd in het systeem, S_0,

$$\mathbb{P}(S_0 > t) \sim \mathbb{P}\left(S_0^{1-\psi_2} > t\right),$$

met

$$\psi_2 := \min\{\rho_2, \phi_2\},$$

waarbij $S_0^{1-\psi_2}$ de verblijftijd is van een klant die arriveert in een Processor Sharing systeem in evenwicht met bedieningssnelheid $1 - \psi_2$, dat alleen gevoed wordt door klasse 1 (zie [6]).

[6] S.C. Borst, R. Núñez-Queija, M.J.G. van Uitert. User-level performance of elastic traffic in a differentiated-services environment. *Performance Evaluation* 49 (2002), 507–519.

VII

Een zanger, hoe professioneel ook, kan nooit buiten de oren van anderen.

VIII

Hoewel het doen van onderzoek op topniveau een behoorlijke dosis enthousiasme vergt, wordt dit zelden weerspiegeld in de manier van presenteren door toponderzoekers.

IX

Positief aan het feit dat de universiteit steeds meer lijkt op een fastfood restaurant is dat daar in ieder geval een minimale kwaliteitsnorm gehanteerd wordt.

X

"Het is beter je eigen taak te verrichten, hoe gebrekkig ook, dan die van een ander, hoe bekwaam ook." [7, Boek III, Vers 35]

[7] *Bhagavad Gītā*. Ars Floreat, Amsterdam (1996).

Generalized Processor Sharing Queues

Cover design by Tobias Baanders
Printed by Ponsen & Looijen BV
ISBN 90 6464 709 7

THOMAS STIELTJES INSTITUTE
FOR MATHEMATICS

The research reported in this thesis has been carried out at the Center for Mathematics and Computer Science (CWI) Amsterdam.

Generalized Processor Sharing Queues

PROEFSCHRIFT

ter verkrijging van de graad van doctor aan de
Technische Universiteit Eindhoven, op gezag van de
Rector Magnificus, prof.dr. R.A. van Santen, voor een
commissie aangewezen door het College voor
Promoties in het openbaar te verdedigen
op maandag 24 november 2003 om 16.00 uur

door

Maria Johanna Gerarda van Uitert

geboren te Kerkdriel

aan Michel

Dit proefschrift is goedgekeurd door de promotoren:

prof.dr.ir. S.C. Borst
en
prof.dr.ir. O.J. Boxma

Dankwoord (Acknowledgements)

Een mengeling van weemoed en trots bekruipt me bij het schrijven van dit dankwoord. Weemoed, omdat ik stiekem toch genoten heb van de 4 jaar en 5 maanden die achter me liggen. Trots, omdat ik vrijwel alle hindernissen op weg naar de verdediging heb weten te nemen.

Het schrijven van dit proefschrift was een voortdurend leerproces, dat zonder de hulp van de nu te noemen personen in een vroegtijdig stadium zou zijn gestopt. Uiteraard begin ik met een grote buiging naar mijn promotor Sem Borst. Hij heeft het toch maar volgehouden om vierenhalf jaar lang al mijn vragen te beantwoorden. Ik heb hard mijn best gedaan om hem te betrappen op een 'zou je straks terug kunnen komen' of 'dit zou je nu toch onderhand moeten weten', maar dat is me niet gelukt. Verder ben ik blij om Onno Boxma in de hoedanigheid van tweede promotor te mogen bedanken. Met name heb ik genoten van de complexe zwerftochten door het Laplace-domein.

Van de reeks kamergenoten wil ik als eerste Sindo Núñez Queija bedanken. Niet in het minst omdat hij het met mij heeft weten uit te houden gedurende mijn gehele CWI-tijd. Als ik weer eens vol vuur een betoog aan het houden was over al het onrecht dat mij was aangedaan, wist hij me op adequate wijze tot bedaren te brengen. De tijd met Bert Zwart was korter, maar uitermate leerzaam voor mijn beheersing van het vakjargon. Van Bert heb ik geleerd dat wij op het CWI geen papers produceren en formules afleiden, maar verhalen schrijven en sommetjes uitrekenen. Krzyś Dębicki heeft mij laten kennis maken met de enorme bureaucratie die buitenlanders in ons landje ten deel valt. Verder bracht hij een behoorlijke portie huiselijkheid met zich mee, die goed bleek te zijn voor de productiviteit; twee van de stellingen berusten op gezamenlijk werk. Tenslotte wil ik Nidhi Hegde noemen die de laatste maanden de broodnodige vrouwelijke touch heeft geleverd. Van de overige CWI-ers wil ik graag als eerste Mervyn Lewis bedanken, die ik heb leren kennen aan het begin van mijn, en toevallig ook zijn, OiO-periode. De cliché-uitspraak 'gedeelde smart is halve smart' en 'gedeeld plezier is dubbel plezier' is hier bijzonder op zijn plaats. Ook Asia Dębicka wil ik hier graag noemen; de combinatie yoga en eigengemaakt gebak was een voltreffer. Tenslotte bedank ik Ton Dieker en René Bekker voor hun jeugdig enthousiasme en het opnieuw in het leven roepen van trakteren, koffie drinken en gezamenlijke etentjes.

Dan wil ik tot slot graag het woord richten aan de mensen die vooral op persoonlijk vlak van belang zijn geweest. Francien, ik ben een bevoorrecht mens dat ik jouw steun heb gehad de afgelopen maanden. Sofie en Marieke, het was altijd fijn om welke beslommeringen dan ook met jullie te delen. Michel, ik heb het proefschrift aan jou opgedragen, vanwege jouw grenzeloze vertrouwen in mij. Pap en mam, jullie ben ik in het bijzonder dankbaar voor alle eigenschappen die ervoor gezorgd hebben dat ik dit proefschrift heb kunnen schrijven. Henry, Corina, Marianne, Annemarie, Gerard, Linda, Agnes, Henk en Paul, bij jullie zijn voelt voor mij als 'thuis zijn' en dat is van grote waarde. Henry en Marieke, tof dat jullie mijn paranimfen willen zijn.

Miranda van Uitert
September 2003

Contents

Chapter 1

Introduction

Generalized Processor Sharing (GPS) has emerged as an important paradigm for achieving *service differentiation* among heterogeneous applications in future communication networks. Over the past decade, GPS has been investigated under an ever increasing variety of assumptions and models. This monograph deals with the performance analysis of GPS for various types of traffic possibly exhibiting *self-similarity* or *long-range dependence.*

With GPS, the service capacity is shared according to *weights* that are assigned to the various traffic classes. GPS is capable of offering protection among traffic classes, while achieving full statistical multiplexing. It is based on a *fluid principle*, and serves as a reference model for packet scheduling mechanisms like *Weighted Fair Queueing* (WFQ).

In this monograph we concentrate on the amount of traffic in the buffer at a network node, the *workload.* We compute the asymptotic behavior of the workload distribution, either by letting the buffer size grow to infinity or the number of traffic flows approach infinity. For these calculations we use probabilistic sample-path techniques, Tauberian theorems and a sample-path large-deviations principle.

In this first chapter we describe the most important concepts of our study. We start in Section 1.1 with an overview of recent developments in communication networks and explain the growing need for scheduling mechanisms. In Section 1.2 we introduce some basic principles of communication networks and relate them to queueing theory. We also describe the traffic models that we use in this monograph. Section 1.3 is devoted to a detailed explanation of GPS. A generic representation of the main results in this monograph is given in Section 1.4. In Section 1.5 we explain briefly which performance measures, other than the workload, could be considered. An overview of the most important literature on GPS is given in Section 1.6. Section 1.7 is devoted to a description of disciplines other than GPS, that are studied in this monograph as well. We end this chapter with an outline of the remainder of this monograph.

1.1 Scheduling in communication networks

In this first section we further motivate the use of scheduling mechanisms, and in particular that of GPS. Then we present a high-level account of its functionalities (a formal description of GPS is postponed to Section 1.3). We comment on some packet-based emulations of GPS.

1.1.1 Role of scheduling mechanisms

The design of the current Internet is based on a paradigm that sufficed for simple data applications, the *best-effort* principle. It means that the network treats every packet equally and does not provide guarantees on the experienced packet losses or packet delays. When there are only data applications, there is indeed no need to differentiate between packets of different applications, or to provide strict loss or delay guarantees. However, various new applications have emerged with different, possibly more stringent requirements in terms of packet losses and delays. Clearly, the best-effort principle does not suffice when these new applications have to be supported over the same network as the existing data applications.

The simplest way to provide each application the quality of service (QoS) it requires, is to build a separate network for each new application, or to strictly dedicate part of an existing network to it. However, one can imagine that this approach has some severe drawbacks. It may take a long time before a new application can be introduced, networks will not be used efficiently, and if the application fails to become a success, or becomes rapidly obsolete, then the network will be useless. Therefore, current communication networks are evolving towards *integrated-services* networks, which support a wide range of applications on a common infrastructure. The question then arises how such networks can provide all demanding applications with the proper QoS. A natural answer to this problem is to *differentiate* within the network between traffic from different applications. Without differentiation, one would have to provide all applications with the QoS of the most demanding one. Obviously, this would lead to an extremely inefficient use of the network capacity.

One of the instruments that can be used to accomplish service differentiation, is the so-called *scheduling mechanism.* This mechanism has to be implemented in the routers or switches of a network. For each packet arriving at the router or switch, it determines at what time it is forwarded to the next router or switch on its route. If no scheduling mechanism is used, then by default packets are sent in the order of arrival. Common scheduling mechanisms in current communication networks are packet-based incarnations of a so-called ideal fluid-based discipline: Generalized Processor Sharing (GPS). The most important packet-based emulations are variants of Weighted Fair Queueing (WFQ), a mechanism that was first proposed in [51]. In this monograph we concentrate however on the performance analysis of GPS, which is ideally fair and has some tractable analytical properties.

1.1.2 GPS

In short, GPS works as follows (in Section 1.3 the discipline will be explained in more detail). All traffic that is transported over the network is divided into so-called traffic classes, which can either be individual flows or consist of several flows with similar QoS requirements. Let us focus on one particular node in a network. If GPS is used, then each class that traverses this node is assigned a positive weight. This weight specifies the guaranteed minimum capacity for a class. If a particular class does not fully use its capacity, then the excess capacity becomes available for other classes passing through this node, on top of their minimum. As the same weights are used to redistribute the excess capacity as well, their role is twofold.

The two-class GPS discipline with the full capacity guaranteed to one class, reduces to a two-class (strict) priority mechanism. When using a strict priority mechanism, there is a danger that the high-priority class will take all capacity, leaving nothing for the low-priority class. The throughput of the low-priority class may suffer severely, especially when it is regulated by an end-to-end flow control mechanism like TCP (transmission control protocol). With GPS, this starvation phenomenon can be prevented by always taking positive minimum capacities.

The importance of a protection mechanism like GPS became even more apparent when measurements showed that data traffic is typically highly variable or bursty over a wide range of time scales (see [82, 112]). These characteristics manifest themselves in long-range dependence (LRD) or self-similarity (see Section 2.1 for a definition of these concepts). Traffic flows with such properties can behave very badly in the sense that they can grab all capacity that is available for a relatively long period. We will explain in Section 2.1 that long-range dependent or self-similar traffic can be modeled using so-called *heavy-tailed* distributions. Especially traffic flows with short-range dependent (SRD) or light-tailed properties will be extremely affected by the heavy-tailed flows, when there is no protection mechanism. In this monograph we study the GPS discipline for both light-tailed and heavy-tailed distributions. In particular, we investigate whether GPS is able to cope with this long-range dependent (or heavy-tailed) traffic and can act as a protection mechanism for SRD (or light-tailed) traffic.

For the sake of completeness we briefly mention another, related, fluid discipline, *Discriminatory Processor Sharing* (DPS), which is studied in [62, 115, 116]. Like in GPS, the traffic flows under the DPS discipline are divided into classes, and each class is assigned a positive weight. However, in contrast to GPS, the service capacity is then shared by *all* flows according to their weights. Recall that with GPS, the weight determines the service rate for the entire class, regardless of the number of active flows within the class. In fact, DPS can be viewed as a generalization of ordinary or egalitarian Processor Sharing (PS) where all flows receive the same service rate.

1.1.3 Implementation issues

An assumption underlying GPS is the *infinite divisibility of capacity.* This means that we assume that traffic of several classes can be served at the same time. Because in reality traffic is composed of discrete packets, which have to be sent sequentially, GPS cannot be implemented in its strict form. However, packet sizes are extremely small when compared to network capacity, meaning that the fluid assumption of GPS is natural over somewhat longer time scales.

Several studies have been conducted to identify packet-based mechanisms that closely resemble GPS. Clearly, a trade-off exists between the accuracy of the approximation and the implementation complexity, as a better approximation of the fluid paradigm will tend to be more complex. Well-known studies on this problem are [107, 108], in which WFQ (the authors use the term packet-based GPS) is analyzed in full detail. With WFQ, it is calculated for every arriving packet when it would have completed service under GPS, if no additional packets were to arrive after it. Packets are then served in the order of completing service under GPS. One can imagine that this scheme is much harder to implement than for instance Weighted Round Robin (WRR), where classes are served in a fixed order. However, it is shown in [107, 108] that WFQ approximates GPS to within one packet transmission time and performs better than WRR. Because WFQ is a good approximation for GPS, the results for the GPS model should hold for WFQ too. Hence, we expect the results in this monograph for GPS to hold for WFQ. Note the similarity with the results for ordinary PS as an approximation for the Round Robin (RR) mechanism, with the difference that PS served as a convenient idealization of RR, whereas WFQ in fact only originated as an emulation of GPS. We refer to [105, 143, 144] for an overview of the literature on PS.

A survey of GPS-based packet scheduling mechanisms that have been proposed in the literature and their performance is given in [129]. We mention also [10], where hierarchical packet-based mechanisms are proposed, and [81], where a packet-based version of GPS is found by rewriting the problem as a zero-sum stochastic game. Finally we refer to [132] where each class is assigned two weights in order to decouple the two roles of a weight in GPS.

1.2 Modeling of communication networks

In this monograph we adhere to the tradition of using queueing theory to evaluate network performance. First, we discuss the relevant notions in communication networks, and relate them to their counterparts in queueing theory. Then we sketch the mathematical models representing the traffic processes involved.

1.2.1 Terminology

We are primarily interested in packet-switched networks, where information travels in packets. The network consists of switches (or routers, but for convenience we will only use the term switch) that are connected through links. The switch contains a scheduling mechanism, which puts packets on the next link towards their destination. Since multiple packets may arrive when the scheduling mechanism is forwarding another packet, the switches contain buffers, where packets can be temporarily stored. The "train" of packets traveling from a particular source to its destination is referred to as *flow*, *connection* or (with abuse of terminology) *source*. Usually the packets of a flow are sent in *bursts*, i.e., intervals in which packets are sent back-to-back are interchanged with intervals in which no packets are sent.

Now we jump to queueing theory and explain how we model a packet-switched network as a queueing system. The server in the queueing model represents the scheduling mechanism, and the service rate corresponds to the speed at which the packets are put on the link. The packets stored in the buffer represent the queue (while the entire switch can be viewed as a queue in the broad sense). We use the terms *workload*, *buffer content* and *queue length* to denote the amount of packets in the queue. This leaves only one important notion in queueing theory unexplained: the customer and its service requirement. In the above description clearly the packets should be seen as customers, with the sizes of the packets being their service requirements. However, it is often convenient to view a cluster of packets as a customer. If we are interested in the performance of the flow as a whole, for instance the perceived throughput, then we take the entire flow as a customer and the total amount of packets as the service requirement. We can also take the number of flows fixed, and view a burst as a customer with as service requirement the packets sent within a burst. Depending on the interpretation of the customers, we speak of flow-level, burst-level and packet-level models (see for instance [119]).

1.2.2 Traffic models

All traffic models that are used in this monograph, assume the number of flows to be fixed, and can thus be placed at the burst level. For the flows we make the assumption that they send traffic as a continuous fluid flow (possibly in bursts) rather than a stream of packets. This assumption is motivated by the fact that packet sizes are so small compared to the service rate, that the stream of packets within a burst can be well approximated by a fluid. We expect the burst-level results for GPS to give a reasonable characterization of the burst-level performance of WFQ.

As explained above, we are interested in the performance of GPS when dealing with SRD traffic, LRD traffic, or both. Below we indicate for each traffic model if it is able to capture SRD or LRD characteristics.

We start with the model where a flow is assumed to generate instantaneous traffic bursts. Assuming the burst sizes and the independent interarrival times to have a general distribution, this model is the classical GI/G/1 queue, where customers with independent generally distributed service requirements enter the system at independent, generally distributed interarrival times. As we explained above, the customers correspond to the bursts of the flow. By varying the distribution of the burst size, we can incorporate both light-tailed and heavy-tailed characteristics.

The previous model assumed that bursts can be sent instantaneously. In practice, however, traffic is sent gradually into the network. A model that takes this feature into account is the *on-off model*. During the on periods traffic is sent at a fixed rate, referred to as the *peak rate*, whereas nothing is sent during the off periods. This model is able to capture both SRD and LRD, simply by choosing the distributions of the on and the off periods appropriately. For instance, when both the on and off period have a light-tailed distribution, the model is SRD. On the other hand, when the on period, or the off period, or both are distributed as a random variable X with a regularly varying distribution, i.e.,

$$\mathbb{P}(X > t) \sim t^{-\nu}, \ t \to \infty,$$

with $\nu \in (1, 2)$ (see Section 2.1 for a more elaborate explanation), it is shown in [35] that this model shows LRD behavior. For any two real functions $g(\cdot)$ and $h(\cdot)$, we use the notation $g(t) \sim h(t)$ (as in the previous formula) to denote

$$\lim_{t\to\infty} \frac{g(t)}{h(t)} = 1, \text{ or equivalently, } g(t) = h(t)(1 + \mathrm{o}(1)) \text{ as } t \to \infty.$$

The superposition of N identical independent on-off flows, with either heavy-tailed on or off periods, exhibits LRD and self-similarity for $N \to \infty$ (see [138]). See also [35] and [150] for a more elaborate discussion.

It is possible that a flow sends its (fluid) traffic at multiple rates, instead of only at rate zero or at peak rate as in the previous model. A model for such flow captures both SRD and LRD by appropriately choosing the distribution of the period that a flow sends at a certain rate. More formally we say that a flow can be in a finite number of states, where each state represents a rate at which traffic is generated. For instance, a popular model for SRD traffic is *Markov-modulated fluid*. Then the transitions between the various states are governed by a Markov process, i.e., the sojourn time in each of the states is exponentially distributed, and the flow jumps with a certain probability to another state. See [123] for a survey on the existing literature and a more thorough explanation of this model. It is also possible to take distributions other than the exponential one for the sojourn time in a certain state. If at least one of the sojourn times is a random variable following a regularly varying distribution, then this model exhibits LRD. For a formal definition of this class of models we refer to [44].

In the last chapters of this monograph we use a somewhat different traffic model. Traffic is again sent as fluid, but now the flows behave as *Gaussian processes* with stationary increments. Such a Gaussian flow is characterized completely by (i) the mean input rate and (ii) the variance function, i.e., the variance of the amount of traffic sent in an interval. As the traffic correlation pattern is captured by the variance, the Gaussian model is well-suited to incorporate both short-range and long-range dependence. Choosing for instance the variance function of a process that exhibits LRD properties will result in a LRD Gaussian process. The most familiar LRD version of Gaussian traffic is fractional Brownian motion (fBm) (see Section 2.1), a process that was first used in a queueing model in [100].

1.3 Formal description of GPS

In this section we give a mathematical description of a model with GPS. We explain the operation of the GPS discipline for the case where m classes are served, with m fixed, and the total service rate is c (see also Figure 1.1 for a graphical illustration). Each of the m classes has its own queue, and may consist of one or several flows. We note that the models considered in subsequent chapters can have different values of m and c, or even allow for networks of GPS nodes (Chapter 3). We say that a class is backlogged if it has a positive buffer content and hence, requires service. Class i is assigned a non-negative weight ϕ_i, $i = 1, \ldots, m$, and without loss of generality we take the weights such that they add up to 1, i.e., $\sum_{j=1}^{m} \phi_j = 1$. Hence, class i has a guaranteed minimum rate of $\phi_i c$, i.e., if all classes are backlogged, then class i receives service at rate $\phi_i c$. If some of the classes do not require service, then the surplus service rate is redistributed among the other classes in proportion to their respective weights. Denoting the set of backlogged classes with $\mathcal{B}$, a backlogged class $i \in \mathcal{B}$ receives service at rate

$$\frac{\phi_i}{\sum_{j \in \mathcal{B}} \phi_j} c \geq \phi_i c. \tag{1.1}$$

In the above description we implicitly assumed that a class only consists of flows that generate instantaneous traffic bursts. We remark that in case of gradual input some subtleties may arise. It is then possible that a class receives service (equal to the input rate) but is not backlogged, so that the above description needs to be somewhat refined.

We introduce some notation in order to explain the most important GPS related concepts that will be used throughout this monograph. Note that we will frequently deviate from this notation in subsequent chapters, as each GPS system that we discuss has its own specific features. Denote by $A_j(s,t)$ the amount of traffic *generated* by flows of class j during time interval $(s,t]$, $s < t$, $s, t \in \mathbb{R}$, $j = 1, \ldots, m$. We define $B_j(s,t)$ to be the amount of service *received* by flows of class j during $(s,t]$. Denote by $V_j(t)$ the workload of class j at time t, $t \in \mathbb{R}$, and by V_j a random variable whose distribution is the limit distribution

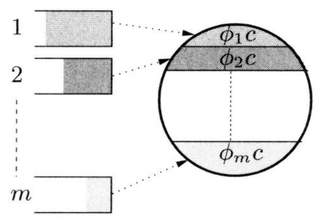

Figure 1.1: The GPS mechanism.

of $V_j(t)$ for $t \to \infty$, assuming it exists (for reasons of completeness we define it here, but we do not use the stationary version V_j until Section 1.4). We need the following identity, which relates the workload of a particular class at time t to its workload at any time s before t,

$$V_j(t) = V_j(s) + A_j(s,t) - B_j(s,t), \quad \text{for } s < t. \tag{1.2}$$

Clearly, in order to compute the workload of class j, it is necessary to know the amount of service that this class has received. The available service capacity for a class is varying as it depends on the service requests of the other classes. Only by knowing the precise state of the system at each point in time, it is possible to keep track of the amount of service that a class receives. Thus, the effective but complicated interaction between classes turns out to be a major difficulty in the analysis of GPS.

Let us first consider a single FIFO queue with input process $A(s,t)$ and constant service rate c. It is shown in the proof of Lemma 5.1 in [117] that the following representation holds for $V(t)$, the workload at time t,

$$V(t) = \sup_{s \le t}\{A(s,t) - c(t-s)\}.$$

In addition, it is proved that the optimizing s denotes the beginning of the busy period that contains time t. (It is easily seen that if it were not true, the supremum could be increased by changing s.) In this monograph we will often use the busy-period interpretation of the optimizing argument.

We can give a similar representation of $V_j(t)$ that only depends on the amount of *class-j* traffic that has arrived and the service rate that was available. Assuming that service rate $c_j(u)$ is available for class j at time u, it can be proved along the same lines as in [117] that

$$V_j(t) = \sup_{s \le t}\left\{A_j(s,t) - \int_s^t c_j(u)\mathrm{d}u\right\}, \tag{1.3}$$

with the optimizing s denoting the beginning of the busy period of class j containing time t.

The GPS discipline is said to be work-conserving, because it never allows idleness as long as there is a queue with positive workload. The total workload in the system is thus the same as that in a single queue, which is fed by the input processes of all classes, i.e.,

$$\sum_{j=1}^{m} V_j(t) = \sup_{s \le t}\left\{\sum_{j=1}^{m} A_j(s,t) - c(t-s)\right\}. \tag{1.4}$$

When proving results for the various GPS models, we will often compare the workload of a certain class with the workload in a system where this class is served in isolation at a fixed rate. We denote with $V_j^c(t)$ the workload of class j at time t in a system where it is served in isolation at rate c and with V_j^c a random variable whose distribution is the limit distribution of $V_j^c(t)$ for $t \to \infty$, assuming it exists. Obviously, analogues of (1.2) and (1.3) continue to hold for $V_j^c(t)$.

The following inequality follows from the protection property of GPS. In this monograph, it is the most frequently used way of comparing workloads in different systems. For $t \in \mathbb{R}$,

$$V_j(t) \le V_j^{\phi_j c}(t). \tag{1.5}$$

It is easily seen that this inequality should hold. In the GPS system, class j has at least service rate $\phi_j c$ if it is backlogged, and can obtain an even higher service rate if some of the other classes do not fully use theirs (see (1.1)). In the isolated system with service rate $\phi_j c$, exactly $\phi_j c$ is available for class j if it has a positive workload. Thus, a GPS system outperforms a corresponding isolated system when taking the workload as performance measure. This immediately implies that also the delay performance will be better in the GPS system than in a corresponding isolated system. Of course, this is not so surprising in view of the role of the weights as minimum capacity guarantees. The more relevant question is how significant the performance gains are and how the results depend on the modeling assumptions.

Throughout this monograph, we assume that the input processes, $A_j(s,t)$ in the model discussed above, are ergodic in the sense that

$$\lim_{t\to\infty} \frac{A_j(0,t)}{t} =: \rho_j$$

with probability 1. Hence, we assume the long-term average input rate ρ_j to exist for all input processes. We use the notation ρ_{-i} for the sum over all average rates except that of flow i, i.e.,

$$\rho_{-i} := \sum_{j \neq i} \rho_j. \tag{1.6}$$

1.4 Main results

In the first part of this section we explain what kind of results we obtain in the next chapters. The second part elaborates on the numerical aspects.

1.4.1 Asymptotic regimes

In this monograph we concentrate on deriving the tail behavior of the workload distribution. Throughout, it is assumed that the buffer size is infinite, meaning that all arriving traffic at a buffer is allowed to enter the buffer (we comment on some results for finite buffers in Section 1.5). We are interested in two asymptotic regimes. Chapters 3, 4 and 5 contain results for the so-called *large-buffer* regime. We calculate the probability that the workload of an individual class in a GPS system exceeds a certain level, and we let this level go to infinity.

In Chapters 6 and 7 we focus on the *many-sources* regime. We determine the probability of the buffer content exceeding a fixed level (scaled with the number of input flows), and we let the number of input flows grow to infinity.

There are several reasons for considering the *asymptotic* behavior of the workload distribution instead of the *exact* distribution. First, and most importantly, the problem of calculating the exact distribution is extremely hard. This is illustrated by the fact that no closed-form results are known for the exact distribution of the workload of an individual class in a GPS system. Only for the special case of a 2-class GPS system, where both classes are fed by a Poisson arrival process, the Laplace Stieltjes transform (LST) is calculated in [42]. Although the LST in principle captures all relevant information regarding the exact distribution, it is already a tedious task to derive the large-buffer asymptotics from this formula (see for instance our calculations in Chapter 5), let alone the exact distribution. Even if one *were* able to find the exact distribution, then the resulting formula would most probably be less insightful than formulas for asymptotic results, as the asymptotic formulas only show the fundamental properties of the system. Another advantage of considering the asymptotic behavior is that the proofs of such results are often insightful, as they reveal how the system must have behaved in order for the buffer level to be exceeded (see Section 2.2 for a detailed explanation). Hence, the complexity of the GPS system motivates the derivation of asymptotic relations from a theoretical point of view. In practice, telecommunication systems have to be designed such that buffer overflow only rarely occurs. Translated into queueing theory, this means that we have to concentrate on asymptotics. Moreover, it is important to gain insight in the events that cause this rare behavior. Not surprisingly, these events are referred to as *rare*, as they happen only with a very small probability. In the large-buffer regime, the event of reaching a large buffer level will be rare, whereas in the many-sources regime, the event of reaching a fixed buffer level (scaled with the number of sources) will be rare.

Asymptotic results appear in two flavours, *exact* asymptotics and *logarithmic* asymptotics.

Exact asymptotics. In Chapters 3, 4 and 5 we derive exact asymptotics for the workload distribution in the large-buffer regime. We briefly mention that it is also possible to calculate exact asymptotics in the many-sources regime, see for instance [46, 84]. An exact asymptotic relation in the large-buffer regime for the workload distribution of class i in a GPS system has the following form:

$$\mathbb{P}(V_i > x) \sim h(x), \quad x \to \infty,$$

for some function $h(x) \in [0, 1]$. For instance in Theorem 3.4.1 we have

$$h(x) = C_1 x^{1-\nu},$$

with C_1 a positive constant and $1 < \nu < 2$. In this example the probability of the workload exceeding level x decays polynomially in x. In Theorem 5.4.2 the function $h(x)$ has a different form, i.e.,

$$h(x) = C_2 \mathrm{e}^{-C_3 x},$$

with C_2 and C_3 positive constants. Now clearly the probability of the workload exceeding x decays exponentially in x. A combination of the previous formats exists as well. See for instance the result in Theorem 4.2.1 (see also the example given in Section 4.2.3), where

$$h(x) = C_4 \mathrm{e}^{-C_5 x} x^{1-\nu},$$

with C_4 and C_5 positive constants and $1 < \nu < 2$. We say that $h(x)$ is a semi-exponential function, as the exponential decay in x dominates.

As will be explained in more detail in Section 2.2, the function $h(x)$ is usually the (large-buffer) asymptotic behavior of the workload in a related system. As an example we mention the first result on exact large-buffer asymptotics for GPS systems with heavy-tailed input [16]. The asymptotic behavior of the workload of class i in an m-class GPS system, with class i exhibiting heavy-tailed characteristics, is as follows:

$$\mathbb{P}(V_i > x) \sim \mathbb{P}(V_i^{c-\rho_{-i}} > x), \quad x \to \infty.$$

Loosely speaking, if the buffer grows large, then the workload of flow i in the GPS system is as if flow i is served in isolation at constant rate $c - \rho_{-i}$. It can be shown that

$$\mathbb{P}(V_i^{c-\rho_{-i}} > x) \sim C_6 x^{1-\nu}, \quad x \to \infty,$$

which is due to [67, 106] (see Sections 3.1.1 and 3.1.2 for a more detailed explanation).

Logarithmic asymptotics. In Chapters 6 and 7 we derive logarithmic asymptotics for the workload distribution in the many-sources regime. With n denoting the number of sources belonging to class i, the results have the following form:

$$\frac{1}{n} \log \mathbb{P}(V_i > nx) \sim -h(x), \quad n \to \infty,$$

with $h(x)$ a function taking values in $\mathbb{R}_+$ (using $\sim$ for $n \to \infty$ according to the same notational convention as previously for $t \to \infty$). The function $h(x)$ is typically known as the *decay rate*. Hence,

$$\mathbb{P}(V_i > nx) \sim g(n,x) \mathrm{e}^{-nh(x)}, \quad n \to \infty,$$

where the function $g(n,x)$ is typically not calculated, but only known to be subexponential (see Section 2.1),

$$\lim_{n \to \infty} \frac{1}{n} \log g(n,x) = 0.$$

One such result is found in Section 6.4 for the two-class priority model, where both classes consist of n input flows. For fBm with $H \in (\frac{1}{2}, 1)$ (see Section 2.1 for an explanation) we find in Theorem 6.4.2 for case (A):

$$h(x) = C_7 x^{2-2H},$$

with C_4 a finite positive constant. For a light-tailed process it typically holds that

$$h(x) - C_8 x - C_9 \to 0, \ \ x \to \infty,$$

with C_8 a positive constant and $C_9 \in \mathbb{R}$, i.e., $h(x)$ is asymptotically linear in x (see [91]). Similarly, logarithmic asymptotics can be found in the large-buffer regime. See for instance the work in [145], which covers a whole range of logarithmic large-buffer results for GPS models (see also Section 1.6).

1.4.2 Numerical aspects

As the situation for the large-buffer regime is different from that for the many-sources regime, we discuss the numerical aspects separately.

Large-buffer regime. The results obtained in Chapters 3, 4 and 5 capture a single scenario that is responsible for the workload exceeding some value x. Thus, the asymptotics tend to provide a lower bound for the true probabilities. If x is very large, then the particular scenario captured by the asymptotic result is said to dominate compared to all others. This means that given that the event under consideration occurs, the dominating scenario occurs with overwhelming probability. However, when x becomes smaller, other potential scenarios can also significantly contribute, and the asymptotics may drastically underestimate the actual probabilities. In order to obtain more accurate results, these other scenarios should thus be taken into account as well. This can be accomplished by calculating more terms in the asymptotic expansion, where each of the terms will be an alternative scenario that contributes to the workload reaching value x. For instance in [137] a three-term asymptotic expansion is obtained for the tail of the waiting-time distribution in an M/G/1 model with heavy-tailed service times. However, it is shown in [2] that even three-term asymptotic expansions do not provide a good approximation for values of x that are of interest in current communication networks. Hence, the results in Chapters 3, 4 and 5 mainly serve to gain qualitative insight in the behavior of the system, especially for the situation that the workload grows large. Moreover, they can serve as a rough guide line for selecting weights. As an example we mention the model in Chapter 4, where the light-tailed flow turns out to be protected against the heavy-tailed flow if the weight of the light-tailed flow is choosen larger than its average traffic intensity.

Summarizing, although large-buffer asymptotics give valuable qualitative insights into the rare-event behavior, they may not provide an accurate approximation of the probability, except for very large values of the buffer level. Thus further study is necessary to obtain accurate approximations for small to moderate buffer levels. We remark that this is a very challenging problem as both numerical and simulation experiments for heavy-tailed distributions are notoriously difficult (see for instance [15]).

Many-sources regime. In Chapters 6 and 7 the situation is somewhat different as we obtain asymptotic expansions that decay exponentially in the number of sources. In several studies, e.g. [38, 57], it has been shown that these expansions could serve as good approximations, even when the number of sources is relatively small. Moreover, for the models in this monograph, the asymptotic results overestimate the true probabilities. We have exploited this knowledge in Section 7.7.2, where we also use the asymptotic expansions as approximations for the true overflow probabilities in a weight-setting algorithm.

1.5 Related performance measures

Closely related to infinite-sized buffer models as considered in this monograph, are models with finite buffers. It is interesting to analyze the asymptotic behavior of the loss probability for such models. Only few results on the asymptotic behavior of the loss probability are available for GPS models. Remarkably enough, in some models the computation of loss probabilities for finite-buffer situations is somewhat easier than that of workload probabilities in infinite-buffer models (see for instance Chapter 7 in [150]).

We mention the results for a GPS system with m classes and a finite buffer of size B in [68]. It is assumed that each class is fed by a single on-off flow with heavy-tailed on times, on rate r_j, and general off times. Under the assumption that $\rho_j < \phi_j c$ for all classes $j \in \{1, \ldots, m\}$, and $r_i > c - \rho_{-i}$ for flow i, it is shown that the loss probability (as $B \to \infty$) of flow i behaves as that in an isolated system with capacity $c - \rho_{-i}$ and buffer size B. This result is very similar to the large-buffer asymptotics obtained in [16], where the same model is considered with infinite buffer size. Note that we need the assumption that $r_i > c - \rho_{-i}$ for the case of on-off input flows in [16]. Both the finite-buffer and infinite-buffer result show that the asymptotic behavior is the same as that in a system where flow i is served in isolation at fixed rate $c - \rho_{-i}$. Denoting by V_j^B and $V_j^{B,c}$ the finite-buffer equivalents of V_j and V_j^c as defined in the previous section, we have for the *finite*-buffer model, see [68],

$$\mathbb{P}(V_i^B > x) \sim \mathbb{P}(V_i^{B,c-\rho_{-i}} > x), \;\; B \to \infty, \;\; x \to \infty.$$

Thus, asymptotically, flow i has the entire buffer B for itself. As mentioned in Section 1.4 (see [16]), and further explained in Section 2.2.1, the following relation holds for the *infinite*-buffer model:

$$\mathbb{P}(V_i > x) \sim \mathbb{P}(V_i^{c-\rho_{-i}} > x), \;\; x \to \infty.$$

The formulas for the behavior of $\mathbb{P}(V_i^{B,c-\rho_{-i}} > x)$ and $\mathbb{P}(V_i^{c-\rho_{-i}} > x)$ as $x \to \infty$, only differ by a constant (see also [69]), i.e., the loss probability is a fraction of the exceedance probability.

Another common performance measure in classical queueing models is the sojourn time or delay of a customer, which is the total time the customer spends in the system. Thus, the sojourn time consists of the time spent in the queue, also referred to as the waiting time, and the time spent in service. In the ordinary PS model for example, it is natural to consider the sojourn time, as the customers are immediately taken into service. In [30] we studied the sojourn time of customers in a two-class GPS system, where we assumed the service discipline within the class under consideration to be PS.

The sojourn time can also be studied in models where traffic is generated through instantaneous bursts, as described in Section 1.2.2. We can for instance view the burst as the customer, and calculate the time it takes for this burst to leave the system. The concept of a customer is not so clear in the other models that we discussed, where traffic is sent gradually. Recall that we used the fluid models to approximate the way packets are sent into the network. As we explained, the sizes of the packets are small compared to the service capacity, so we could view a fictitious particle as a representation for a packet. The delay of such a particle in fluid models can then be seen as a measure for the delay of a packet.

In the literature, two notions of delay can be found, which are quite different: *virtual* delay and *actual* delay. Virtual delay is defined as the time it takes for the buffer to be emptied at an arbitrary instant in time. We remark that the virtual delay is closely related to the workload. Because all time epochs are taken into account, also those at which no fluid arrives, it does not provide much insight into the actual delay of a fluid particle. Results on the virtual delay, or the virtual waiting time, can be found in, e.g., [78, 79, 102, 104, 111]. All results indeed follow easily from the results that are known for the workload.

The other notion is the actual delay, which is defined as the time between the arrival of an arbitrary fluid particle and its departure from the system. To determine the actual delay, we have to condition on the fact that an arrival occurs, and thus weigh the virtual delay with the arrival rate. Results on the actual delay in fluid models are scarce. In [124] it is shown that the logarithmic many-sources asymptotics in a single discrete-time FIFO queue for the actual delay are equal to those for the virtual delay, under some technical statistical assumption on the input process. The result is also shown to hold for two-queue priority models. Also in [3] results can be found on the actual delay for a fluid particle in a special kind of on-off model.

As a final remark we mention that one has to be careful with the notion of actual delay when the input rate is not well-defined. For instance in the case of fBm, the notion of an input rate does not make sense as the sample paths of $A(s,t)$ are not differentiable with probability 1.

1.6 Relevant literature

In this section we review the most important results that have appeared on GPS. We use some concepts that have not been introduced before, but will be explained in detail in the next chapter.

The first *mathematical* analyses were performed in [107] (the single-node model) and [108] (the multiple-node case). Under the assumption that traffic is shaped by a so-called leaky bucket mechanism, a *deterministic* upper bound for both the delay and workload per class is derived. Recently this approach was revisited in [113], where a different type of leaky bucket is used to shape traffic, the fractal leaky bucket. It is argued in [113] that such a leaky bucket is a more effective mechanism for traffic exhibiting long-range dependence, especially for traffic behaving as fBm. The most important results are again a deterministic upper bound on the delay and workload per class.

The *probabilistic* analysis of GPS was initiated in [141, 142]. The inputs are assumed to be (β, C, α) exponentially bounded burst (EBB) processes (see [140] for more results), i.e.,

$$\mathbb{P}(A(s, s+t) \geq \beta t + x) \leq Ce^{-\alpha x}.$$

Under some assumptions, it is shown that the output process of an EBB GPS class is again EBB. In [149] this work is continued by calculating upper bounds on the tail distribution for the backlog and delay of an EBB GPS class, where the input processes of the classes can be dependent. Similar results are shown for a network of GPS servers and EBB input processes.

The analysis for a broader class of light-tailed processes is pursued in [48]. Under the assumption that the log-moment generating function of the arrival processes is finite, an upper bound on the logarithmic large-buffer asymptotics in a discrete-time GPS system is calculated. The results in [48] are extended in [146], where under a similar assumption on the log-moment generating function, the exact logarithmic large-buffer asymptotics for a two-queue GPS model are derived. In [147] lower and upper bounds are given for the logarithmic large-buffer asymptotics of the multiple-queue model (tightness of the lower and upper bound only holds for the case with two queues). For processes that have a finite scaled log-moment generating function, upper bounds on the logarithmic large-buffer asymptotics are derived in [145]. Note that a broad class of heavy-tailed arrival processes fits in this framework. Although the results are quite different from the ones that will be derived in this monograph, we would like to mention that in this work sample-path bounds (see Section 2.2.1) were used as well (see [145] for an overview). These results are then used in [148] to develop and compare several admission control schemes. Related results are derived in [98] for a multiple-class GPS system. Under the assumption that the input process satisfies a large-deviations principle in the space of cadlag functions (right-continuous with left-limits), logarithmic large-buffer asymptotics are calculated.

In [12] the two-queue GPS model of [145] is considered in discrete time. The logarithmic large-buffer asymptotics are calculated for a system where the service capacity is a random variable. In addition, it is shown that the set of possible overflow paths can be reduced to two scenarios. We will show in Chapter 7 that it is a complicated task to determine under which conditions each of the two scenarios occurs.

In [56] the GPS model is considered from a different perspective. By reformulating the multiple-class GPS model in terms of a Skorokhod problem, the authors show that the mapping from the input processes to the buffer content exists, is unique and is continuous. This justifies the use of fluid inputs with GPS, and the use of fluid inputs as an approximation for discrete inputs.

The performance analysis of a multiple-class GPS system with service rate c, and the input of one of the classes, flow i, having heavy-tailed characteristics, is initiated in [16]. Assuming that the average input rates of the classes are smaller than the GPS weights, the exact large-buffer asymptotics for flow i are shown to be equal to those in a system where flow i is served in isolation at constant rate $c - \rho_{-i}$. This result is known as a so-called *reduced-load equivalence*, as explained and observed for other models in [6, 67]. The assumption that $\rho_j < \phi_j c$ for each class j is sufficient but not necessary for class j to be stable. In [18] a necessary and sufficient condition is established for the stability of a class and it is shown that the exact large-buffer asymptotics as derived in [16] continue to hold under further statistical assumptions. In case these assumptions do not hold, flow i may be strongly affected by other flows with heavier tails, see for instance [17]. Extensions of the results in [16, 17, 18] can be found in [21]. The network model in Chapter 3 is an extension of the model in [16].

In [77] the order of the large-buffer asymptotics is calculated for a two-class GPS system in continuous time, where both queues are fed by an M/G/∞ process (see Chapter 6 and [110] for more details on such input processes). The goal of [78] is to obtain many-sources asymptotics for GPS systems. Bounds on the many-sources asymptotics are derived for a discrete-time GPS system fed by a large number of sources that have a finite moment generating function. It extends the results in [79], where a GPS system with only two queues is studied.

Recently in [83] results have been proved for the instantaneous burst input model in [21] under weaker distributional assumptions. The exact large-buffer asymptotics are calculated for the workload of an individual queue with a subexponential input process (see Section 2.1 for a definition of subexponentiality), or an input process that is lighter than subexponential and interacts with processes that are dominatedly varying (a subclass of subexponential distributions see [73]).

In [96] the study of a multiple-class GPS system with a large number of Gaussian inputs is initiated. The authors obtain the decay rate and the optimal path of the total workload exceeding a threshold. They give some heuristics for the workload behavior of the individual classes in a two-class GPS system. We will extend their results in Chapter 7.

So far, not much work has appeared on the topic of the choice of the GPS weights. Under the assumption of finite log-moment generating functions, the Chernoff large-deviations approximation is used to calculate the loss probabilities in [60]. Weights are then chosen such that the loss probability of the sources remains below a predefined threshold. In [80] some algorithms are developed for on-line weight adaptation.

1.7 Related disciplines

This monograph not only considers the performance of the GPS discipline, but also studies two related models, the *strict-priority* mechanism and a model where two servers operate in a coupled manner, referred to as the *coupled-processors* model. As briefly mentioned in Section 1.1.1, the two-class strict-priority mechanism can be viewed as a special, somewhat degenerate, case of GPS, taking 0 for the weight of the low-priority class and 1 for the weight of the high-priority class. In this respect one could argue that studying the GPS mechanism makes a direct analysis of the priority mechanism superfluous. However, analyzing the workload behavior in a GPS system is typically hard, especially that of a class whose traffic intensity exceeds its GPS weight, whereas a direct analysis for the low-priority class in the two-class priority model can be simpler. In the priority system, the high-priority class does not notice the presence of the low-priority class, so both the total workload and the workload of the high-priority class are known. Taking their difference, we obtain the workload of the low-priority class. This approach does not work in the GPS system, where only the total workload is known, because both classes mutually influence each other. In the last two chapters of this monograph we first consider the logarithmic many-sources asymptotic of the two-queue priority mechanism and then that of the GPS discipline. Finally, we mention that the two-class priority mechanism has some useful similarities with the two-node tandem model. We will exploit these similarities in Chapter 6.

The priority model has been studied extensively in the literature. We briefly mention some work that is related to our work in Chapter 6. In [49] bounds on the exact many-sources asymptotics are calculated for a two-queue strict-priority model under some technical assumptions. Bounds on the logarithmic many-sources asymptotics are obtained in [139]. Assuming the input traffic to have Gaussian characteristics, heuristics on the logarithmic many-sources asymptotics are given in [95].

The two-class GPS system with instantaneous input can be seen as a special case of the two-class coupled-processors model that is analyzed in Chapter 5. The two-class coupled-processors model can be described as follows. Consider two classes, both with their own queue *and* their own server. Whenever there is traffic of both classes, each server serves its own queue at speed 1. However, the speed of server i increases to $r_i > 1$ when the other server is idle, $i = 1, 2$. Suppose we scale the input process in the coupled-processors model of class i by $\phi_i c$ such that we have input process $A_i(\cdot,\cdot)/(\phi_i c)$ in queue i, and we choose $r_i = 1/\phi_i$, such that $1/r_1 + 1/r_2 = 1$. Then this coupled-processors model is equivalent to a two-class GPS model, with weights ϕ_i assigned to class i, input process $A_i(\cdot,\cdot)$, and total service rate c. For instance, the two-class coupled-processors model with $r_1 = r_2 = 2$ corresponds to the two-class GPS model with $\phi_1 = \phi_2 = 0.5$ and $c = 2$.

We review some literature on models with two coupled processors. This model was first studied in [61, 75] for the case of exponentially distributed job sizes. In [42] the Laplace Stieltjes transform of the joint distribution of the workload was derived for the case of generally distributed job sizes. This result was used in [17] to derive the exact large-buffer asymptotics for the case where at least one of the queues has customers with heavy-tailed service requirements. An extension can be found in [24]. The approach of [42] was taken in [65] to derive Laplace Stieltjes transforms for the joint queue length distribution in a two-class GPS system, but no explicit results were found. Recently the heavy-traffic approximation for two coupled M/M/1 queues was derived in [74].

Unfortunately, analyzing the coupled-processors model with more than two classes turns out to be extremely hard. Only partial results on the Laplace Stieltjes transform of the joint distribution in a three-class coupled-processors model are obtained in [41]. In this model class i, $i = 1, 2, 3$, is served at rate 1 if all queues are backlogged, class i is served at rate $r_{i,j}$ if only queue j, $j \neq i$, is idle, and class i is served at rate $r_{i,j,k}$ if both other queues are idle. Again it is easy to scale this model such that a three-class GPS model with weights ϕ_i, input process $A_i(\cdot,\cdot)$ and total service rate c can be seen as a special case. Take in the three-class coupled-processors model input processes $A_i(\cdot,\cdot)/(\phi_i c)$, $r_{1,3} = r_{2,3} = 1/(\phi_1+\phi_2)$, $r_{1,2} = r_{3,2} = 1/(\phi_1+\phi_3)$, $r_{2,1} = r_{3,1} = 1/(\phi_2+\phi_3)$, and $r_{i,j,k} = 1/\phi_i$, such that $1/r_{1,j,k} + 1/r_{2,j,k} + 1/r_{3,j,k} = 1$. In [41] the analysis is simplified by assuming all $r_{i,j}$ to be 1, but under this assumption the three-class GPS system is not included as a special case. Hence, it is difficult to obtain results for multiple-class GPS systems by analyzing multiple-class coupled-processor models.

Also, the analysis of the two-class coupled-processors model with on-off inputs or Markov modulated fluid inputs is complicated, so GPS models with such input processes cannot be easily solved using coupled-processors results. Moreover, the analysis for $1/r_1 + 1/r_2 = 1$ is typically a non-straightforward extension of that for $1/r_1 + 1/r_2 \neq 1$, which is already hard. In this monograph we have used different tools to study both models, see Section 2.2 and Chapters 4 and 5.

1.8 Outline

In Chapter 2 we give a brief introduction of the mathematical concepts that are used in this monograph. We start with a description of the concepts of self-similarity and LRD. Then we introduce the class of heavy-tailed distributions and explain the relation between heavy-tailed distributions and the notions of self-similarity and LRD. The second part of this chapter is devoted to a brief description of the methods that have been used in this monograph to analyze the various models.

In Chapter 3 we consider various feed-forward networks of GPS nodes. We derive the exact large-buffer asymptotics for the workload distribution of a particular flow with heavy-tailed characteristics, at the bottleneck node on its path through the network. We show that the asymptotic behavior of its workload distribution is equivalent to that in a two-node tandem network, where it is served in isolation at *constant rates*. The protective characteristics of GPS are reflected in these rates, as they only contain the average rates of the other, possibly heavier-tailed, flows. This chapter is completely based on [134]; a short version has appeared in [133].

In Chapter 4 we consider a single GPS node, which is fed by a flow with light-tailed characteristics and a flow with heavy-tailed properties. We derive exact large-buffer asymptotics for the light-tailed flow, under the assumption that it generates traffic according to a Markov-modulated fluid process. For the heavy-tailed flow we consider three cases: (i) it generates instantaneous bursts with a regularly varying distribution, (ii) it behaves as an on-off process with regularly varying on times, with the peak rate smaller than $1 - \rho_1$, and (iii) it behaves as an on-off process with regularly varying on times, with the peak rate larger than $1 - \rho_1$. This chapter is completely based on [28]. A summary of this work is published in [26]. The results for the case where flow 2 generates instantaneous bursts are presented in [27] and [29].

In Chapter 5 we consider a model that is closely related to that of Chapter 4: the two-queue coupled-processors model. We assume one class to exhibit heavy-tailed properties, and the other class to have either heavy-tailed or light-tailed characteristics. We calculate the exact large-buffer asymptotics for both classes under different assumptions on the average rates ρ_j relative to the service rate 1. The results for the scenario with one class having light-tailed properties qualitatively resemble those of the previous section when $1/r_1 + 1/r_2 \neq 1$. We show that the results exactly agree for the case $1/r_1 + 1/r_2 = 1$. This chapter is based on the results in [23] and [24]. Note that the results in Sections 5.3 and 5.5 have already appeared in Section 6 of [20]. We include these results and their proofs because the proofs of the other results either use them or closely resemble them.

In Chapter 6 we consider two models: a two-node tandem network and a two-class priority system. We assume the input traffic to consist of a large number of Gaussian processes. Using Schilder's sample-path large-deviations principle, we calculate the logarithmic many-sources asymptotics for the second node in the tandem model, and for the low-priority class in the priority model. This chapter is completely based on [90]. The results on the tandem network are published in [93].

In Chapter 7 we extend the model of Chapter 6 to a two-class GPS model. We derive an upper and a lower bound on the logarithmic many-sources asymptotics of the workload distribution. Again we make use of Schilder's sample-path large-deviations principle. Based on extensive numerical work, we conjecture the exact value of the asymptotic decay rate. We conclude the chapter with some weight-setting procedures. The content of this chapter is completely based on [92].

CHAPTER 2

Heavy tails and methods

In the previous chapter we mentioned that phenomena such as LRD and self-similarity have been observed in communication networks. We use the first section in this chapter to give a more formal elaboration on these concepts. In the second section we sketch the methods that we use in this monograph to calculate the desired performance measures.

2.1 Heavy tails, self-similarity and long-range dependence

Although the terms heavy tails, self-similarity and long-range dependence are connected and often used interchangeably, they refer to fundamentally different concepts. In this section we give definitions of each of these notions and explain how they are related. We refer the interested reader to [52, 109, 150] for more background on these concepts in relation to network traffic, and to [122] for additional mathematical details. We end this section with a short discussion on the practical relevance of these phenomena for network performance.

2.1.1 Self-similarity and long-range dependence

Self-similarity. A real-valued *continuous-time* stochastic process $Z = \{Z(t), t \geq 0\}$ is *self-similar* with Hurst parameter H, $H \in (0, 1]$, if

$$a^{-H} Z(at) \stackrel{\mathrm{d}}{=} Z(t), \ \forall a > 0.$$

We use the notation $\stackrel{\mathrm{d}}{=}$ here to indicate that the processes have the same finite dimensional distributions. Observe that this definition implies that the process Z is invariant under appropriate rescaling. Moreover, a process Z with this property can never be stationary (apart from the degenerate case with $Z(t) = 0$ for all t). We are typically interested in self-similar processes Z that have *stationary increments.*

Remark 2.1.1 *A weaker version of self-similarity is* asymptotic second-order self-similarity*, see [109] for a detailed explanation.*

Long-range dependence. Let $Y = \{Y(t), t \geq 0\}$ be a real-valued *stationary* continuous-time stochastic process with finite variance and autocorrelation function $\gamma(t)$, i.e.,

$$\gamma(t) = \frac{\mathbb{C}\mathrm{ov}[Y(s), Y(s+t)]}{\mathbb{V}\mathrm{ar}[Y(s)]}, \quad t \geq 0.$$

We say that the process Y exhibits *LRD* if the autocorrelation function is not integrable, i.e.,

$$\int_0^\infty |\gamma(t)| \mathrm{d}t = \infty, \tag{2.1}$$

and that Y is *SRD* if

$$\int_0^\infty |\gamma(t)| \mathrm{d}t < \infty.$$

The following definition of LRD is used in the literature as well. It describes the rate at which $\gamma(t)$ decreases to zero as $t \to \infty$. The process Y exhibits LRD if

$$\gamma(t) \sim c_\gamma t^{-\beta}, \quad t \to \infty, \tag{2.2}$$

with c_γ a positive constant and $\beta \in (0, 1)$. It is easily seen that (2.2) indeed implies (2.1).

It is worth mentioning that one should be very cautious when testing a process on LRD, as only *stationary* processes can exhibit LRD. Slowly decreasing autocorrelations might be caused by non-stationarities instead of LRD.

As mentioned in the beginning of this section, the notions of self-similarity and LRD are sometimes used synonymously. Obviously they are not, as only stationary processes can exhibit LRD, and self-similar processes are by definition not stationary. We explain the relation between LRD and self-similarity using some specific examples.

Fractional Brownian motion. We start with the class of fBm, which is an important class of self-similar processes. The process $B_H = \{B_H(t), t \in \mathbb{R}\}$ with Hurst parameter $H \in (0, 1]$ is standard fBm if it has stationary increments, $B_H(0) = 0$, $\mathbb{E}[B_H(t)] = 0$, and $\mathbb{V}\mathrm{ar}[B_H(t)] = |t|^{2H}$ for all t, and all finite-dimensional distributions are multivariate Gaussian. It follows from this definition that B_H has continuous sample paths with probability (w.p.) 1, that B_H has non-differentiable sample paths for $H \in (0, 1)$ w.p. 1, and that $B_H(t)$ is a Gaussian or normally distributed random variable for all t. Its discrete-time increment process $X = \{X_n = B_H(n) - B_H(n-1), n \geq 1\}$ is referred to as fractional Gaussian noise, which is a stationary process with $\mathbb{E}[X_n] = 0$ and $\mathbb{V}\mathrm{ar}[X_n] = 1$, for all $n \geq 1$ and $H \in (0, 1]$. Using a second-order Taylor approximation, it is easily seen that for $H \in (\frac{1}{2}, 1)$, the discrete-time autocorrelation function of X behaves as

$$\gamma(t) \sim H(2H-1)t^{2H-2}, \quad t \to \infty.$$

Hence, X exhibits LRD for $H \in (\frac{1}{2}, 1)$, whereas for $H \in (0, \frac{1}{2}]$, X exhibits SRD. For the special case of $H = \frac{1}{2}$ the process B_H reduces to ordinary Brownian motion (and thus has independent increments). The case $H = 1$ gives $B_1(t) = tB_1(1)$. Note that we abuse the terminology LRD in the sense that we often refer to fBm as exhibiting LRD, whereas only its discrete-time increment process exhibits LRD.

Lévy process. Another frequently-used example of a self-similar process is the so-called α-stable Lévy motion (see [122]), which has stationary *independent* increments. Hence, its increment process does not exhibit LRD.

2.1.2 Heavy-tailed distributions

Before explaining the relation of heavy-tailed distributions with self-similarity and long-range dependence, we first give some definitions that will frequently be used in this monograph. We refer to Chapter 2 in [52, 150] for a more thorough review of properties and intuitive insights concerning heavy-tailed distributions. See also [127] for a short survey on heavy-tailed distributions. Recall the notation $g(t) \sim h(t)$ as $t \to \infty$, which denotes $\lim_{t\to\infty} g(t)/h(t) = 1$. Let $F^{n\star}(\cdot)$ denote the n-fold convolution of $F(\cdot)$, i.e., for example $F^{2\star}(x) = \int_0^x F(x-y)F(\mathrm{d}y)$.

Definition 2.1.1 (Subexponential) *A distribution function $F(\cdot)$ on $[0, \infty)$ is called subexponential if, for some $n \geq 2$,*

$$1 - F^{n\star}(x) \sim n(1 - F(x)), \quad \textit{as } x \to \infty.$$

We use $\mathcal{S}$ to denote the class of subexponential distribution functions. For a random variable X with distribution function $F(\cdot)$, we use both $X \in \mathcal{S}$ and $F(\cdot) \in \mathcal{S}$ to indicate that $F(\cdot)$ is a subexponential distribution function. We also say that X is subexponential. In order to provide some intuition on subexponentiality, we rewrite the above definition as

$$\mathbb{P}(X_1 + \ldots + X_n > x) \sim n\mathbb{P}(X > x),$$

with X_i, $i = 1, \ldots, n$, i.i.d. distributed as $X \in \mathcal{S}$. Equivalently, we can restate this property as

$$\mathbb{P}(X_1 + \ldots + X_n > x) \sim \mathbb{P}(\max\{X_1, \ldots, X_n\} > x).$$

Hence, a sum of subexponential random variables most likely exceeds a large level x, because *one* of the random variables does so. This is in stark contrast with the sum of light-tailed random variables. Consider for instance the case where X follows an exponential distribution. Then we have

$$\mathbb{P}(X_1 + \ldots + X_n > x) = \sum_{j=0}^{n-1} e^{-\lambda x} \frac{(\lambda x)^j}{j!}, \quad x \geq 0,$$

reflecting that each random variable contributes.

A relevant subclass of $\mathcal{S}$ is the class $\mathcal{R}$ of regularly varying distributions.

Definition 2.1.2 (Regularly varying) *A distribution function $F(\cdot)$ is called regularly varying with index $-\nu$, $\nu \geq 0$ ($F \in \mathcal{R}_{-\nu}$) if*

$$1 - F(x) = l(x)x^{-\nu}, \quad x > 0,$$

where $l(\cdot) : \mathbb{R}_+ \to \mathbb{R}_+$ is a function of slow variation, i.e., $l(\eta x) \sim l(x)$ as $x \to \infty$ for $\eta > 0$.

A technical extension of $\mathcal{R}$ is the class $\mathcal{IR}$ of intermediately regularly varying distributions. Quite often, the class of $\mathcal{IR}$ emerges as a natural sufficient condition from a probabilistic proof geared towards the class $\mathcal{R}$, see for instance Chapter 3, but sometimes the extension is not immediate, as in Chapter 4 for instance.

Definition 2.1.3 (Intermediately regularly varying) *A distribution function $F(\cdot)$ on $[0, \infty)$ is called intermediately regularly varying if*

$$\lim_{\eta \uparrow 1} \limsup_{x \to \infty} \frac{1 - F(\eta x)}{1 - F(x)} = 1.$$

A well-known distribution in $\mathcal{R}$ is the Pareto distribution. Examples of subexponential distributions that do not belong to $\mathcal{IR}$ are the Weibull and lognormal distributions.

Unfortunately there is no unique definition of a heavy-tailed distribution in the literature. Sometimes $F(\cdot)$ is said to be heavy-tailed if it decays polynomially and has infinite variance. Another common definition is that $F(\cdot)$ is heavy-tailed if it decays non-exponentially, i.e., if there does not exist any constant $\eta > 0$ such that

$$1 - F(x) \sim \mathrm{O}(\mathrm{e}^{-\eta x}), \quad \text{as } x \to \infty.$$

In this monograph, we simply use the term heavy-tailed distribution to indicate that it belongs to either $\mathcal{S}, \mathcal{R}$ or $\mathcal{IR}$.

2.1.3 Relating the concepts

Now we consider the relation between heavy-tailed distributions and the phenomena of self-similarity and LRD. It turns out that heavy-tailed distributions frequently lie at the root of these phenomena. In [66] (see also [35, 128]) it is shown for instance that a single stationary on-off flow with on-period distribution and off-period distribution denoted by F_{on} and F_{off} respectively, exhibits LRD if either

$$F_{\mathrm{on}}(x) \sim \eta_1 x^{-\nu_1} \text{ as } x \to \infty,$$

or

$$F_{\mathrm{off}}(x) \sim \eta_2 x^{-\nu_2} \text{ as } x \to \infty,$$

with $\nu_1, \nu_2 \in (1, 2)$ and $\eta_1, \eta_2 > 0$ constant. Note that such an on-off process is clearly not self-similar. In [138] the superposition of such on-off processes is considered, when appropriately rescaled in space, and rescaled in time with constant T. First letting the number of processes go to infinity, and then taking the limit for $T \to \infty$, the superposition weakly converges to fBm with $H = (3 - \min\{\nu_1, \nu_2\})/2$. Hence, the rescaled superposition is self-similar and its increment process exhibits LRD. The formal proof of this result can be found in [131], but appeared to contain a small gap, which was filled in Theorem 7.2.5 in [136]. When taking the limits in reversed order, i.e., first taking the limit for $T \to \infty$, and then letting the number of processes go to infinity, it is shown in [131] that the finite-dimensional distributions converge to those of an α-stable Lévy motion, which is self-similar but does not exhibit LRD. In [128], also the limiting process is calculated when both the number of sources and time go to infinity together, i.e., the number of sources is assumed to be a function of time. This result is extended in [99], where the rate of growth of the number of sources appears to play a crucial role. If the number of sources grows "slowly" with respect to time, then the Lévy motion arises in the limit (only convergence of finite-dimensional distributions) and otherwise fBm (weak convergence).

2.1.4 Using Gaussian processes

In Chapters 6 and 7 we assume the amount of traffic arriving in an interval $(s,t]$, $A(s,t) := A(t) - A(s)$, to follow a Gaussian distribution, with mean input rate μ, i.e., $\mathbb{E}A(s,t) = \mu(t-s)$, and variance function $v(\cdot)$, i.e, $\mathbb{V}\mathrm{ar}A(s,t) = v(t-s)$. Observe that the mean $\mathbb{E}A(s,t)$ and variance $\mathbb{V}\mathrm{ar}A(s,t)$ depend on the interval length $t - s$ only. Hence, the process $A := \{A(t), t \geq 0\}$ is not necessarily stationary, but the increment process *is*.

In principle, the process A can be endowed with any correlation structure through the variance function $v(\cdot)$. For instance, taking $v(t) = t^{2H}$ with $H \in (\frac{1}{2}, 1)$ the input process A reduces to fBm, a self-similar process.

It is well-known that this Gaussian framework can be used to approximate other input models. The Gaussian counterpart of an M/G/∞ input process is found by calculating the mean input rate and the variance function of the amount of traffic generated by the M/G/∞ model, μ_∞ and $v_\infty(\cdot)$, and using these in the Gaussian distribution. We note that the M/G/∞ input model is especially nice, since both light-tailed and heavy-tailed distributions can be used for the service time.

2.1.5 Performance impact

Ever since the discovery of heavy-tailed or non-stationary traffic characteristics, both leading to LRD in measurement data, there has been an intense debate about the relevance of these phenomena for network performance. Conclusions are drawn ranging from "heavy tails have a severe impact" to "Markovian models are still suitable for engineering purposes". In a certain sense, both statements are correct, and it depends completely on the kind of issues that one is interested in, to which philosophy one should adhere. Both [64] and [121] argue that, in the framework of logarithmic asymptotics, there exists a critical time scale beyond which the correlation (LRD) in the traffic process does not influence the overflow probabilities. See Chapter 5 in [15] for a more elaborate discussion. The results in [37] demonstrate that it is important to carefully choose the scaling of the parameters in the model when trying to investigate the impact of LRD. They consider overflow probabilities in a large-deviations framework (many sources) when the buffer is not scaled with the number of sources, as they argue that it is not likely that in the future the buffer capacity will grow proportionally to the link speed. Not surprisingly, their results are very different from the results in [86] where the buffer *is* scaled with the link speed. Finally we mention [8] which shows the effect of feedback mechanisms like TCP, which prevent buffer contents from getting too large or having a heavy-tailed distribution. We remark that the transfer times of files with heavy-tailed sizes will still have heavy-tailed characteristics when transmitted with TCP. It is not clear however whether these heavy-tailed properties can adversely affect the transfer times of other, non-heavy-tailed, files, or files of a given size.

2.2 Methods

In this section, we give a description of the methods that we have used to obtain our results. These methods can roughly be divided into three categories: (i) in Chapters 3 and 4 we exploit a probabilistic sample-path approach, (ii) in Chapter 5 we use Tauberian theorems, and (iii) in Chapters 6 and 7 we apply a sample-path large-deviations principle. Below we explain the basic ideas of these methods, and introduce some preliminaries. More detailed explanations of (i) and (ii) can be found in [22, 150]. We refer to [135] for a general introduction on many-sources asymptotics and to [94, 103] for a detailed explanation of (iii).

2.2.1 Probabilistic sample-path approach

We start with the method that we have used to calculate exact large-buffer asymptotics for the workload in a GPS network (Chapter 3), and the workload in a 2-class GPS system with both heavy-tailed and light-tailed input (Chapter 4). In fact, it is not a method in the sense that there is a systematic recipe or generic theorem available. It is rather a special way of looking at the system and identifying the most important aspects that determine the probability under consideration. At first sight it is therefore hard to recognize the similarities between the proofs in Chapters 3 and 4, as they only manifest themselves on a more intuitive level.

The method is very powerful in that complex systems like multiple-class GPS systems can be handled. A difficulty could be that it requires some a priori intuition on the way a rare event, like an extremely high workload, occurs. Applying the method typically reinforces the intuition, and provides a more detailed interpretation of the rare-event scenario. A nice feature of traffic flows with heavy-tailed characteristics is that rare-event scenarios are often remarkably simple. For example the workload in a queue fed by a single on-off flow with heavy-tailed on periods most likely gets large due to a single event: one long on period of the flow (see [67]). If the queue is fed by a single flow that generates instantaneous heavy-tailed traffic bursts, then the workload gets large because of a single large burst (see [106]).

Mathematically, the method is relatively simple, and usually involves only first principles. The complexity of the method rather lies in the question how to show that the overflow scenarios other than that responsible for the asymptotic behavior, can be asymptotically neglected. As an example (see [16]), we consider the GPS system as introduced in Section 1.3. We assume one of the flows, say i, to send instantaneous traffic bursts with heavy-tailed characteristics. We also make the assumption that for all flows j we have $\rho_j < \phi_j c$, which guarantees stability for all flows. Recall that we are interested in $\mathbb{P}(V_i(t) > x)$ for $x \to \infty$. A major difficulty of the workload analysis in GPS systems arises immediately when representing the workload $V_i(t)$ in terms of arrival processes, weights and the service rate, without using workloads $V_j(s)$ of flows j at time instants $s < t$. It is impossible to give such an expression. Therefore, we have to work with bounds for the process $V_i(t)$, which are expressed in terms of the arrival processes, weights and service rate only. The ultimate goal is to find bounds that coincide when calculating the probabilities $\mathbb{P}(V_i(t) > x)$ for $x \to \infty$.

As mentioned above, the method relies on a priori insight into the dominant overflow scenario. From [106, 7] we know that the workload of flow i gets large in isolation due to a single large burst, a scenario that is also sufficient for flow i to reach a large workload level in the GPS system. This leads to the following overflow scenario. Flow i generates a very large burst, while the other flows show average behavior, leaving $c - \rho_{-i}$ for flow i, i.e.,

$$\mathbb{P}(V_i(t) > x) \sim \mathbb{P}(V_i^{c-\rho_{-i}}(t) > x), \quad x \to \infty.$$

Hence, the leading term in both the upper and lower bound for $V_i(t)$ should be the workload in a system where flow i is served in isolation at rate $c - \rho_{-i}$, i.e., $V_i^{c-\rho_{-i}}(t)$.

The next step is to bound $V_i(t)$ from above and from below by $V_i^{c-\rho_{-i}}(t)$ and some additional terms that account for scenarios that occur with a smaller probability than the most likely one. The derivation of the upper bound roughly goes as follows. Obviously, we have

$$V_i(t) \leq \sum_{j=1}^{m} V_j(t) = \sup_{s \leq t} \left\{ \sum_{j=1}^{n} A_j(s,t) - c(t-s) \right\}.$$

Now using a basic property of the sup operator,

$$\begin{aligned} V_i(t) &\leq \sup_{s \leq t} \{A_i(s,t) - (c - \rho_{-i} - \delta)(t-s)\} \\ &\quad + \sup_{s \leq t} \left\{ \sum_{j \neq i} A_j(s,t) - (\rho_{-i} + \delta)(t-s) \right\} \\ &= V_i^{c-\rho_{-i}-\delta}(t) + \sup_{s \leq t} \left\{ \sum_{j \neq i} A_j(s,t) - (\rho_{-i} + \delta)(t-s) \right\}, \end{aligned}$$

with $\delta > 0$ a small constant. For the lower bound the following trick applies. It holds that

$$\begin{aligned} V_i(t) &= \sum_{j=1}^{m} V_j(t) - \sum_{j\neq i} V_j(t) \\ &= \sup_{s\leq t}\left\{\sum_{j=1}^{n} A_j(s,t) - c(t-s)\right\} - \sum_{j\neq i} V_j(t). \end{aligned}$$

Using again a basic property of the sup operator and the trivial GPS inequality as given in (1.5),

$$\begin{aligned} V_i(t) &\geq \sup_{s\leq t}\{A_i(s,t) - (c - \rho_{-i} + \delta)(t-s)\} \\ &\quad + \inf_{s\leq t}\left\{\sum_{j\neq i} A_j(s,t) - (\rho_{-i} - \delta)(t-s)\right\} - \sum_{j\neq i} V_j^{\phi_j c}(t) \\ &= V_i^{c-\rho_{-i}+\delta}(t) + \inf_{s\leq t}\left\{\sum_{j\neq i} A_j(s,t) - (\rho_{-i} - \delta)(t-s)\right\} - \sum_{j\neq i} V_j^{\phi_j c}(t), \end{aligned}$$

with again $\delta > 0$ a small constant. Indeed, both the lower and upper bound now contain $V_i^{c-\rho_{-i}}(t)$ up to some small perturbation δ. The last step is to show that the other terms can be neglected when combining this upper bound with the trivial GPS inequality (1.5). This is straightforward probabilistic calculus.

Above we found that the asymptotic behavior of the workload in an individual queue in the GPS system is equivalent to that in an isolated system with a constant service rate. This holds in far greater generality as we will show in Chapters 3 and 4. In Chapter 3 the asymptotic behavior of the workload of flow i in a general feed-forward GPS network is related to that of the workload of flow i in a *two-node* tandem network, where it is served in isolation at constant rates. In Chapter 4 the asymptotic behavior of the workload of a light-tailed flow (flow 1), when served according to GPS together with a flow generating instantaneous heavy-tailed bursts (flow 2), is shown to be (assume $c = 1$),

$$\mathbb{P}(V_1 > x) \sim \mathbb{P}(V_1^{\phi_1} > x) C_7 \mathbb{P}(B_2^{\mathrm{r}} > C_8 x), \quad x \to \infty,$$

with C_7, C_8 positive finite constants that depend on the parameters of flow 2, and B_2 the burst size of flow 2. Note that for any non-negative random variable X with $\mathbb{E}X < \infty$, we denote by X^{r} a random variable representing the residual lifetime of X, i.e., $\mathbb{P}(X^{\mathrm{r}} > x) = (1/\mathbb{E}X)\int_x^\infty \mathbb{P}(X > u)\mathrm{d}u$. Hence, the workload of the light-tailed flow gets large due to a combination of two rare events, reflected by the product of the probabilities. (i) The light-tailed flow itself shows similar rare-event behavior as in a system with service rate ϕ_i where it is served in isolation, and (ii) the heavy-tailed flow generates a single large burst. Note that the proofs in these chapters are similar to the proof that we illustrated above, but considerably more involved.

2.2.2 Tauberian approach

In Chapter 5 we derive exact large-buffer asymptotics using a Tauberian approach. This approach relies on the relationship between the asymptotic behavior of a distribution function and the behavior of its Laplace (Stieltjes) transform near the origin or the poles. Obviously, such a method requires a rather explicit expression for the Laplace (Stieltjes) transform, which is not always available. For the case of two coupled queues, the Laplace Stieltjes transform (see [42]) has a manageable format, but for three coupled queues the expression is already exceedingly complicated (see [41]).

Like the probabilistic method described in the above section, this method requires some a priori insight into the form of the results as well. Depending on this form, we can choose one of the two theorems presented below, that relate the Laplace (Stieltjes) transform of a distribution function to its asymptotic behavior. The first is presented in Lemma 2.2 in [34] as an extension of Theorem 8.1.6 in [14], which is due to [13]. It relates the behavior of $\mathbb{E}e^{-sY}$ for $s \downarrow 0$ to the behavior of $\mathbb{P}(Y > x)$ for $x \to \infty$ for a random variable Y with a regularly varying distribution.

Theorem 2.2.1 *Let Y be a non-negative random variable, $l(x)$ a slowly varying function, $\nu \in (n, n+1)$ for $n \in \mathbb{N}$, and $D \geq 0$. Then the following two statements are equivalent:*

(i) $\mathbb{P}(Y > x) = (D + o(1))x^{-\nu}l(x)$ *as* $x \to \infty$*;*

(ii) $\mathbb{E}[Y^n] < \infty$ *and*

$$\mathbb{E}e^{-sY} - \sum_{j=0}^{n} \frac{\mathbb{E}[Y^j](-s)^j}{j!} = (-1)^n \Gamma(1-\nu)(D + o(1))s^{\nu} l(1/s)$$

as $s \downarrow 0$.

Note that $\Gamma(\cdot)$ is the ordinary Gamma function. Conceptually, the procedure of using this theorem is fairly straightforward, although the technical details can be quite involved. Suppose the Laplace Stieltjes transform $\mathbb{E}e^{-sV_i}$ of the workload distribution $\mathbb{P}(V_i > x)$ is known or can be derived.Then calculating the series expansion of $\mathbb{E}e^{-sV_i}$ for $s \downarrow 0$ gives the exact asymptotics.

The second theorem is established in [130]. It relates the asymptotic behavior of a function $\psi(x)$ for $x \to \infty$ to that of its Laplace transform $\phi(s)$ for s near its poles.

Theorem 2.2.2 *Define for $d > 0$,*

$$\psi(x) = \frac{1}{2\pi i} \int_{d-i\infty}^{d+i\infty} e^{sx} \phi(s) ds,$$

where $s = z + iy$ and the path of integration is the straight line $z = d$, chosen so that $\phi(s)$ is analytic for $z \geq d$. If

(i) *$\phi(s)$ is analytic for $z \geq a - \delta$ $(\delta > 0)$, except at k points $s_1, \ldots, s_p, \ldots, s_k$ on $z = a$;*

(ii) *near each such point s_p, we have*

$$(s - s_p)\phi(s) = \sum_{n=0}^{\infty} a_{np}(s - s_p)^n + (s - s_p)^{\gamma_p} \sum_{n=0}^{\infty} b_{np}(s - s_p)^n,$$

where $0 < \gamma_p < 1$, and the series converge for $|s - s_p| < l$ $(l > 0)$;

(iii) *$\phi(s) \to 0$ as $y \to \pm\infty$, uniformly in z for $a - \delta \leq z \leq d$ $(d > a)$, and in such a manner that $\int |\phi(s)| \mathrm{d}y$ converges at $y = \pm\infty$,*

then,

$$\psi(x) \sim \sum_{p=1}^{k} \mathrm{e}^{s_p x} \left(a_{0p} + \frac{\sin(\pi\gamma_p)}{\pi} \sum_{n=0}^{\infty} (-1)^n b_{np} \Gamma(\gamma_p + n) x^{-\gamma_p - n} \right), \qquad x \to \infty.$$

The procedure of using the above theorem is similar to that of using Theorem 2.2.1, but slighty more involved. Calculating the series expansion of $\mathbb{E}\mathrm{e}^{-sV_i}$ for s around its poles gives the result. The above theorem is more general, in that it may be applied if the asymptotic behavior for $x \to \infty$ of $\mathbb{P}(V_i > x)$ is exponential (the constants b_{np} are 0), polynomial (the constant a_{0p} is 0) or a mixture of a function that is exponential in x and a function that is polynomial in x. However, the above theorem cannot be applied if the asymptotic behavior involves a slowly varying function. To the best of our knowledge, an extension to functions with a slowly varying component does not exist, but we expect that a similar theorem should hold that *does* incorporate slowly varying behavior.

2.2.3 Sample-path large deviations approach

In Chapters 6 and 7 we use a sample-path large-deviations principle to derive logarithmic many-sources asymptotics for the workload in a two-node tandem network, a priority model and a GPS system. As the results of the tandem network follow immediately from those of the priority model, and the priority model is a special case of the GPS system, we focus in this section on a description of the method that is used for the GPS system.

When deriving many-sources asymptotics, the service rate and the buffer threshold both have to be scaled by the number of sources n, i.e., we take service rate nc and buffer threshold nx. Recall that we are interested in the behavior of $(1/n) \log \mathbb{P}(V_i > nx)$ for $n \to \infty$. Like the method described in Section 2.2.1, this method too requires some intuition on the rare-event behavior. In a rare-event scenario the n flows will typically conspire.

Given that the rare event takes place, it is as if the amount of traffic that is sent by each of the flows is i.i.d. distributed according to a distribution function that differs from the actual distribution. We say that the rare-event distribution is a *twisted* version of the actual distribution, see e.g. [15, 139].

Because we let the number of flows go to infinity, intuitively the law of large numbers applies. Not surprisingly, sample-path large-deviations principles are stated in terms of the *sample mean* of the input processes, i.e., $(1/n)\sum_{i=1}^{n} A_i(t)$, where $A_i(t) = A_i(t,0)$ for $t < 0$ and $A_i(t) = A_i(0,t)$ for $t > 0$. The term *sample path* then refers to a realization of the sample mean process

$$\left\{\frac{1}{n}\sum_{i=1}^{n} A_i(t),\ t \in \mathbb{R}\right\}.$$

Intuitively, we can thus say that sample-path large deviations are concerned with deviations from the 'average sample path'.

A sample-path large-deviations principle is simply a tool that transforms the stochastic problem into a deterministic variational problem. Unfortunately, the variational problem can be a hard optimization problem, for which no analytical solution is available (as is the case in Chapter 7). The solution of this variational problem provides two important results. (i) It gives the value of $(1/n)\log \mathbb{P}(V_i > nx)$ for $n \to \infty$ in terms of a function that depends on x. (ii) It gives the most-probable sample path, i.e., given that the event $\{V_i > nx\}$ occurs, the sample mean will follow this most-probable path with probability 1, for $n \to \infty$.

We will make use of a sample-path large-deviations principle that is a generalization of Schilder's theorem, which is a powerful theorem for i.i.d. centered Gaussian processes. We define $\bar{A}_i(\cdot)$ to be the centered version of $A_i(\cdot)$, and $f(\cdot)$ to be the realization of $(1/n)\sum_{i=1}^{n} \bar{A}_i(\cdot)$. Denoting by Ω the total space of sample paths and by F the space of sample paths that imply the rare event, Schilder's theorem is as follows. Details can be found in Section 6.1.4.

Theorem 2.2.3 *(i) for any closed set $\bar{F} \subset \Omega$,*

$$\limsup_{n\to\infty} \frac{1}{n} \log \mathbb{P}\left(\frac{1}{n}\sum_{j=1}^{n} \bar{A}_j(0,\cdot) \in \bar{F}\right) \leq -\inf_{f\in\bar{F}} I(f);$$

(ii) for any open set $F^o \subset \Omega$,

$$\liminf_{n\to\infty} \frac{1}{n} \log \mathbb{P}\left(\frac{1}{n}\sum_{j=1}^{n} \bar{A}_j(0,\cdot) \in F^o\right) \geq -\inf_{f\in F^o} I(f).$$

The function $I(\cdot)$, which is a function of the sample paths, is known as the rate function. It measures how likely the occurrence of each of the sample paths in Ω is. The higher the value of the rate function, the less likely the particular sample path will be. Intuitively, we could view the value of $I(f)$ as the *cost* that is incurred for generating an average amount of traffic according to the process f. Using this interpretation, the variational problem of determining the most probable path then boils down to finding the cheapest path that induces the rare event.

It is not always a straightforward task to transform a queueing problem into the Schilder setting, i.e., to find an explicit set $F \subset \Omega$. For the tandem network and the priority model the set F can be found easily. However, an explicit expression for F in the GPS system in terms of arrival processes, weights and the service rate only, is not available. This is due to the fact that no such expression exists for the workload $V_i(t)$, as explained in Section 2.2.1. Hence, we have to follow the same approach, and calculate bounds for $\mathbb{P}(V_i > nx)$. Then we have to apply the sample-path large-deviations principle to both bounds in order to obtain the cheapest or most-probable path in the upper bound *and* in the lower bound. If these paths coincide, then we have found *the* most probable path for the overflow probability, and hence the asymptotic behavior of the overflow probability.

Chapter 3

GPS networks

While *single*-node models can provide useful insights into the GPS mechanism, in reality one normally encounters networks of multiple interacting nodes. When considering a network of GPS nodes, this interaction can obscure the results, and make it difficult to gain insights into the GPS mechanism. Therefore, it is important to initiate the investigation of GPS through single-node models, see e.g. [21], and to gradually increase the complexity of the model, once the model under consideration is fully understood. In this chapter, we present the results of [134] (a short version has appeared in [133]), and consider a network of nodes which are all equipped with the GPS mechanism. We start with a very simple network, and consider increasingly complex network scenarios in each subsequent section.

Existing work on GPS networks is largely restricted to a *deterministic* setting. In [108] it is shown that the GPS property of minimum guaranteed rates translates into worst-case bounds on delay and workload for leaky bucket controlled traffic flows. In a stochastic setting, the work-conserving property of GPS should yield statistical multiplexing gains. In order to quantify these gains however, and to examine how they are possibly influenced by the occurrence of heavy-tailed traffic, a *stochastic* analysis of GPS networks is required. Networks of fluid flows seem to defy exact analysis for all but a few specific cases, and in particular we are not familiar with any stochastic analysis of GPS networks. In [114] a FIFO network is studied, which is fed by fluid flows defined in terms of finite-state Markov processes. The buffer content distribution in a tandem queue fed by independent on-off flows with exponential off periods and generally distributed on periods is determined in [1].

In this chapter we specifically focus on GPS networks fed by several traffic flows, of which at least one has heavy-tailed traffic characteristics. Under certain conditions, we show that the tail distribution of the workload of the heavy-tailed flow at the bottleneck node on its path, is equivalent to that in a *two*-node tandem network where it is served in isolation at *constant* rates. These rates are the service rates of the nodes that have on average the least amount of service for the heavy-tailed flow in the original network. Hence, we show that the heavy-tailed flow is only affected by the traffic characteristics of the other flows through their average rates. Moreover, it is not influenced by excessive behavior of any of the other flows. This extends the results in [21] for a single GPS node fed by traffic with heavy-tailed characteristics. In [6] a similar *reduced-load equivalence* result is established for a fluid queue fed by a flow with subexponentially distributed on periods and a general light-tailed flow.

The remainder of this chapter is organized as follows. In Section 3.1 we consider a simple two-node tandem network, which is fed by a single flow. As alluded to above, this model will play a key role in analyzing more complex network configurations. We first relate the tail behavior of the busy-period distribution at node 1 to the arrival process. Then we determine the tail behavior of the workload distribution at the second node in terms of the residual busy-period distribution at node 1. Two traffic processes are considered: (i) a traffic flow generating instantaneous bursts and (ii) a traffic flow behaving according to an on-off process. Preliminaries will be given in Sections 3.2 and 3.6. In Sections 3.3 and 3.5 we extend the model in Section 3.1 to a GPS tandem network which is fed by multiple flows. We consider two network configurations: in Section 3.3 we assume that all flows which are served at node 1 proceed to node 2, while in Section 3.5 we allow for flows which are only served at node 1. In both sections, we determine an upper and a lower bound for the workload distribution of the ed flow at node 2. In Section 3.4 we prove a general lemma which shows that the lower and upper bounds for the workload distribution asymptotically coincide. We use this lemma to derive the asymptotics for the other models in this chapter as well. In the subsequent sections, we extend the analysis to more general GPS networks with the heavy-tailed flow traversing more than two nodes. In particular, in Sections 3.7 and 3.8 we consider extensions of the GPS networks in Sections 3.3 and 3.5 respectively. For both network configurations we determine an upper and a lower bound on the workload distribution of the heavy-tailed flow at the bottleneck node on its path in order to obtain the tail behavior.

3.1 Two-node tandem network fed by a single flow

In this section we consider a simple tandem network, which is fed by a single flow. We analyze the tail behavior of the workload distribution at the first and second node. Admittedly, this model represents the simplest possible network scenario, but it plays a central role in the further analysis. We need the results concerning the tail behavior of the workload distribution in this tandem network to analyze more general networks, where multiple flows share the capacity according to the GPS principle. Surprisingly, it turns out that in the GPS networks that we consider, the tail behavior of the workload distribution of an individual flow is equivalent to that in a tandem network where the flow is served in isolation at constant rates. We consider two traffic scenarios: (i) the flow generates instantaneous traffic bursts and (ii) the flow behaves according to an on-off process. In Sections 3.1.1 and 3.1.2 we give for both traffic scenarios the tail behavior of the busy-period distribution at node 1. In Section 3.1.3 we derive the tail behavior of the workload distribution at node 2 for both traffic scenarios using the busy-period characteristics at node 1.

First we introduce some notation. Because this chapter is notationally involved, we deviate from the notation introduced in Section 1.3. In this chapter we use $V(t)$ for workloads in the GPS model, like in Section 1.3, but $W(t)$ for workloads of flows served in isolation at a constant rate. Denote by c_1 and c_2 the constant service rates at node 1 and node 2, respectively. We assume $c_1 > c_2$ to exclude the trivial case where the workload at node 2 is always zero. We define ρ to be the traffic intensity, i.e., the mean amount of traffic offered to the network per unit of time. For stability, we assume $\rho < c_2$. Denote by $A(s,t)$ the amount of traffic generated during the time interval $(s,t]$. We define $W^c(t)$ to be the workload at time t if the flow were fed into a queue of rate c,

$$W^c(t) := \sup_{0\le s\le t} \{A(s,t) - c(t-s)\},$$

assuming $W^c(0) = 0$. Note that only in this chapter it is assumed that all systems considered are empty at time 0. For $c > \rho$, W^c is a random variable with the limiting distribution of $W^c(t)$ for $t \to \infty$. We define P to be the busy period in this queue. Observe that the total workload in the tandem network at time t is $W^{c_2}(t)$, because node 2 is the bottleneck node. The workload at node 1 is $W^{c_1}(t)$. Thus the workload at node 2 at time t is

$$\begin{aligned} W^{c_1,c_2}(t) &:= W^{c_2}(t) - W^{c_1}(t) \\ &= \sup_{0\le s\le t} \{A(s,t) - c_2(t-s)\} - \sup_{0\le s\le t} \{A(s,t) - c_1(t-s)\}, \qquad (3.1) \end{aligned}$$

assuming the system is empty at time 0. For $c_2 > \rho$, let W^{c_1,c_2} be a random variable with the limiting distribution of $W^{c_1,c_2}(t)$ for $t \to \infty$.

We now state some results for the distribution of the workload and the busy period at a single node. In Section 3.1.1 we consider the case where the flow generates instantaneous traffic bursts, and in Section 3.1.2 the case where the flow behaves according to an on-off process. We need these results to determine the asymptotic behavior of W^{c_1,c_2}, and later that of the workload in more general networks.

3.1.1 Instantaneous arrivals

Suppose the flow generates instantaneous traffic bursts according to a Poisson process with rate λ. Let K be the random variable representing the burst size. We assume that the burst size distribution $K(\cdot)$ is intermediately regularly varying (see Definition 2.1.3) with mean $\kappa < \infty$. The traffic intensity is $\rho = \lambda\kappa$. The following results play a crucial role in the analysis in subsequent sections. Theorem 3.1.1 is due to [106] and Theorem 3.1.2 is due to [151].

Theorem 3.1.1 *If $K^{\mathrm{r}}(\cdot) \in \mathcal{S}$ and $\rho < c$, then*

$$\mathbb{P}(W^c > x) \sim \frac{\rho}{c-\rho}\mathbb{P}(K^{\mathrm{r}} > x).$$

Recall that for a random variable X with distribution function $F(\cdot)$ we use X^{r} to denote its residual lifetime and $F^{\mathrm{r}}(\cdot)$ to denote the distribution function of X^{r}.

Theorem 3.1.2 *If $K(\cdot) \in \mathcal{IR}$ and $\rho < c$, then*

$$\mathbb{P}(P > x) \sim \frac{c}{c-\rho}\mathbb{P}(K > x(c-\rho)).$$

The above theorem immediately gives the tail distribution of the residual busy period.

Theorem 3.1.3 *If $K(\cdot) \in \mathcal{IR}$ and $\rho < c$, then*

$$\mathbb{P}(P^{\mathrm{r}} > x) \sim \frac{c}{c-\rho}\mathbb{P}(K^{\mathrm{r}} > x(c-\rho)).$$

Remark 3.1.1 *The assumption that $K(\cdot) \in \mathcal{IR}$ is in fact not necessary for Theorem 3.1.3 to hold. In [21] it is shown that the weaker condition $K^{\mathrm{r}}(\cdot) \in \mathcal{IR}$ is also sufficient. Since we only need Theorem 3.1.3 together with $K^{\mathrm{r}}(\cdot) \in \mathcal{IR}$ in our analysis, the results in this chapter still hold under this weaker condition.*

3.1.2 On-off processes

Suppose the flow generates traffic according to an on-off process. We assume the off periods to be exponentially distributed with mean $1/\lambda$. While on, the flow produces traffic at a constant rate r. Assume the random variable representing the on period K to have an intermediately regularly varying distribution with mean κ. Because the fraction of off-time is equal to $p = 1/(1+\lambda\kappa)$, the traffic intensity is equal to $\rho = \lambda\kappa r/(1+\lambda\kappa)$.

The following results are the analogues of Theorems 3.1.1, 3.1.2 and 3.1.3, respectively. Theorem 3.1.4 is due to [67] and Theorem 3.1.5 is due to [34, 151].

Theorem 3.1.4 *If $K^{\mathrm{r}}(\cdot) \in \mathcal{S}$ and $\rho < c < r$, then*

$$\mathbb{P}(W^c > x) \sim p\frac{\rho}{c-\rho}\mathbb{P}\left(K^{\mathrm{r}} > \frac{x}{r-c}\right).$$

Theorem 3.1.5 *If $K(\cdot) \in \mathcal{IR}$ and $\rho < c < r$, then*

$$\mathbb{P}(P > x) \sim p\frac{c}{c-\rho}\mathbb{P}\left(K > \frac{x(c-\rho)}{r-\rho}\right).$$

The following theorem immediately follows from Theorem 3.1.5.

Theorem 3.1.6 *If $K(\cdot) \in \mathcal{IR}$ and $\rho < c < r$, then*

$$\mathbb{P}(P^{\mathrm{r}} > x) \sim p\frac{c}{c-\rho}\mathbb{P}\left(K^{\mathrm{r}} > \frac{x(c-\rho)}{r-\rho}\right).$$

Remark 3.1.2 *Again the assumption $K(\cdot) \in \mathcal{IR}$ is sufficient but not necessary for the above theorem to hold. In [21] it is shown that the weaker condition $K^{\mathrm{r}}(\cdot) \in \mathcal{IR}$ is also sufficient, implying that under this weaker condition the results in this chapter still hold.*

3.1.3 Workload distribution

The above results completely specify the tail behavior of the workload distribution at node 1. Moreover, we can use them to analyze the workload distribution at node 2. Observe that the input process at node 2 is an on-off process with as on periods the busy periods at node 1. The on-rate is equal to the service rate at node 1, c_1. The off periods correspond to the idle periods at node 1, which are exponentially distributed. Because of this, the on and the off periods at node 2 are independent.

For both traffic scenarios the tail distribution of the residual busy period at node 1 is intermediately regularly varying. Hence, we can apply Theorem 3.1.4 to determine the tail behavior of the workload distribution at node 2, which is given in the following lemma. Note that if the interarrival times or the off periods at node 1 are non-exponential, then the on and the off periods at node 2 are not independent, so we cannot apply Theorem 3.1.4. Based on the intuitive explanation of this theorem however, we expect that the following lemma continues to hold under the mild form of dependence that would arise under non-exponential interarrival times or off periods. We conjecture that this can be formally shown through a probabilistic proof of the result in Theorem 3.1.4.

Lemma 3.1.1 *If $K(\cdot) \in \mathcal{IR}$, then*

$$\mathbb{P}\left(W^{c_1,c_2} > x\right) \sim p' \frac{\rho}{c_2 - \rho} \mathbb{P}\left(P^{\mathrm{r}} > \frac{x}{c_1 - c_2}\right),$$

with the fraction of off-time $p' = (c_1 - \rho)/c_1$.

In Section 3.4 we give our main theorem concerning the tail behavior of the workload distribution. In the proof of that theorem we need three properties which are satisfied for the two traffic scenarios that we described in the previous subsections. In the following lemma these properties are given.

Lemma 3.1.2 *For the traffic scenarios as described in Sections 3.1.1 and 3.1.2 the following three properties hold:*

(i) for α, β sufficiently small,

$$\lim_{x\to\infty} \frac{\mathbb{P}\left(W^{c_1+\alpha,c_2+\beta} > x\right)}{\mathbb{P}\left(W^{c_1,c_2} > x\right)} = G(\alpha,\beta), \text{ with } \lim_{\alpha,\beta\to 0} G(\alpha,\beta) = 1; \tag{3.2}$$

(ii) for any real y,

$$\lim_{x\to\infty} \frac{\mathbb{P}\left(W^{c_1,c_2} > x - y\right)}{\mathbb{P}\left(W^{c_1,c_2} > x\right)} = 1; \tag{3.3}$$

(iii) for each $c > \rho$ there exists a finite constant C such that,

$$\limsup_{x\to\infty} \frac{\mathbb{P}(W^c > x)}{\mathbb{P}\left(W^{c_1,c_2} > x\right)} = C < \infty. \tag{3.4}$$

Proof Theorems 3.1.3 (instantaneous arrivals), 3.1.6 (on-off processes) and Lemma 3.1.1 have to be used for all properties. In addition, we use for (ii) that $P^{\mathrm{r}}(\cdot) \in \mathcal{IR} \subset \mathcal{L}$ for both traffic scenarios. Finally, for (iii) we obtain, using Theorems 3.1.1 (instantaneous arrivals), 3.1.4 (on-off processes) and Lemma 3.1.1,

$$\limsup_{x\to\infty} \frac{\mathbb{P}(W^c > x)}{\mathbb{P}\left(W^{c_1,c_2} > x\right)} = \frac{g\frac{\rho}{c-\rho}}{\frac{c_1-\rho}{c_1}\frac{\rho}{c_2-\rho}} \limsup_{x\to\infty} \frac{\mathbb{P}(K^{\mathrm{r}} > \frac{x}{h})}{\mathbb{P}(P^{\mathrm{r}} > \frac{x}{c_1-c_2})},$$

with $g = 1$, $h = 1$ and K^{r} denoting the residual burst size for instantaneous arrivals, and $g = 1/(1+\lambda\kappa)$, $h = r - c$ and K^{r} denoting the residual on period for on-off processes. Because $K^{\mathrm{r}}(\cdot) \in \mathcal{IR}$, (3.4) follows. □

3.2 Preliminaries

In the next sections we extend the model which we described in the previous section. We consider again a two-node tandem network, but now fed by multiple flows, where traffic is scheduled according to the GPS mechanism. We focus on the workload distribution of a particular flow i which passes through both nodes. We assume this flow i to either generate instantaneous traffic bursts, precisely as described in Section 3.1.1, or to behave according to an on-off process, with the same characteristics as described in Section 3.1.2.

In this section we introduce the notation which we use throughout this chapter. Although the network that we consider in Sections 3.3 and 3.5 has only two nodes, we introduce notation for networks where flow i traverses N nodes. We conclude with two lemmas which we use in our analysis.

At each node of the network, traffic is served according to the GPS mechanism. Define c_n to be the service rate of node n and $S^{(n)}$ to be the set of all flows that receive service at node n, $n = 1, \ldots, N$. Each flow $q \in S^{(n)}$ is assigned a positive weight $\hat{\phi}_{q,n}$. Note that flows may have different weights at different nodes. In this chapter we do not make the assumption that the weights add up to 1. If every flow at node n is backlogged at time t, then flow $q \in S^{(n)}$ is served at node n at rate

$$\phi_{q,n} := \frac{\hat{\phi}_{q,n}}{\sum_{q\in S^{(n)}} \hat{\phi}_{q,n}} c_n.$$

Denote by $A_q(s,t)$ the amount of traffic generated by flow $q \in Q$ in the time interval $(s,t]$, and denote by $A_{q,n}(s,t)$ the amount of traffic that arrives at node n originating from flow q during $(s,t]$. In particular, $A_{q,n}(s,t) = A_q(s,t)$ if node n is the first node that flow q feeds into. We define

$$A_Q(s,t) := \sum_{q\in Q} A_q(s,t) \text{ and } A_{Q,n}(s,t) := \sum_{q\in Q} A_{q,n}(s,t).$$

Let $B_{q,n}(s,t)$ be the amount of service that flow q receives at node n during the time interval $(s,t]$. We denote by $V_{q,n}(t)$ the workload of flow q at node n at time t, and by $V_{q,n}$ a random variable with the limiting distribution of $V_{q,n}(t)$ for $t \to \infty$ (assuming it exists). We define

$$V_{Q,n}(t) := \sum_{q \in Q} V_{q,n}(t) \text{ and } V_n(t) := \sum_{q \in S^{(n)}} V_{q,n}(t),$$

where $V_n(t)$ is the total workload at node n at time t.

Using the above definitions, the identity given in (1.2) yields,

$$V_{q,n}(t) = A_{q,n}(s,t) + V_{q,n}(s) - B_{q,n}(s,t), \text{ for } 0 \leq s \leq t. \tag{3.5}$$

It implies the following relation between the arrival processes at two successive nodes,

$$A_{q,n+1}(s,t) = B_{q,n}(s,t) = A_{q,n}(s,t) + V_{q,n}(s) - V_{q,n}(t). \tag{3.6}$$

Following (1.4), the total workload at node n at time t is given by,

$$V_n(t) = \sup_{0 \leq s \leq t} \left\{ A_{S^{(n)},n}(s,t) - c_n(t-s) \right\}, \tag{3.7}$$

assuming $V_n(0) = 0$.

We define ρ_q to be the average rate of flow q and

$$\rho_Q := \sum_{q \in Q} \rho_q$$

to be the aggregate average rate of all flows $q \in Q$. Let $W_Q^c(t)$ be the workload at time t in a queue with service rate $c \geq 0$ which is fed by the flows $q \in Q$. Then, for $c > \rho_Q$, W_Q^c is a random variable with the limiting distribution of $W_Q^c(t)$ for $t \to \infty$. Analogously, we denote by $W_Q^{c_1,c_2}(t)$ the workload at time t at node 2 of a tandem network with service rates c_1 and c_2, fed by the flows $q \in Q$. For $c_2 > \rho_Q$, $W_Q^{c_1,c_2}$ is a random variable with the limiting distribution of $W_Q^{c_1,c_2}(t)$ for $t \to \infty$.

We make the following crucial assumptions throughout the remainder of this chapter.

Assumption 3.2.1 *We assume for each flow q, $\phi_{q,n} > \rho_q$ for all $n = 1, \ldots, N$.*

The above assumption ensures that each of the flows is guaranteed a higher minimum service rate than the average input rate. This is clearly sufficient for stability, regardless of the traffic parameters of the other flows, but certainly not necessary. Define $\tilde{c}_n := c_n - \rho_{S^{(n)} \setminus \{i\}}$ as the average service rate available at node n for flow i, i.e., the service rate at node n minus the aggregate average rate of all flows in $S^{(n)}$ other than i.

Assumption 3.2.2 *We assume $\tilde{c}_N < \tilde{c}_n$ for all $n = 1, \ldots, N-1$.*

The above assumption implies that node N can be viewed as the bottleneck node for flow i. In the following lemma we express the workload of the set of flows Q at node n in terms of the amount of traffic served of the other flows. The proof can be found in the appendix.

Lemma 3.2.1 (workload at time t) *Assuming $V_{Q,n}(0) = 0$,*

$$\begin{aligned} V_{Q,n}(t) &= \sup_{0 \le s \le t} \{A_{Q,n}(s,t) - B_{Q,n}(s,t)\} \\ &= \sup_{0 \le s \le t} \left\{A_{Q,n}(s,t) - (c_n(t-s) - B_{S^{(n)} \setminus Q,n}(s,t))\right\}. \end{aligned}$$

In the next lemma we present an upper bound for $V_{q,n}(t)$ which follows immediately from the GPS discipline. The result is trivial for the workload at node 1, see (1.5). It will be used in deriving the upper and lower bound for the workload of flow i at node 2 in Sections 3.3 and 3.5. Since this lemma is a special case of Lemma 3.6.3, we omit the proof.

Lemma 3.2.2 *For $n \in \{1, 2\}$,*

$$V_{q,n}(t) \le W_q^{\tilde{\phi}_q}(t),$$

with $\tilde{\phi}_q = \phi_{q,1}$ if $n = 1$ and $\tilde{\phi}_q = \min\{\phi_{q,1}, \phi_{q,2}\}$ if $n = 2$.

3.3 Merging flows

We distinguish between the following two scenarios. In this section we assume the other flows which feed into the network to join the path of flow i, i.e., they are not allowed to leave this path (see Fig. 3.1). In Section 3.5 flows *are* allowed to leave the path of flow i. The latter model includes cross traffic as a special case. We treat these scenarios separately because the model described in Section 3.5 requires a more difficult approach. At first sight, only the flows requesting service at node 2 seem to influence the workload of flow i at this node. However, the input rate of flow i plays an important role as well and depends on the behavior of all flows passing through node 1. Hence, it is not enough to consider only the flows which directly affect flow i at node 2. We also need to account for the behavior of all flows that once interacted with flow i.

In particular, we consider the following scenario in this section. We assume the GPS network to be fed by flow i and by two additional sets of flows. The set S_1 and flow i feed into node 1 and are served both at nodes 1 and 2, while the set of flows S_2 feeds into node 2 and receives only service at this node. We are interested in the distribution of the workload of flow i at node 2, $V_{i,2}$.

In this section we derive both a lower and an upper bound for $\mathbb{P}(V_{i,2} > x)$. The idea can be described as follows. If the flows other than i always showed exactly average behavior, then $V_{i,2}$ would be equal in distribution to $W_i^{\tilde{c}_1,\tilde{c}_2}$. In

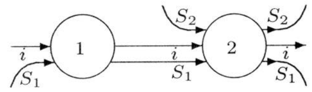

Figure 3.1: Two-node network with merging.

reality, stochastic fluctuations in the activity of the other flows will cause $V_{i,2}$ to deviate somewhat from $W_i^{\tilde{c}_1,\tilde{c}_2}$. Accordingly, the bounds will relate $V_{i,2}$ to $W_i^{\tilde{c}_1,\tilde{c}_2}$ with some additional correction terms. In the subsequent section, we will then show that these terms can be neglected asymptotically, resulting in the exact workload asymptotics.

In both the upper and lower bound for $V_{i,2}(t)$ we need a manageable expression for the total workload at node 2. The following lemma provides such an expression.

Lemma 3.3.1 *An alternative expression for $V_2(t)$ is:*

$$\sup_{0\le s\le t}\{A_i(s,t)+A_{S_1}(s,t)+A_{S_2}(s,t)+V_1(s)-c_2(t-s)\}$$
$$-\sup_{0\le s\le t}\{A_i(s,t)+A_{S_1}(s,t)-c_1(t-s)\}.$$

Proof Using (3.7) the total workload at node 2 is given by

$$V_2(t)=\sup_{0\le s\le t}\{A_{i,2}(s,t)+A_{S_1,2}(s,t)+A_{S_2}(s,t)-c_2(t-s)\}.$$

Using (3.6) to substitute for $A_{i,2}(s,t)+A_{S_1,2}(s,t)$ and then using (3.7) to substitute for $V_1(t)$ completes the proof. □

Before presenting the lower and upper bound, we introduce an additional variable. For $c<\rho_Q$, U_Q^c is defined to be a random variable with the limiting distribution of $U_Q^c(t)$ for $t\to\infty$, with

$$U_Q^c(t)=\sup_{0\le s\le t}\{c(t-s)-A_Q(s,t)\}. \tag{3.8}$$

In words, $U_Q^c(t)$ is the workload at a node of a flow which feeds this node at constant rate c and receives an amount of service $A_Q(s,t)$ during a time interval $(s,t]$.

Throughout the analysis, we use the following property of the sup operator,

$$\sup_t\{f(t)+g(t)\}\le\sup_t\{f(t)\}+\sup_t\{g(t)\}, \tag{3.9}$$

which also implies

$$\sup_t\{f(t)+g(t)\}\ge\sup_t\{f(t)\}-\sup_t\{-g(t)\}. \tag{3.10}$$

The lower bound for $\mathbb{P}(V_{i,2}>x)$ is given in the following lemma.

Lemma 3.3.2 *For any $\delta,\epsilon>0$ sufficiently small and any y,*

$$\mathbb{P}(V_{i,2}>x)\ge\mathbb{P}\left(W_i^{\tilde{c}_1-\epsilon,\tilde{c}_2+2\delta}>x+y\right)\mathbb{P}\left(Y^{\delta,\epsilon}\le y\right),$$

with $Y^{\delta,\epsilon}$ a stochastic variable with the limiting distribution of $Y^{\delta,\epsilon}(t)$ for $t\to\infty$, where

$$Y^{\delta,\epsilon}(t):=U_{S_1}^{\rho_{S_1}-\delta}(t)+U_{S_2}^{\rho_{S_2}-\delta}(t)+W_{S_1}^{\rho_{S_1}+\epsilon}(t)+\sum_{q\in S_1}W_q^{\tilde{\phi}_q}(t)+\sum_{q\in S_2}W_q^{\tilde{\phi}_q}(t). \tag{3.11}$$

The random variable $Y^{\delta,\epsilon}$ can be seen as the 'correction term' mentioned earlier, accounting for scenarios where $V_{i,2}(t)$ is smaller than $W_i^{\tilde{c}_1-\epsilon,\tilde{c}_2+2\delta}(t)$.

Proof By definition,

$$V_{i,2}(t) = V_2(t) - V_{S_1,2}(t) - V_{S_2,2}(t).$$

According to Lemma 3.3.1,

$$\begin{aligned} V_2(t) &= \sup_{0\leq s\leq t}\{A_i(s,t) + A_{S_1}(s,t) + A_{S_2}(s,t) + V_1(s) - c_2(t-s)\} \\ &\quad - \sup_{0\leq s\leq t}\{A_i(s,t) + A_{S_1}(s,t) - c_1(t-s)\}. \end{aligned} \tag{3.12}$$

Using (3.10), the first supremum in (3.12) can be lower bounded by

$$\sup_{0\leq s\leq t}\{A_i(s,t) - (\tilde{c}_2 + 2\delta)(t-s)\} - \sup_{0\leq s\leq t}\{(\rho_{S_1} - \delta)(t-s) - A_{S_1}(s,t)\}$$
$$- \sup_{0\leq s\leq t}\{(\rho_{S_2} - \delta)(t-s) - A_{S_2}(s,t)\}.$$

By definition, this is equal to

$$W_i^{\tilde{c}_2+2\delta}(t) - U_{S_1}^{\rho_{S_1}-\delta}(t) - U_{S_2}^{\rho_{S_2}-\delta}(t). \tag{3.13}$$

Using (3.9), the second supremum in (3.12) is upper bounded by

$$\begin{aligned} &\sup_{0\leq s\leq t}\{A_i(s,t) - (\tilde{c}_1 - \epsilon)(t-s)\} + \sup_{0\leq s\leq t}\{A_{S_1}(s,t) - (\rho_{S_1} + \epsilon)(t-s)\} \\ =\ & W_i^{\tilde{c}_1-\epsilon}(t) + W_{S_1}^{\rho_{S_1}+\epsilon}(t). \end{aligned} \tag{3.14}$$

Finally, we have to find an upper bound for $V_{S_1}(t) + V_{S_2}(t)$. Using Lemma 3.2.2,

$$V_{S_1,2}(t) + V_{S_2,2}(t) \leq \sum_{q\in S_1} W_q^{\tilde{\phi}_q}(t) + \sum_{q\in S_2} W_q^{\tilde{\phi}_q}(t). \tag{3.15}$$

Arranging the terms in (3.13), (3.14) and (3.15), we obtain, using both (3.1) and (3.11),

$$V_{i,2}(t) \geq W_i^{\tilde{c}_1-\epsilon,\tilde{c}_2+2\delta}(t) - Y^{\delta,\epsilon}(t).$$

Hence, a lower bound is given by

$$\mathbb{P}(V_{i,2} > x) \geq \mathbb{P}\left(\left\{W_i^{\tilde{c}_1-\epsilon,\tilde{c}_2+2\delta} > x + y\right\} \cap \left\{Y^{\delta,\epsilon} \leq y\right\}\right),$$

for any y. Because $Y^{\delta,\epsilon}$ is independent of the traffic process of flow i, the proof is completed. □

The next lemma provides an upper bound for $\mathbb{P}(V_{i,2} > x)$.

Lemma 3.3.3 *For any $\eta, \nu > 0$ sufficiently small and any y,*

$$\mathbb{P}(V_{i,2} > x) \leq \mathbb{P}\left(W_i^{\tilde{c}_1+\eta,\tilde{c}_2-2\nu} > x - y\right) + \mathbb{P}\left(W_i^{\tilde{\phi}_i} > x\right)\mathbb{P}\left(Z^{\eta,\nu} > y\right),$$

with $Z^{\eta,\nu}$ a random variable with the limiting distribution of $Z^{\eta,\nu}(t)$ for $t \to \infty$, where

$$Z^{\eta,\nu}(t) := U_{S_1}^{\rho_{S_1}-\eta}(t) + W_{S_1}^{\rho_{S_1}+\nu}(t) + W_{S_2}^{\rho_{S_2}+\nu}(t). \tag{3.16}$$

Analogously to $Y^{\delta,\epsilon}$ in the lower bound, the random variable $Z^{\eta,\nu}$ can be seen as the correction term, corresponding to situations where $V_{i,2}(t)$ is larger than $W_i^{\tilde{c}_1+\eta,\tilde{c}_2-2\nu}(t)$.

Proof By definition,

$$V_{i,2}(t) \leq V_2(t).$$

According to Lemma 3.3.1,

$$\begin{aligned} V_2(t) &= \sup_{0\leq s\leq t}\{A_i(s,t) + A_{S_1}(s,t) + A_{S_2}(s,t) + V_1(s) - c_2(t-s)\} \\ &\quad - \sup_{0\leq s\leq t}\{A_i(s,t) + A_{S_1}(s,t) - c_1(t-s)\}. \end{aligned} \tag{3.17}$$

Using (3.7) to substitute for $V_1(s)$, we obtain for the first supremum in (3.17),

$$\sup_{0\leq u\leq s\leq t}\{A_i(u,t) + A_{S_1}(u,t) + A_{S_2}(s,t) - c_1(s-u) - c_2(t-s)\},$$

which is upper bounded by, using (3.9),

$$\begin{aligned} &\sup_{0\leq s\leq t}\{A_{S_2}(s,t) - (\rho_{S_2}+\nu)(t-s)\} + \sup_{0\leq u\leq t}\{A_{S_1}(u,t) - (\rho_{S_1}+\nu)(t-u)\} \\ &+ \sup_{0\leq u\leq s\leq t}\{A_i(u,t) - (\tilde{c}_1-\nu)(s-u) - (\tilde{c}_2-2\nu)(t-s)\}. \end{aligned} \tag{3.18}$$

The first two suprema in (3.18) are equal to

$$W_{S_1}^{\rho_{S_1}+\nu}(t) + W_{S_2}^{\rho_{S_2}+\nu}(t). \tag{3.19}$$

Because $\tilde{c}_2 < \tilde{c}_1$ (Assumption 3.2.2), the third supremum in (3.18) is upper bounded by

$$\sup_{0\leq s\leq t}\{A_i(s,t) - (\tilde{c}_2-2\nu)(t-s)\} = W_i^{\tilde{c}_2-2\nu}(t). \tag{3.20}$$

Next we have to find a lower bound for the second supremum in (3.17). Using (3.10), we obtain as lower bound,

$$\begin{aligned} &\sup_{0\leq s\leq t}\{A_i(s,t) - (\tilde{c}_1+\eta)(t-s)\} - \sup_{0\leq s\leq t}\{(\rho_{S_1}-\eta)(t-s) - A_{S_1}(s,t)\} \\ &= W_i^{\tilde{c}_1+\eta}(t) - U_{S_1}^{\rho_{S_1}-\eta}(t). \end{aligned} \tag{3.21}$$

Arranging the terms in (3.19), (3.20) and (3.21), we obtain, using both (3.1) and (3.16),

$$V_{i,2}(t) \leq W_i^{\tilde{c}_1+\eta,\tilde{c}_2-2\nu}(t) + Z^{\eta,\nu}(t).$$

Combining the above bound with the upper bound in Lemma 3.2.2,

$$V_{i,2}(t) \leq \min\left\{W_i^{\tilde{\phi}_i}(t), W_i^{\tilde{c}_1+\eta,\tilde{c}_2-2\nu}(t) + Z^{\eta,\nu}(t)\right\}.$$

Hence, an upper bound is given by

$$\mathbb{P}(V_{i,2} > x) \leq \mathbb{P}\left(\left\{W_i^{\tilde{\phi}_i} > x\right\} \cap \left(\left\{W_i^{\tilde{c}_1+\eta,\tilde{c}_2-2\nu} > x-y\right\} \cup \{Z^{\eta,\nu} > y\}\right)\right),$$

for any y, which gives the result because $Z^{\eta,\nu}$ is independent of the traffic process of flow i. □

3.4 Tail behavior of the workload distribution

We now state our key theorem concerning the tail behavior of the workload distribution. It follows immediately from Lemma 3.4.1, which is given below.

Theorem 3.4.1 *In the two-node network with merging flows, under Assumptions 3.2.1 and 3.2.2,*

$$\mathbb{P}(V_{i,2} > x) \sim \mathbb{P}\left(W_i^{\tilde{c}_1,\tilde{c}_2} > x\right).$$

According to this theorem, the workload distribution of flow i at node 2 is asymptotically equivalent to that in a tandem network where flow i is served in isolation at rates $\tilde{c}_1$ and $\tilde{c}_2$. Hence, the workload of flow i at node 2 is only affected by the characteristics of the other flows through their average rates, even when the other flows are 'heavier-tailed'. This suggests that an extremely large workload of flow i is most likely due to either a long on period or a large burst size of flow i itself. During the subsequent congestion period, the other flows continue to receive service at approximately their average rates. In the

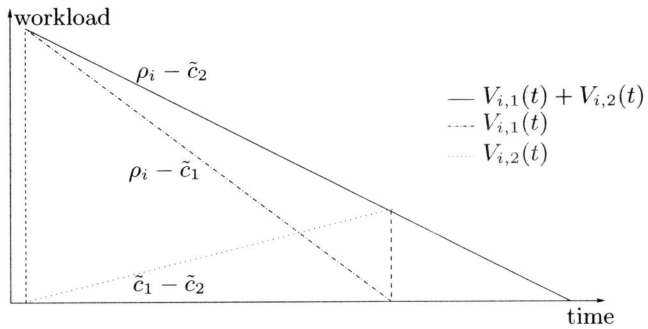

Figure 3.2: Overflow scenario for instantaneous traffic bursts.

theorem this is represented by the constant rates $\tilde{c}_1$ and $\tilde{c}_2$. This extends the result of Theorem 3.1 in [21] for the single-node case and shows that GPS is capable of isolating flows in networks as well.

The typical overflow scenario is schematically depicted in Fig. 3.2. At some point, flow i generates a large burst, causing $V_{i,1}(t)$ to reach some large value. After that, flow i returns to its average behavior, producing traffic at rate ρ_i. Consequently, $V_{i,1}(t)$ will start to decrease roughly at rate $\rho_i - \tilde{c}_1$, and $V_{i,2}(t)$ will start to increase approximately at rate $\tilde{c}_1 - \tilde{c}_2$, until $V_{i,1}(t)$ reduces to zero at some point. From then on, $V_{i,1}(t)$ will remain relatively small, and $V_{i,2}(t)$ will also start to decrease, roughly at rate $\rho_i - \tilde{c}_2$, until $V_{i,2}(t)$ becomes zero as well. The corresponding behavior for an on-off process is illustrated in Fig. 3.3. Note that we do not need any assumptions on the characteristics of the flows other than flow i. They can have lighter-tailed characteristics, or even more heavy-tailed characteristics. We should stress though that Assumption 3.2.1 is crucial. If $\rho_q > \phi_{q,n}$ for some n then flow i may not receive service at a stable rate when other flows generate a large amount of traffic. Flows with an on-period distribution or a burst-size distribution which is heavier-tailed than that of flow i will then potentially affect the workload of flow i, see [21].

The above theorem follows from Lemma 3.4.1, which is a general result showing that the bounds of Lemmas 3.3.2 and 3.3.3 asymptotically coincide. For this general lemma, we use some additional generic notation. Let a_1 and a_2 be some positive constants which play the role of $\tilde{c}_j$ and $\tilde{c}_N$ respectively, where $j < N$, and N is the bottleneck node on the path of flow i. Let a be a positive constant representing in our results the $\tilde{\phi}_i$. We introduce a random variable R_i that plays the role of the workload of flow i at the bottleneck node on its path. Finally, for δ, ϵ, η and $\nu > 0$ let $C_{-i}^{\delta,\epsilon}$ and $D_{-i}^{\eta,\nu}$ also be random variables.

Lemma 3.4.1 *If for δ, ϵ, η and $\nu > 0$ sufficiently small and any y,*

$$\mathbb{P}(R_i > x) \geq \mathbb{P}\left(W_i^{a_1-\epsilon, a_2+\delta} > x + y\right) \mathbb{P}(C_{-i}^{\delta,\epsilon} \leq y), \tag{3.22}$$

$$\mathbb{P}(R_i > x) \leq \mathbb{P}\left(W_i^{a_1+\eta, a_2-\nu} > x - y\right) + \mathbb{P}(W_i^a > x)\mathbb{P}\left(D_{-i}^{\eta,\nu} > y\right), \tag{3.23}$$

and both $\mathbb{P}(W_i^a > x)$ and $\mathbb{P}(W_i^{a_1,a_2} > x)$ satisfy Properties (3.2), (3.3) and (3.4), then

$$\mathbb{P}(R_i > x) \sim \mathbb{P}\left(W_i^{a_1,a_2} > x\right). \tag{3.24}$$

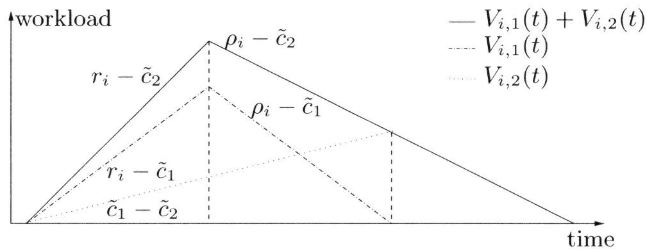

Figure 3.3: Overflow scenario for an on-off process.

Proof The lower bound (3.22) implies, for any $\delta, \epsilon > 0$ sufficiently small and any y,

$$\frac{\mathbb{P}(R_i > x)}{\mathbb{P}(W_i^{a_1,a_2} > x)} \geq \frac{\mathbb{P}(W_i^{a_1-\epsilon,a_2+\delta} > x+y)}{\mathbb{P}(W_i^{a_1,a_2} > x+y)} \frac{\mathbb{P}(W_i^{a_1,a_2} > x+y)}{\mathbb{P}(W_i^{a_1,a_2} > x)} \mathbb{P}(C_{-i}^{\delta,\epsilon} \leq y).$$

Using Properties (3.2) and (3.3), we obtain

$$\liminf_{x\to\infty} \frac{\mathbb{P}(R_i > x)}{\mathbb{P}(W_i^{a_1,a_2} > x)} \geq G_i(-\epsilon, \delta)\mathbb{P}(C_{-i}^{\delta,\epsilon} \leq y).$$

Letting $y \to \infty$ and then $\delta, \epsilon \downarrow 0$,

$$\liminf_{x\to\infty} \frac{\mathbb{P}(R_i > x)}{\mathbb{P}(W_i^{a_1,a_2} > x)} \geq 1. \tag{3.25}$$

Analogously, the upper bound (3.23) implies, for any η, $\nu > 0$ sufficiently small and any y,

$$\begin{aligned}\frac{\mathbb{P}(R_i > x)}{\mathbb{P}(W_i^{a_1,a_2} > x)} \leq & \frac{\mathbb{P}(W_i^{a_1+\eta,a_2-\nu} > x-y)}{\mathbb{P}(W_i^{a_1,a_2} > x-y)} \frac{\mathbb{P}(W_i^{a_1,a_2} > x-y)}{\mathbb{P}(W_i^{a_1,a_2} > x)} \\ & + \frac{\mathbb{P}(W_i^{a} > x)}{\mathbb{P}(W_i^{a_1,a_2} > x)} \mathbb{P}(D_{-i}^{\eta,\nu} > y).\end{aligned}$$

Using Properties (3.2), (3.3) and (3.4), we have

$$\limsup_{x\to\infty} \frac{\mathbb{P}(R_i > x)}{\mathbb{P}(W_i^{a_1,a_2} > x)} \leq G_i(\eta, -\nu) + C\mathbb{P}(D_{-i}^{\eta,\nu} > y),$$

for some constant $C < \infty$. Letting $y \to \infty$ and $\eta, \nu \downarrow 0$,

$$\limsup_{x\to\infty} \frac{\mathbb{P}(R_i > x)}{\mathbb{P}(W_i^{a_1,a_2} > x)} \leq 1. \tag{3.26}$$

Combining (3.25) and (3.26) gives the desired result. □

Proof of Theorem 3.4.1 Combine Lemmas 3.3.2, 3.3.3 and 3.4.1. □

3.5 Splitting flows

Consider again a tandem network in which the following flows are served according to the GPS principle (see Fig. 3.4). As in Section 3.3, flow i and the set of flows S_1 feed into node 1 and are served both at nodes 1 and 2, and the set of flows S_2 feeds into node 2. In addition, we consider in this section the set of flows S_3 which feeds into node 1 but does not move on to node 2 after receiving service at node 1. We first derive a lower bound and an upper bound for the workload distribution of flow i at node 2, $\mathbb{P}(V_{i,2} > x)$. Then we use Lemma 3.4.1 to determine the tail behavior of $\mathbb{P}(V_{i,2} > x)$.

In the following lemma we give an alternative expression for $V_2(t)$ which we need in the proof of the lower and upper bound for $\mathbb{P}(V_{i,2} > x)$.

Lemma 3.5.1 *An alternative expression for $V_2(t)$ is:*

$$\sup_{0\le s\le t}\{A_i(s,t)+A_{S_1}(s,t)+A_{S_2}(s,t)+V_{i,1}(s)+V_{S_1,1}(s)-c_2(t-s)\}$$
$$-\sup_{0\le s\le t}\{A_i(s,t)+A_{S_1}(s,t)+A_{S_3}(s,t)-c_1(t-s)\}+V_{S_3,1}(t).$$

Proof Because of (3.7),

$$V_2(t)=\sup_{0\le s\le t}\{A_{i,2}(s,t)+A_{S_1,2}(s,t)+A_{S_2}(s,t)-c_2(t-s)\}.$$

Using (3.6) to substitute for $A_{i,2}(s,t)+A_{S_1,2}(s,t)$, we obtain for $V_2(t)$

$$\sup_{0\le s\le t}\{A_i(s,t)+A_{S_1}(s,t)+A_{S_2}(s,t)+V_{i,1}(s)+V_{S_1,1}(s)-c_2(t-s)\}$$
$$-V_{i,1}(t)-V_{S_1,1}(t).$$

As $V_1(t)=V_{i,1}(t)+V_{S_1,1}(t)+V_{S_3,1}(t)$, the proof is completed using (3.7) to substitute for $V_1(t)$. □

Analogously to Section 3.3, we introduce some additional variables. Due to the presence of the additional set of flows S_3, these variables are more complicated. For $\delta,\epsilon>0$, redefine $Y^{\delta,\epsilon}$ to be a random variable with the limiting distribution of $Y^{\delta,\epsilon}(t)$ for $t\to\infty$, with

$$\begin{aligned}Y^{\delta,\epsilon}(t) \quad := \quad & W_{S_1}^{\rho_{S_1}+\epsilon}(t)+W_{S_3}^{\rho_{S_3}+\epsilon}(t)+\sum_{q\in S_1}W_q^{\tilde{\phi}_q}(t)+\sum_{q\in S_2}W_q^{\tilde{\phi}_q}(t)\\ &+U_{S_1}^{\rho_{S_1}-\delta}(t)+U_{S_2}^{\rho_{S_2}-\delta}(t).\end{aligned} \tag{3.27}$$

For $\eta,\nu>0$, redefine $Z^{\eta,\nu}$ to be a random variable with the limiting distribution of $Z^{\eta,\nu}(t)$ for $t\to\infty$, with

$$\begin{aligned}Z^{\eta,\nu}(t) \quad := \quad & U_{S_1}^{\rho_{S_1}-\eta}(t)+U_{S_3}^{\rho_{S_3}-\nu}(t)+U_{S_3}^{\rho_{S_3}-\eta}(t)\\ &+\sum_{j=1}^{3}W_{S_j}^{\rho_{S_j}+\nu}(t)+\sum_{q\in S_3}W_q^{\tilde{\phi}_q}(t).\end{aligned} \tag{3.28}$$

Now we derive both an upper and a lower bound for $\mathbb{P}(V_{i,2}>x)$. These bounds are similar to the bounds in Lemmas 3.3.2 and 3.3.3, except for the structure of the correction terms $Y^{\delta,\epsilon}$ and $Z^{\eta,\nu}$. In the following lemma we give a lower bound for $\mathbb{P}(V_{i,2}>x)$.

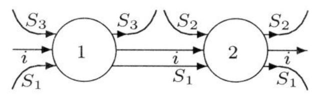

Figure 3.4: Two-node network with splitting.

Lemma 3.5.2 *For any $\delta, \epsilon > 0$ sufficiently small and any y,*

$$\mathbb{P}(V_{i,2} > x) \geq \mathbb{P}\left(W_i^{\tilde{c}_1 - 2\epsilon, \tilde{c}_2 + 2\delta} > x + y\right) \mathbb{P}(Y^{\delta,\epsilon} \leq y).$$

Proof By definition,

$$V_{i,2}(t) = V_2(t) - V_{S_1,2}(t) - V_{S_2,2}(t). \tag{3.29}$$

According to Lemma 3.5.1,

$$\begin{aligned} V_2(t) \geq & \sup_{0 \leq s \leq t} \{A_i(s,t) + A_{S_1}(s,t) + A_{S_2}(s,t) - c_2(t-s)\} \\ & - \sup_{0 \leq s \leq t} \{A_i(s,t) + A_{S_1}(s,t) + A_{S_3}(s,t) - c_1(t-s)\}. \end{aligned} \tag{3.30}$$

Using (3.10), the first supremum in (3.30) is lower bounded by

$$\begin{aligned} & \sup_{0 \leq s \leq t} \{A_i(s,t) - (\tilde{c}_2 + 2\delta)(t-s)\} - \sup_{0 \leq s \leq t} \{(\rho_{S_1} - \delta)(t-s) - A_{S_1}(s,t)\} \\ & - \sup_{0 \leq s \leq t} \{(\rho_{S_2} - \delta)(t-s) - A_{S_2}(s,t)\}, \end{aligned}$$

which is equal to (using (3.8))

$$W_i^{\tilde{c}_2 + 2\delta}(t) - U_{S_1}^{\rho_{S_1} - \delta}(t) - U_{S_2}^{\rho_{S_2} - \delta}(t).$$

Next we need an upper bound for the second supremum in (3.30). Using (3.9) it is upper bounded by

$$W_i^{\tilde{c}_1 - 2\epsilon}(t) + W_{S_1}^{\rho_{S_1} + \epsilon}(t) + W_{S_3}^{\rho_{S_3} + \epsilon}(t).$$

Finally, using Lemma 3.2.2, we find a similar upper bound for $V_{S_1,2}(t)$ and $V_{S_2,2}(t)$ as in (3.15). Adding the three bounds and using (3.1) and (3.27),

$$V_{i,2}(t) \geq W_i^{\tilde{c}_1 - 2\epsilon, \tilde{c}_2 + 2\delta}(t) - Y^{\delta,\epsilon}(t).$$

Because $Y^{\delta,\epsilon}$ is independent of the traffic process of flow i, the proof is completed. □

The following lemma provides an upper bound for $\mathbb{P}(V_{i,2} > x)$.

Lemma 3.5.3 *For any $\eta, \nu > 0$ sufficiently small and any y,*

$$\mathbb{P}(V_{i,2} > x) \leq \mathbb{P}\left(W_i^{\tilde{c}_1 + 2\eta, \tilde{c}_2 - 4\nu} > x - y\right) + \mathbb{P}\left(W_i^{\tilde{\phi}_i} > x\right) \mathbb{P}(Z^{\eta,\nu} > y).$$

Proof By definition,

$$V_{i,2}(t) \leq V_2(t).$$

According to Lemma 3.5.1, $V_2(t)$ is equal to

$$\sup_{0 \leq s \leq t} \{A_i(s,t) + A_{S_1}(s,t) + A_{S_2}(s,t) + V_{i,1}(s) + V_{S_1,1}(s) - c_2(t-s)\}$$

$$- \sup_{0\le s\le t} \{A_i(s,t) + A_{S_1}(s,t) + A_{S_3}(s,t) - c_1(t-s)\} + V_{S_3,1}(t). \qquad (3.31)$$

First observe that $V_{i,1}(s)+V_{S_1,1}(s) \le V_1(s)$. Using (3.7) to substitute for $V_1(s)$, the first supremum in (3.31) is thus upper bounded by

$$\sup_{0\le r\le s\le t} \{A_i(r,t)+A_{S_1}(r,t)+A_{S_3}(r,s)+A_{S_2}(s,t)-c_1(s-r)-c_2(t-s)\}. \quad (3.32)$$

Note that (3.32) can be written as

$$\begin{aligned}\sup_{0\le r\le s\le t} \{&A_i(r,t) - (\tilde{c}_1 - 2\nu)(s-r) - (\tilde{c}_2 - 4\nu)(t-s) \\ &+A_{S_1}(r,t) - (\rho_{S_1} + \nu)(t-r) + A_{S_2}(s,t) - (\rho_{S_2} + \nu)(t-s) \\ &+A_{S_3}(r,t) - (\rho_{S_3} + \nu)(t-r) + (\rho_{S_3} - \nu)(t-s) - A_{S_3}(s,t)\}.\end{aligned}$$

Using (3.9) and $\tilde{c}_1 > \tilde{c}_2$ (Assumption 3.2.2), this is upper bounded by

$$\begin{aligned}&\sup_{0\le r\le t} \{A_i(r,t) - (\tilde{c}_2 - 4\nu)(t-r)\} + \sup_{0\le r\le s\le t} \{-2\nu(s-r)\} \\ &+ \sup_{0\le r\le t} \{A_{S_1}(r,t) - (\rho_{S_1} + \nu)(t-r)\} \\ &+ \sup_{0\le s\le t} \{A_{S_2}(s,t) - (\rho_{S_2} + \nu)(t-s)\} \\ &+ \sup_{0\le r\le t} \{A_{S_3}(r,t) - (\rho_{S_3} + \nu)(t-r)\} \\ &+ \sup_{0\le s\le t} \{(\rho_{S_3} - \nu)(t-s) - A_{S_3}(s,t)\},\end{aligned}$$

which by definition is equal to

$$W_i^{\tilde{c}_2-4\nu}(t) + W_{S_1}^{\rho_{S_1}+\nu}(t) + W_{S_2}^{\rho_{S_2}+\nu}(t) + W_{S_3}^{\rho_{S_3}+\nu}(t) + U_{S_3}^{\rho_{S_3}-\nu}(t).$$

Now we have to find a lower bound for the second supremum in (3.31). Using (3.10), this lower bound is given by

$$W_i^{\tilde{c}_1+2\eta}(t) - U_{S_1}^{\rho_{S_1}-\eta}(t) - U_{S_3}^{\rho_{S_3}-\eta}(t).$$

Finally, because of Lemma 3.2.2, we obtain for the third term in (3.31)

$$V_{S_3,1}(t) \le \sum_{q\in S_3} W_q^{\tilde{\phi}_q}(t).$$

Adding the three bounds and using (3.1) and (3.28),

$$V_{i,2}(t) \le W_i^{\tilde{c}_1+2\eta,\tilde{c}_2-4\nu} + Z^{\eta,\nu}(t).$$

Combining the above bound with the upper bound in Lemma 3.2.2, we obtain the following upper bound,

$$V_{i,2}(t) \le \min\left\{W_i^{\tilde{\phi}_i}(t), W_i^{\tilde{c}_1+2\eta,\tilde{c}_2-4\nu}(t) + Z^{\eta,\nu}(t)\right\}.$$

Because $Z^{\eta,\nu}$ is independent of the traffic process of flow i, the proof is completed. □

Now we have all the ingredients to use Lemma 3.4.1, which gives the main result of this section.

Theorem 3.5.1 *Under Assumptions 3.2.1 and 3.2.2,*

$$\mathbb{P}(V_{i,2} > x) \sim \mathbb{P}\left(W_i^{\tilde{c}_1,\tilde{c}_2} > x\right).$$

3.6 Preliminaries general networks

In the next two sections we extend the model of Section 3.5 and focus on the Nth node on the path of flow i. We assume this node to be the bottleneck node for flow i. Again we assume the flows to be served at each node according to the GPS mechanism. First we introduce some additional notation and present a number of lemmas which we use in the next sections. Then we analyze the behavior of the workload of flow i at the bottleneck node on its path, if no other flows feed into any of the nodes on this path. Although this model is quite simple, it provides some useful intuition for the results in Sections 3.7 and 3.8.

We define S_j to be the set of flows that feed into node j and S_m^p to be the set of flows that feed into node m and leave the path of flow i at node p (so flows in S_m^p receive service at node p). For $q \in S_m^p$ we define $\tilde{\phi}_q := \min\{\phi_{q,m}, \ldots, \phi_{q,p}\}$, which is the minimum rate guaranteed to flow q on its path along node m up to and until p.

We now present some lemmas which we use in the next sections. The proofs can be found in the appendix of this chapter. The first lemma gives a lower bound for the amount of service that flow q receives at node n during the time interval $(s, t]$. The time s_m in the lower bound expression represents the last epoch before time t when flow q is non-backlogged, i.e., at which it starts a busy period that lasts at least until time t. Consequently, the amount of service received in $(s, s_m]$ is at least equal to $A_q(s, s_m)$ and the amount of service received in $(s_m, t]$ is at least $\tilde{\phi}_q(t - s_m) \geq \gamma_q(t - s_m)$.

Lemma 3.6.1 *For $q \in S_m^p$, $1 \leq m \leq n \leq p$ and $\gamma_q \leq \tilde{\phi}_q$,*

$$\begin{aligned} B_{q,n}(s,t) &\geq \gamma_q(t-s) - \sup_{s \leq s_m \leq t} \{\gamma_q(s_m - s) - A_q(s, s_m)\} \\ &= \inf_{s \leq s_m \leq t} \{A_q(s, s_m) + \gamma_q(t - s_m)\}. \end{aligned}$$

Using this lemma, we can derive an upper bound for the total workload of flow $q \in S_m^p$ at nodes $m, \ldots, n$, $n \leq p$. This upper bound is presented in the next lemma.

Lemma 3.6.2 *For $q \in S_m^p$, $1 \leq m \leq n \leq p$ and $\gamma_q \leq \tilde{\phi}_q$,*

$$\sum_{j=m}^{n} V_{q,j}(t) \leq W_q^{\gamma_q}(t).$$

In fact, the above lemma indicates that the total amount of work in the GPS network, where traffic from flow q has to traverse nodes m through n, is upper bounded by the amount of work in a single-node system with the same input traffic. Crucial for this lemma to hold is that the traffic in the single-node system is served at a rate which is at most the minimum guaranteed rate for flow q along nodes m through n. It immediately implies the following lemma, which includes Lemma 3.2.2 as a special case.

Lemma 3.6.3 *For* $q \in S_m^p$, $1 \leq m \leq n \leq p$,

$$V_{q,n}(t) \leq W_q^{\tilde{\phi}_q}(t).$$

From Lemma 3.6.2 we can derive an upper bound for the amount of service that flow q receives during the interval $(s, t]$ as well. This upper bound is given in the following lemma.

Lemma 3.6.4 *For* $q \in S_m^p$, $1 \leq m \leq n \leq p$ *and* $\gamma_q \leq \tilde{\phi}_q$,

$$B_{q,n}(s,t) \leq \gamma_q(t-s) + \sup_{0 \leq s_m \leq t} \{A_q(s_m, t) - \gamma_q(t - s_m)\}.$$

We now briefly discuss the workload behavior at the Nth node of a network which is fed *only* by flow i. Take $m^\star \in \arg\min_{n=1,\ldots,N-1}\{\tilde{c}_n\}$. In Section 3.2 we assumed that $\tilde{c}_n > \tilde{c}_N$ (Assumption 3.2.2) for all $n = 1, \ldots, N-1$, so that $\tilde{c}_{m^\star} > \tilde{c}_N$. The workload distribution, $\mathbb{P}(V_{i,N} > x)$, is given in the following theorem.

Theorem 3.6.1

$$\mathbb{P}(V_{i,N} > x) = \mathbb{P}\left(W_i^{c_{m^\star}, c_N} > x\right).$$

Proof Observe that, because of the definition of $m^\star$, the total workload at nodes $1, \ldots, m^\star$ is equal to that at a node with service rate $c_{m^\star}$ which is fed by the original traffic process of flow i. Hence,

$$\sum_{j=1}^{m^\star} V_{i,j}(t) = W_i^{c_{m^\star}}(t), \tag{3.33}$$

for which a formal proof is given in the appendix. Since $c_N < c_{m^\star}$ (Assumption 3.2.2), we can apply the same reasoning to the total workload at nodes $1, \ldots, N$, and we have

$$\sum_{j=1}^{N} V_{i,j}(t) = W_i^{c_N}(t).$$

In [71] the following observation is made. If $c_k > c_j$ for $k > j$ then the backlog at node k will always be zero in stationarity and this node can be removed from the network. Because the nodes succeeding node $m^\star$ (except N) have a service rate which is larger than $c_{m^\star}$, $\sum_{j=m^\star}^{N-1} V_{ij}(t) = 0$, and we have, using (3.1),

$$V_{i,N}(t) = \sum_{j=1}^{N} V_{i,j}(t) - \sum_{j=1}^{m^\star} V_{i,j}(t) = W_i^{c_{m^\star}, c_N}(t),$$

which completes the proof. □

Hence, the workload at node N in this network is equal to that at node 2 in a *two*-node tandem network serving flow i at rates $c_{m^\star}$ and c_N. Thus the distribution of the workload is entirely determined by the bottleneck nodes. Intuitively, this can be explained as follows. The workload at a particular node depends on two rates, the rate at which traffic is sent into the node and the rate at which traffic is served by the node. The first rate is the rate at which traffic is served by the bottleneck node on the path to the relevant node, i.e., $c_{m^\star}$. The other rate is the service rate for flow i, which is c_N. Asymptotically, this is still true for the more general networks which we discuss in the next sections. For more results on the reduction of networks we refer to [120, 125].

3.7 General network with merging

Analogously to Sections 3.3 and 3.5, we distinguish between two network scenarios. In this section we consider an extension of the network described in Section 3.3 and assume that each node on the path of flow i in the GPS network is fed by an additional set of flows (see Fig. 3.5 for an example where flow i traverses $N = 4$ nodes). These sets follow the path of flow i and do not leave before node N, the bottleneck node. In Section 3.8 we consider an extension of this network and the network described in Section 3.5, and allow the flows feeding into a node on the path of flow i to leave this path before the bottleneck node. We first derive bounds for $\mathbb{P}(V_{i,N} > x)$. The idea is similar to that in Section 3.3. If the flows other than i always showed exactly average behavior, then $V_{i,n}$ would in distribution be equal to $W_{i,N}^{\tilde{c}_1,\ldots,\tilde{c}_N}$. In Section 3.6 we showed that $W_{i,N}^{\tilde{c}_1,\ldots,\tilde{c}_N}$ has the same distribution as $W_i^{\tilde{c}_{m^\star},\tilde{c}_N}$. In addition to $W_i^{\tilde{c}_{m^\star},\tilde{c}_N}$, the bounds contain some correction terms accounting for the stochastic fluctuations of the flows other than flow i, which we later show can be asymptotically neglected. Note that the proofs of some lemmas are moved again to the appendix.

Recall that in the two-node model the upper and lower bounds for $V_{i,2}(t)$ were derived from bounds for $V_1(t)$ and $V_2(t)$. Similarly, in the N-node case,

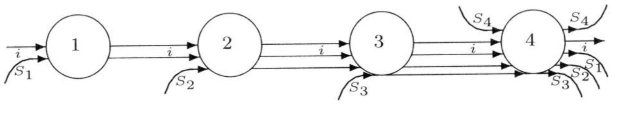

Figure 3.5: Example of general network with merging where $N = 4$.

the lower and upper bounds for $V_{i,N}(t)$ rely on bounds for the total workload at each node $n \in \{1,\ldots,N\}$. Define

$$X_n(t) := \sup_{0 \le s_1 \le \ldots \le s_n \le s_{n+1} = t} \left\{ A_i(s_1,t) + \sum_{j=1}^{n} \left[A_{S_j}(s_j,t) - c_j(s_{j+1}-s_j) \right] \right\}.$$

In the next lemma we give an expression for $V_n(t)$ in terms of $X_n(t)$. This expression will be used in deriving the upper and lower bounds for $V_{i,N}(t)$.

Lemma 3.7.1 *For $n \ge 2$,*

$$V_n(t) = X_n(t) - X_{n-1}(t).$$

In order to determine a lower and an upper bound for $V_n(t)$ we have to find a lower and an upper bound for $X_n(t)$. In the next lemma the lower bound for $X_n(t)$ is presented.

Lemma 3.7.2 *For any $\theta_1,\ldots,\theta_n$,*

$$X_n(t) \ge W_i^e(t) - \sum_{j=1}^{n} U_{S_j}^{\theta_j}(t),$$

with $e := \min_{m=1,\ldots,n}\{c_m - \sum_{j=1}^{m} \theta_j\}$.

The upper bound for $X_n(t)$ is given in the following lemma.

Lemma 3.7.3 *For any $\xi_1,\ldots,\xi_n$,*

$$X_n(t) \le W_i^d(t) + \sum_{j=1}^{n} W_{S_j}^{\xi_j}(t),$$

with $d := \min_{m=1,\ldots,n}\{c_m - \sum_{j=1}^{m} \xi_j\}$.

We now introduce some additional notation similar to Section 3.3. For $\delta,\epsilon > 0$, redefine $Y^{\delta,\epsilon}$ as a random variable with the limiting distribution of $Y^{\delta,\epsilon}(t)$ for $t \to \infty$, where

$$Y^{\delta,\epsilon}(t) := \sum_{j=1}^{N} U_{S_j}^{\rho_{S_j}-\delta}(t) + \sum_{j=1}^{N-1} W_{S_j}^{\rho_{S_j}+\epsilon}(t) + \sum_{j=1}^{N} \sum_{q \in S_j} W_q^{\tilde{\phi}_q}(t). \tag{3.34}$$

For $\eta,\nu > 0$, redefine $Z^{\eta,\nu}$ as a random variable with the limiting distribution of $Z^{\eta,\nu}(t)$ for $t \to \infty$, where

$$Z^{\eta,\nu}(t) := \sum_{j=1}^{N} W_{S_j}^{\rho_{S_j}+\nu}(t) + \sum_{j=1}^{N-1} U_{S_j}^{\rho_{S_j}-\eta}(t). \tag{3.35}$$

We use the bounds for $X_n(t)$ to construct bounds for $\mathbb{P}(V_{i,N} > x)$. The lower bound is given in the following lemma.

Lemma 3.7.4 *For any $\delta, \epsilon > 0$ sufficiently small and any y,*

$$\mathbb{P}(V_{i,N} > x) \geq \mathbb{P}\left(W_i^{\tilde{c}_{m^\star} - m^\star \epsilon, \tilde{c}_N + N\delta} > x + y\right) \mathbb{P}\left(Y^{\delta,\epsilon} \leq y\right).$$

Proof By definition,

$$V_{i,N}(t) = V_N(t) - \sum_{j=1}^{N} \sum_{q \in S_j} V_{q,N}(t).$$

Using Lemmas 3.6.3 and 3.7.1, $V_{i,N}(t)$ is lower bounded by

$$X_N(t) - X_{N-1}(t) - \sum_{j=1}^{N} \sum_{q \in S_j} W_q^{\tilde{\phi}_q}(t).$$

Now we can use the lower bound in Lemma 3.7.2 for $X_N(t)$ and the upper bound in Lemma 3.7.3 for $X_{N-1}(t)$. Taking $\theta_j = \rho_{S_j} - \delta$ in Lemma 3.7.2 and $\xi_j = \rho_{S_j} + \epsilon$ in Lemma 3.7.3, we obtain for $\delta, \epsilon > 0$ sufficiently small,

$$\begin{aligned} V_{i,N}(t) \quad &\geq \quad W_i^{\tilde{c}_N + N\delta}(t) - \sum_{j=1}^{N} U_{S_j}^{\rho_{S_j} - \delta}(t) \\ &\quad - W_i^{\tilde{c}_{m^\star} - m^\star \epsilon}(t) - \sum_{j=1}^{N-1} W_{S_j}^{\rho_{S_j} + \epsilon}(t) - \sum_{j=1}^{N} \sum_{q \in S_j} W_q^{\tilde{\phi}_q}(t). \end{aligned}$$

Using (3.1) and (3.34) yields,

$$V_{i,N}(t) \geq W_i^{\tilde{c}_{m^\star} - m^\star \epsilon, \tilde{c}_N + N\delta}(t) - Y^{\delta,\epsilon}(t).$$

Hence, the lower bound is given by

$$\mathbb{P}(V_{i,N} > x) \geq \mathbb{P}\left(\left\{W_i^{\tilde{c}_{m^\star} - m^\star \epsilon, \tilde{c}_N + N\delta} > x + y\right\} \cap \left\{Y^{\delta,\epsilon} \leq y\right\}\right).$$

Because $Y^{\delta,\epsilon}$ is independent of the traffic process of flow i, the proof is completed. □

Note that the lower bound which we found for $V_{i,2}(t)$ in Lemma 3.3.2 is indeed a special case of the lower bound for $V_{i,N}(t)$.

The upper bound for $\mathbb{P}(V_{i,N} > x)$ is given in the following lemma.

Lemma 3.7.5 *For any $\eta, \nu > 0$ sufficiently small and any y,*

$$\mathbb{P}(V_{i,N} > x) \leq \mathbb{P}\left(W_i^{\tilde{c}_{m^\star} + m^\star \eta, \tilde{c}_N - N\nu} > x - y\right) + \mathbb{P}\left(W_i^{\tilde{\phi}_i} > x\right) \mathbb{P}\left(Z^{\eta,\nu} > y\right).$$

Proof By definition,

$$V_{i,N}(t) \leq V_N(t).$$

Thus, because of Lemma 3.7.1,

$$V_{i,N}(t) \leq X_N(t) - X_{N-1}(t).$$

We use the upper bound in Lemma 3.7.3 for $X_N(t)$ and the lower bound in Lemma 3.7.2 for $X_{N-1}(t)$. Choosing $\xi_j = \rho_{S_j} + \nu$ in Lemma 3.7.3, and $\theta_j = \rho_{S_j} - \eta$ in Lemma 3.7.2, we obtain for $\eta, \nu > 0$ sufficiently small,

$$V_{i,N}(t) \leq W_i^{\tilde{c}_N - N\nu}(t) + \sum_{j=1}^{N} W_{S_j}^{\rho_{S_j}+\nu}(t) - W_i^{\tilde{c}_{m^\star}+m^\star\eta}(t) + \sum_{j=1}^{N-1} U_{S_j}^{\rho_{S_j}-\eta}(t).$$

Using (3.1) and (3.35) yields,

$$V_{i,N}(t) \leq W_i^{\tilde{c}_{m^\star}+m^\star\eta, \tilde{c}_N - N\nu}(t) + Z^{\eta,\nu}(t).$$

Combining the above bound with the upper bound in Lemma 3.6.3, we obtain the following upper bound for $\mathbb{P}(V_{i,N} > x)$:

$$\mathbb{P}\left(\left\{W_i^{\tilde{\phi}_i} > x\right\} \cap \left(\left\{W_i^{\tilde{c}_{m^\star}+m^\star\eta,\tilde{c}_N-N\nu} > x - y\right\} \cup \{Z^{\eta,\nu} > y\}\right)\right).$$

Finally, we use that $Z^{\eta,\nu}$ is independent of the traffic process of flow i. □

Again note that the upper bound for $V_{i,2}(t)$ in Lemma 3.3.3 is a special case of the upper bound for $V_{i,N}(t)$.

We are now able to characterize the tail behavior of $\mathbb{P}(V_{i,N} > x)$. It follows immediately from Lemma 3.4.1 and the lower and upper bound given in Lemmas 3.7.4 and 3.7.5.

Theorem 3.7.1 *Under Assumptions 3.2.1 and 3.2.2,*

$$\mathbb{P}(V_{i,N} > x) \sim \mathbb{P}(W_i^{\tilde{c}_{m^\star},\tilde{c}_N} > x).$$

Remarkably, the workload distribution of flow i at the bottleneck node is asymptotically equivalent to that in a *two*-node tandem network where flow i is served in isolation at constant rates. In Sections 3.4 and 3.5 these rates are simply $\tilde{c}_1$ and $\tilde{c}_2$. For the N-node network we have to take the two smallest service rates for flow i when reduced by the aggregate average rates of the other flows, $\tilde{c}_{m^\star}$ and $\tilde{c}_N$. Hence, for the network described in this section as well, the workload of flow i at the bottleneck node is only affected by the characteristics of the other flows through their average rates. This suggests that an extremely large workload of flow i at its bottleneck node is most likely due to either a long on period or a large burst of the flow itself, while the other flows show roughly their average behavior. Consequently, we can consider flow i to be served in isolation at constant rates $\tilde{c}_1, \ldots, \tilde{c}_N$. Following the reasoning of [71] as in the proof of Theorem 3.6.1 we can then remove all nodes with capacity $\tilde{c}_n > \tilde{c}_{m^\star}$, after which we are left with a two-node tandem network.

3.8 General network with splitting

In this section we extend the model of the previous section and assume that each node on the path of flow i is fed by an additional set of flows, which can leave this path before node N (see Fig. 3.6 for the case where flow i traverses 4 nodes). As before, we derive an upper and a lower bound for $\mathbb{P}(V_{i,N} > x)$ and we use Lemma 3.4.1 to determine the tail behavior of this distribution. Again we defer most of the proofs to the appendix.

We first introduce some additional notation. Define $\hat{A}_k^p(s,t)$ to be the amount of work arriving at node k during the interval $(s,t]$ associated with flows entering the path of flow i at node k and passing through node $p \geq k$, i.e.,

$$\hat{A}_k^p(s,t) := \sum_{m=p}^{N} A_{S_k^m}(s,t).$$

Define $\hat{A}_{k,n}^p(s,t)$ to be the amount of work arriving at node n during the interval $(s,t]$ associated with flows entering the path of flow i *before or at* node k and passing through node $p \geq n \geq k$, i.e.,

$$\hat{A}_{k,n}^p(s,t) := \sum_{j=1}^{k} \sum_{m=p}^{N} A_{S_j^m,n}(s,t). \tag{3.36}$$

Similarly, we define $V_k^p(t)$ to be the workload at node k at time t associated with flows passing through node $p \geq k$ (including flow i), i.e.,

$$V_k^p(t) := \sum_{j=1}^{k} \sum_{m=p}^{N} V_{S_j^m,k}(t) + V_{i,k}(t).$$

Note that we deviate here from the notation as introduced in Section 1.3 in the introduction. Finally, we define $c_k^p(s,t)$ to be the amount of service available in node k during the interval $(s,t]$ for flows passing through node $p \geq k$, i.e.,

$$c_k^p(s,t) := c_k(t-s) - \sum_{j=1}^{k} \sum_{m=k}^{p-1} B_{S_j^m,k}(s,t). \tag{3.37}$$

The following lemma expresses the workload at node n at time t associated with

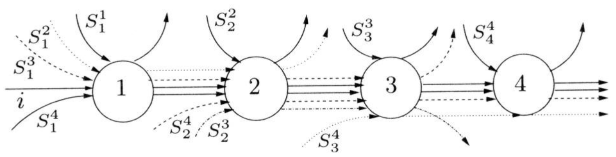

Figure 3.6: General network with splitting.

the flows passing through node p, in terms of $X_n^p(t)$, with

$$X_n^p(t) := \sup_{0 \le s_1 \le \ldots \le s_{n+1} = t} \left\{ A_i(s_1, t) + \sum_{k=1}^{n} \left(\hat{A}_k^p(s_k, t) - c_k^p(s_k, s_{k+1}) \right) \right\}.$$

Lemma 3.8.1 *For $2 \le n \le p$,*

$$V_n^p(t) = X_n^p(t) - X_{n-1}^p(t).$$

Taking p equal to N in the previous lemma, and

$$\sum_{j=1}^{k} \sum_{m=k}^{N-1} B_{S_j^m, k}(s, t) = 0$$

in (3.37) (which yields $c_k^p(s_k, s_{k+1}) = c_k(s_{k+1} - s_k)$ for $k = 1, \ldots, N-1$), the result reduces to that in Lemma 3.7.1, where we assumed that flows cannot leave the path of flow i before node N.

Before presenting the bounds for $X_n^p(t)$ we first introduce some additional notation. Let R be the index set of the flows and γ, ζ and $\psi \in \mathbb{R}^R$. For any vector $x \in \mathbb{R}^R$, denote $x_{S_j^m} = \sum_{q \in S_j^m} x_q$. Define

$$d_k^p := c_k - \sum_{j=1}^{k} \sum_{m=k}^{N} \gamma_{S_j^m} - \sum_{f=1}^{k-1} \sum_{j=1}^{f} \sum_{m=f}^{p-1} \left(\gamma_{S_j^m} - \psi_{S_j^m} \right),$$

and

$$e_k^p := c_k - \sum_{j=1}^{k} \sum_{m=k}^{N} \zeta_{S_j^m} - \sum_{f=1}^{k} \sum_{j=1}^{f} \sum_{m=f}^{p-1} \left(\zeta_{S_j^m} - \gamma_{S_j^m} \right).$$

In the next lemma we present the lower bound for $X_n^p(t)$.

Lemma 3.8.2 *For $n \le p$ and $\gamma_q \le \tilde{\phi}_q$,*

$$X_n^p(t) \ge W_i^{(e_n^p)^\star}(t) - \sum_{k=1}^{n} \sum_{m=p}^{N} U_{S_k^m}^{\zeta_{S_k^m}}(t) - \sum_{k=1}^{n} \sum_{j=1}^{k} \sum_{m=k}^{p-1} \sum_{q \in S_j^m} \left(U_q^{\zeta_q}(t) + W_q^{\gamma_q}(t) \right),$$

with $(e_n^p)^\star := \min_{k=1,\ldots,n} \{e_k^p\}$.

In the following lemma an upper bound is given for $X_n^p(t)$.

Lemma 3.8.3 *For $n \le p$ and $\gamma_q \le \tilde{\phi}_q$,*

$$\begin{aligned} X_n^p(t) \quad &\le \quad W_i^{(d_n^p)^\star}(t) + \sum_{k=1}^{n} \sum_{m=p}^{N} W_{S_k^m}^{\gamma_{S_k^m}}(t) \\ &+ \sum_{k=1}^{n} \sum_{j=1}^{k} \sum_{m=k}^{p-1} \sum_{q \in S_j^m} \left(U_q^{\psi_q}(t) + W_q^{\gamma_q}(t) \right), \end{aligned}$$

with $(d_n^p)^\star := \min_{k=1,\ldots,n} \{d_k^p\}$.

Note that if we take $p = N$ in Lemmas 3.8.2 and 3.8.3 and omit all terms concerning flows that leave before node N, then we obtain Lemmas 3.7.2 and 3.7.3, respectively. The additional terms reflect the fluctuations of the service capacities available for flow i. Before deriving bounds for the workload of flow i in node N at time t, we first introduce some additional notation. For $\delta, \epsilon > 0$, redefine $Y^{\delta,\epsilon}$ as a random variable with the limiting distribution of $Y^{\delta,\epsilon}(t)$ for $t \to \infty$, with

$$Y^{\delta,\epsilon}(t) := \sum_{k=1}^{N-1}\sum_{j=1}^{k}\sum_{m=k}^{N-1}\sum_{q\in S_j^m}\left(U_q^{\rho_q-\delta}(t) + U_q^{\rho_q-\epsilon}(t) + W_q^{\rho_q+\delta}(t) + W_q^{\rho_q+\epsilon}(t)\right)$$

$$+\sum_{k=1}^{N} U_{S_k^N}^{\rho_{S_k^N}-\delta|S_k^N|}(t) + \sum_{k=1}^{N-1} W_{S_k^N}^{\rho_{S_k^N}+\epsilon|S_k^N|}(t) + \sum_{j=1}^{N}\sum_{q\in S_j^N} W_q^{\tilde{\phi}_q}(t). \quad (3.38)$$

For $\eta, \nu > 0$, redefine $Z^{\eta,\nu}$ as a random variable with the limiting distribution of $Z^{\eta,\nu}(t)$ for $t \to \infty$, with

$$Z^{\eta,\nu}(t) := \sum_{k=1}^{N-1}\sum_{j=1}^{k}\sum_{m=k}^{N-1}\sum_{q\in S_j^m}\left(U_q^{\rho_q-\eta}(t) + U_q^{\rho_q-\nu}(t) + W_q^{\rho_q+\eta}(t) + W_q^{\rho_q+\nu}(t)\right)$$
$$+\sum_{k=1}^{N} W_{S_k^N}^{\rho_{S_k^N}+\nu|S_k^N|}(t) + \sum_{k=1}^{N-1} U_{S_k^N}^{\rho_{S_k^N}-\eta|S_k^N|}(t). \quad (3.39)$$

Also define

$$\sigma_k := \sum_{j=1}^{k}\sum_{m=k}^{N}|S_j^m| + 2\sum_{f=1}^{k-1}\sum_{j=1}^{f}\sum_{m=f}^{N-1}|S_j^m|. \quad (3.40)$$

In the next lemma the lower bound for $\mathbb{P}(V_{i,N} > x)$ is given.

Lemma 3.8.4 *For any $\delta, \epsilon > 0$ sufficiently small and any y,*

$$\mathbb{P}(V_{i,N} > x) \geq \mathbb{P}\left(W_i^{\tilde{c}_{m^\star}-\epsilon\sigma_{m^\star},\tilde{c}_N+\delta\sigma_N} > x + y\right)\mathbb{P}\left(Y^{\delta,\epsilon} \leq y\right).$$

Proof By definition,

$$V_{i,N}(t) = V_N(t) - \sum_{j=1}^{N}\sum_{q\in S_j^N} V_{q,N}(t) = V_N^N(t) - \sum_{j=1}^{N}\sum_{q\in S_j^N} V_{q,N}(t).$$

Using Lemmas 3.6.3 and 3.8.1 this is lower bounded by

$$X_N^N(t) - X_{N-1}^N(t) - \sum_{j=1}^{N}\sum_{q\in S_j^N} W_q^{\tilde{\phi}_q}(t).$$

Now we can use the lower bound for $X_N^N(t)$ as given in Lemma 3.8.2 and the upper bound for $X_{N-1}^N(t)$ as given in Lemma 3.8.3. In Lemma 3.8.2 we take $\zeta_q = \rho_q - \delta$ and $\gamma_q = \rho_q + \delta$, hence,

$$\begin{aligned} X_N^N(t) &\geq W_i^{e_{N^\star}^N}(t) - \sum_{k=1}^{N} U_{S_k^N}^{\rho_{S_k^N} - \delta|S_k^N|}(t) \\ &\quad - \sum_{k=1}^{N}\sum_{j=1}^{k}\sum_{m=k}^{N-1}\sum_{q\in S_j^m} \left(U_q^{\rho_q-\delta}(t) + W_q^{\rho_q+\delta}(t)\right), \end{aligned}$$

with for $k = 1, \ldots, N$,

$$\begin{aligned} e_k^N &= c_k - \sum_{j=1}^{k}\sum_{m=k}^{N}(\rho_{S_j^m} - |S_j^m|\delta) + 2\delta\sum_{f=1}^{k-1}\sum_{j=1}^{f}\sum_{m=f}^{N-1}|S_j^m| \\ &= \tilde{c}_k + \delta\sigma_k, \end{aligned}$$

and thus $e_{N^\star}^N = \tilde{c}_N + \delta\sigma_N$ for $\delta > 0$ sufficiently small. Analogously, we take $\gamma_q = \rho_q + \epsilon$ and $\psi_q = \rho_q - \epsilon$ in Lemma 3.8.3, hence,

$$\begin{aligned} X_{N-1}^N(t) &\leq W_i^{d_{N-1^\star}^N}(t) + \sum_{k=1}^{N-1} W_{S_k^N}^{\rho_{S_k^N} + \epsilon|S_k^N|}(t) \\ &\quad + \sum_{k=1}^{N-1}\sum_{j=1}^{k}\sum_{m=k}^{N-1}\sum_{q\in S_j^m} \left(U_q^{\rho_q-\epsilon}(t) + W_q^{\rho_q+\epsilon}(t)\right), \end{aligned}$$

with for $k = 1, \ldots, N-1$,

$$\begin{aligned} d_k^N &= c_k - \sum_{j=1}^{k}\sum_{m=k}^{N}(\rho_{S_j^m} + |S_j^m|\epsilon) - 2\epsilon\sum_{f=1}^{k-1}\sum_{j=1}^{f}\sum_{m=f}^{N-1}|S_j^m| \\ &= \tilde{c}_k - \epsilon\sigma_k, \end{aligned}$$

and thus $d_{N-1^\star}^N = \tilde{c}_{m^\star} - \epsilon\sigma_{m^\star}$ for $\epsilon > 0$ sufficiently small. Then using (3.1) and (3.38) we obtain,

$$V_{i,N}(t) \geq W_i^{\tilde{c}_{m^\star} - \epsilon\sigma_{m^\star}, \tilde{c}_N + \delta\sigma_N}(t) - Y^{\delta,\epsilon}(t).$$

Hence, the lower bound is given by,

$$\mathbb{P}(V_{i,N} > x) \geq \mathbb{P}\left(\left\{W_i^{\tilde{c}_{m^\star} - \epsilon\sigma_{m^\star}, \tilde{c}_N + \delta\sigma_N} > x + y\right\} \cap \left\{Y^{\delta,\epsilon} \leq y\right\}\right).$$

Because $Y^{\delta,\epsilon}$ is independent of the traffic process of flow i, the result follows immediately. □

The upper bound for $\mathbb{P}(V_{i,N} > x)$ is given in the following lemma.

Lemma 3.8.5 *For any $\eta, \nu > 0$ sufficiently small and any y,*

$$\mathbb{P}(V_{i,N} > x) \leq \mathbb{P}\left(W_i^{\tilde{c}_{m^\star}+\eta\sigma_{m^\star}, \tilde{c}_N - \nu\sigma_N} > x - y\right) + \mathbb{P}\left(W_i^{\tilde{\phi}_i} > x\right) \mathbb{P}(Z^{\eta,\nu} > y).$$

Proof By definition,

$$V_{i,N}(t) \leq V_N(t) = V_N^N(t).$$

Using Lemma 3.8.1,

$$V_{i,N}(t) \leq X_N^N(t) - X_{N-1}^N(t).$$

Analogously to the proof of Lemma 3.8.4, we use the upper bound for $X_N^N(t)$ as given in Lemma 3.8.3 and the lower bound for $X_{N-1}^N(t)$ as given in Lemma 3.8.2. In Lemma 3.8.3 we take $\gamma_q = \rho_q + \nu$ and $\psi_q = \rho_q - \nu$. In Lemma 3.8.2 we take $\zeta_q = \rho_q - \eta$ and $\gamma_q = \rho_q + \eta$. Using (3.1) and (3.39) yields,

$$V_{i,N}(t) \leq W_i^{\tilde{c}_{m^\star}+\eta\sigma_{m^\star}, \tilde{c}_N - \nu\sigma_N}(t) + Z^{\eta,\nu}(t).$$

Combining the above bound with the upper bound in Lemma 3.6.3, we obtain the following upper bound for $\mathbb{P}(V_{i,N} > x)$:

$$\mathbb{P}\left(\left\{W_i^{\tilde{\phi}_i} > x\right\} \cap \left(\left\{W_i^{\tilde{c}_{m^\star}+\eta\sigma_{m^\star}, \tilde{c}_N - \nu\sigma_N} > x - y\right\} \cup \{Z^{\eta,\nu} > y\}\right)\right).$$

Finally, we use that $Z^{\eta,\nu}$ is independent of the traffic process of flow i. □

Note that the bounds for $V_{i,N}(t)$ as given in Lemmas 3.8.4 and 3.8.5 reduce to the bounds in Lemmas 3.7.4 and 3.7.5, if we assume that no flows leave the path of flow i, i.e., if $S_j^m = \emptyset$ for $m < N$.

We have now gathered all the elements to characterize the tail behavior of the workload distribution in the most general class of networks that we consider.

Theorem 3.8.1 *Under Assumptions 3.2.1 and 3.2.2,*

$$\mathbb{P}(V_{i,N} > x) \sim \mathbb{P}\left(W_i^{\tilde{c}_{m^\star}, \tilde{c}_N} > x\right).$$

Again the workload distribution of flow i at the bottleneck node is asymptotically equivalent to that in a *two-node* tandem network where flow i is served in isolation at constant rates.

3.9 Concluding remarks

In this chapter we analyzed the workload behavior under the GPS mechanism in networks fed by multiple flows. Specifically, we considered a particular flow i traversing the network and assumed it to have heavy-tailed traffic characteristics. We distinguished between two configurations of feed-forward networks: (i) other flows follow the path of flow i when they feed into any of the nodes on this path and (ii) other flows can leave the path of flow i. In addition, we considered two traffic scenarios for flow i: (i) flow i generates instantaneous traffic bursts and (ii) flow i generates traffic according to an on-off process. Under these conditions we showed that the tail behavior of the workload distribution of flow i at its bottleneck node is equivalent to that in a *two*-node tandem network where flow i is served in isolation at *constant* rates. In case flow i traverses only two nodes and the second node is the bottleneck node, these rates are the service rates in the original network reduced by the average rates of the other flows. However, when flow i traverses more than two nodes, we have to take the rates from the nodes which are bottleneck when the service rate is reduced by the average rates of the other flows. Hence, asymptotically flow i is only affected by the characteristics of the other flows through their average rates. This suggests that the GPS mechanism is capable of isolating individual flows in networks, even when they have heavy-tailed traffic characteristics, while achieving significant multiplexing gains.

The results of this chapter may be extended in several directions. First of all, we assumed for each of the flows the minimum guaranteed service rate by the GPS mechanism to exceed the average input rate (Assumption 3.2.1). This assumption is somewhat restrictive for best-effort flows, and may actually be considerably weakened. If Assumption 3.2.1 is relaxed for flows other than i, then these flows may be potentially unstable, or temporarily become so when flow i is backlogged. We expect that the results for the workload behavior of flow i do not drastically change in that case, but are only influenced through a larger value of $\tilde{c}_n$. In contrast, if Assumption 3.2.1 is not satisfied for flow i itself, then flow i relies on surplus capacity from other flows for stability, and may thus be potentially vulnerable. We conjecture that flow i may then be strongly affected by interfering flows with 'heavier'-tailed characteristics, as is suggested by the single-node results in [21]. Secondly, we only considered in this chapter the workload distribution at 'bottleneck' nodes with the minimum average service rate for flow i on its path (Assumption 3.2.2). The tail behavior of the workload distribution of flow i at non-bottleneck nodes remains an interesting topic for further research.

Appendix

3.A Proof of Lemma 3.2.1

Lemma *Assuming* $V_{Q,n}(0) = 0$,

$$\begin{aligned} V_{Q,n}(t) &= \sup_{0 \le s \le t} \{A_{Q,n}(s,t) - B_{Q,n}(s,t)\} \\ &= \sup_{0 \le s \le t} \left\{A_{Q,n}(s,t) - (c_n(t-s) - B_{S^{(n)}\setminus Q,n}(s,t))\right\}. \end{aligned}$$

Proof We show

(i)

$$V_{Q,n}(t) \le \sup_{0 \le s \le t} \left\{A_{Q,n}(s,t) - (c_n(t-s) - B_{S^{(n)}\setminus Q,n}(s,t))\right\},$$

(ii)

$$\begin{aligned} & \sup_{0\le s\le t} \left\{A_{Q,n}(s,t) - (c_n(t-s) - B_{S^{(n)}\setminus Q,n}(s,t))\right\} \\ \le & \sup_{0\le s\le t} \left\{A_{Q,n}(s,t) - B_{Q,n}(s,t)\right\} \end{aligned}$$

and (iii)

$$\sup_{0\le s\le t} \left\{A_{Q,n}(s,t) - B_{Q,n}(s,t)\right\} \le V_{Q,n}(t).$$

Ad (i). Define

$$s^\star := \max\left\{s \mid V_{Q,n}(s) = 0, 0 \le s \le t\right\},$$

i.e., $s^\star$ is the last time before t at which the workload of all the flows $q \in Q$ at node n was 0. Note that $s^\star$ is well-defined since $V_{Q,n}(0) = 0$. Because of the definition of $s^\star$, $V_{Q,n}(s) > 0$ for all $s \in (s^\star, t]$. Recall that the GPS mechanism is work-conserving, so that

$$B_{Q,n}(s^\star,t) + B_{S^{(n)}\setminus Q,n}(s^\star,t) = c_n(t-s^\star),$$

and hence,

$$\begin{aligned} V_{Q,n}(t) &= A_{Q,n}(s^\star,t) + V_{Q,n}(s^\star) - B_{Q,n}(s^\star,t) \\ &= A_{Q,n}(s^\star,t) - (c_n(t-s^\star) - B_{S^{(n)}\setminus Q,n}(s^\star,t)) \\ &\le \sup_{0\le s\le t} \left\{A_{Q,n}(s,t) - (c_n(t-s) - B_{S^{(n)}\setminus Q,n}(s,t))\right\}. \end{aligned}$$

Ad (ii). By definition,

$$B_{Q,n}(s,t) \le c_n(t-s) - B_{S^{(n)}\setminus Q,n}(s,t),$$

for all $s \in [0,t]$.
Ad (iii). From (3.5),

$$V_{Q,n}(t) \ge A_{Q,n}(s,t) - B_{Q,n}(s,t),$$

for all $s \in [0,t]$. Hence,

$$V_{Q,n}(t) \ge \sup_{0\le s\le t} \left\{A_{Q,n}(s,t) - B_{Q,n}(s,t)\right\},$$

for all $t \ge 0$. □

3.B Proof of Lemma 3.6.1

Lemma *For $q \in S^p_m$, $1 \le m \le n \le p$ and $\gamma_q \le \tilde{\phi}_q$,*

$$B_{q,n}(s,t) \ge \gamma_q(t-s) - \sup_{s\le s_m\le t} \left\{\gamma_q(s_m - s) - A_q(s,s_m)\right\}.$$

Proof We will prove by induction on r that for each $r \in \{0, \ldots, n-m\}$,

$$B_{q,n}(s,t) \geq \gamma_q(t-s) - \sup_{s \leq s_{n-r} \leq t} \{\gamma_q(s_{n-r} - s) - A_{q,n-r}(s, s_{n-r})\}, \tag{3.41}$$

which gives immediately the desired result for $r = n - m$.

For $r = 0$, (3.41) reduces to

$$B_{q,n}(s,t) \geq \gamma_q(t-s) - \sup_{s \leq s_n \leq t} \{\gamma_q(s_n - s) - A_{q,n}(s, s_n)\}, \tag{3.42}$$

which can be verified as follows. We distinguish between two cases. Case (i). We have $V_{q,n}(s_n) > 0$ for all $s_n \in [s,t]$, i.e., flow q is continuously backlogged at node n during $[s,t]$. Hence,

$$B_{q,n}(s,t) \geq \phi_{q,n}(t-s) \geq \tilde{\phi}_q(t-s) \geq \gamma_q(t-s),$$

which immediately gives (3.42). Case (ii). Now $V_{q,n}(s_n)$ is equal to 0 for some $s_n \in [s,t]$. Defining

$$s_n^\star := \max\{s_n \mid V_{q,n}(s_n) = 0, 0 \leq s_n \leq t\},$$

we have,

$$\begin{aligned} B_{q,n}(s,t) &= B_{q,n}(s, s_n^\star) + B_{q,n}(s_n^\star, t) \\ &= V_{q,n}(s) + A_{q,n}(s, s_n^\star) - V_{q,n}(s_n^\star) + B_{q,n}(s_n^\star, t). \end{aligned}$$

Since $V_{q,n}(s_n^\star) = 0$ and flow q is continuously backlogged at node n during $(s_n^\star, t]$, this is lower bounded by

$$\begin{aligned} & A_{q,n}(s, s_n^\star) + \phi_{q,n}(t - s_n^\star) \geq A_{q,n}(s, s_n^\star) + \gamma_q(t - s_n^\star) \\ = \; & \gamma_q(t-s) - (\gamma_q(s_n^\star - s) - A_{q,n}(s, s_n^\star)) \\ \geq \; & \gamma_q(t-s) - \sup_{s \leq s_n \leq t} \{\gamma_q(s_n - s) - A_{q,n}(s, s_n)\}. \end{aligned}$$

Now assume (3.41) to hold for $r-1$, i.e.,

$$B_{q,n}(s,t) \geq \gamma_q(t-s) - \sup_{s \leq s_{n-r+1} \leq t} \{\gamma_q(s_{n-r+1} - s) - A_{q,n-r+1}(s, s_{n-r+1})\}. \tag{3.43}$$

As in (3.42),

$$\begin{aligned} B_{q,n-r}(s, s_{n-r+1}) \geq \; & \gamma_q(s_{n-r+1} - s) \\ & - \sup_{s \leq s_{n-r} \leq s_{n-r+1}} \{\gamma_q(s_{n-r} - s) - A_{q,n-r}(s, s_{n-r})\}. \end{aligned}$$

Using (3.6) to substitute $B_{q,n-r}(s, s_{n-r+1})$ for $A_{q,n-r+1}(s_n, s_{n-r+1})$ in (3.43) yields (3.41). □

3.C Proof of Lemma 3.6.2

Lemma *For $q \in S_m^p$, $1 \le m \le n \le p$ and $\gamma_q \le \tilde{\phi}_q$,*

$$\sum_{j=m}^{n} V_{q,j}(t) \le W_q^{\gamma_q}(t).$$

Proof By induction on r we prove that for each $r \in \{0, \dots, n-m\}$,

$$V_{q,n}(t) = \sum_{j=n-r}^{n} V_{q,j}(s) + A_{q,n-r}(s,t) - B_{q,n}(s,t) - \sum_{j=n-r}^{n-1} V_{q,j}(t), \qquad (3.44)$$

for all $s \in [0, t]$.

For $r = 0$, (3.44) reduces to (3.5).

Now assume (3.44) to hold for $r - 1$. Substituting (3.6) for $A_{q,n-r+1}(s,t)$, we immediately obtain (3.44).

Taking $r = n - m$ in (3.44) and choosing time s such that $\sum_{j=m}^{n} V_{q,j}(s) = 0$ (for example $s = 0$) yields,

$$\sum_{j=m}^{n} V_{q,j}(t) = A_q(s,t) - B_{q,n}(s,t).$$

Rewriting the lower bound for $B_{q,n}(s,t)$ in Lemma 3.6.1 to

$$- \sup_{s \le s_m \le t} \{-A_q(s, s_m) - \gamma_q(t - s_m)\},$$

we obtain,

$$\begin{aligned} \sum_{j=m}^{n} V_{q,j}(t) &\le A_q(s,t) + \sup_{s \le s_m \le t} \{-A_q(s, s_m) - \gamma_q(t - s_m)\} \\ &= \sup_{s \le s_m \le t} \{A_q(s_m, t) - \gamma_q(t - s_m)\}, \end{aligned}$$

and the proof is completed. □

3.D Proof of Lemma 3.6.4

Lemma *For $q \in S_m^p$, $1 \le m \le n \le p$ and $\gamma_q \le \tilde{\phi}_q$,*

$$B_{q,n}(s,t) \le \gamma_q(t-s) + \sup_{0 \le s_m \le t} \{A_q(s_m, t) - \gamma_q(t - s_m)\}.$$

Proof We first prove by induction on r that for each $r \in \{0, \dots, n-m\}$,

$$B_{q,n}(s,t) \le A_{q,n-r}(s,t) + \sum_{j=n-r}^{n} V_{q,j}(s). \qquad (3.45)$$

For $r = 0$, (3.45) reduces to the upper bound which immediately follows from (3.5).

Now assume (3.45) to hold for $r-1$. Substituting (3.6) for $A_{q,n-r}(s,t)$ yields (3.45).

Finally, taking $r = n-m$ in (3.45) and using Lemma 3.6.2, we obtain

$$B_{q,n}(s,t) \leq A_q(s,t) + W_q^{\gamma_q}(s) \leq A_q(s,t) + \sup_{0 \leq s_m \leq t} \{A_q(s_m, s) - \gamma_q(s - s_m)\}.$$

□

3.E Proof of Equation (3.33)

For convenience we restate Equation (3.33) as a lemma.

Lemma

$$\sum_{j=1}^{m^\star} V_{i,j}(t) = \sup_{0 \leq s \leq t} \{A_i(s,t) - c_{m^\star}(t-s)\}.$$

Proof As the next formula is a special case of (3.47), which is proved in the proof of Lemma 3.7.1, we present it here without proof. (Note that we do not use Equation (3.33) in the proof of Lemma 3.7.1.) For any $n \geq 1$,

$$\sum_{j=1}^{n} V_{i,j}(t) = \sup_{0 \leq s_1 \leq \ldots \leq s_n \leq s_{n+1} = t} \left\{ A_i(s_1, t) - \sum_{j=1}^{n} c_j (s_{j+1} - s_j) \right\} =: W_i^{(n)}(t).$$

We now show with a lower and upper bound that

$$W_i^{(n)}(t) = W_i^{c_{n^\star}}(t), \tag{3.46}$$

with $c_{n^\star} := \min_{j=1,\ldots,n}\{c_j\}$. We first prove that the right-hand side is a lower bound for the left-hand side. Imposing a restriction on the optimizing arguments, the supremum becomes smaller. Hence, choosing $s = s_1 = \ldots = s_{j^\star}$ and $s_{j^\star+1} = \ldots = t$,

$$\sup_{0 \leq s_1 \leq \ldots \leq s_{n+1} = t} \left\{ A_i(s_1, t) - \sum_{j=1}^{n} c_j (s_{j+1} - s_j) \right\} \geq \sup_{0 \leq s \leq t} \{A_i(s,t) - c_{n^\star}(t-s)\}.$$

Next we show that the right-hand side is in fact also an upper bound. Because $c_j \geq c_{n^\star}$ for all $j = 1, \ldots, n$,

$$\sum_{j=1}^{n} c_j (s_{j+1} - s_j) \geq \sum_{j=1}^{n} c_{n^\star} (s_{j+1} - s_j) = c_{n^\star}(t - s_1),$$

and the proof is completed. □

3.F Proof of Lemma 3.7.1

Lemma *For $n \geq 2$,*

$$V_n(t) = X_n(t) - X_{n-1}(t).$$

Proof Note that $S^{(n-m)} = S^{(n-m-1)} \cup S_{n-m}$ and $S^{(n-m-1)} \cap S_{n-m} = \emptyset$. We prove by induction on m that for each $m \in \{0, \ldots, n-1\}$,

$$\begin{aligned} V_n(t) &= \sup_{0 \leq s_{n-m} \leq \cdots \leq s_n \leq s_{n+1} = t} \Bigg\{ A_{S^{(n-m-1)},n-m}(s_{n-m}, t) \\ &\quad + \sum_{j=n-m}^{n} (A_{S_j}(s_j, t) - c_j(s_{j+1} - s_j)) \Bigg\} - \sum_{j=n-m}^{n-1} V_j(t), \quad (3.47) \end{aligned}$$

with the notational convention that $S^{(0)} = \{i\}$.

If $m = 0$, then (3.47) reduces to (3.7).

Now assume (3.47) to hold for $m-1$. Using the relation in (3.6) to substitute for $A_{S^{(n-m)},n-m+1}(s_{n-m+1}, t)$, $V_n(t)$ is equal to

$$\begin{aligned} &\sup_{0 \leq s_{n-m+1} \leq \cdots \leq s_n \leq s_{n+1} = t} \Bigg\{ A_{S^{(n-m)},n-m}(s_{n-m+1}, t) + V_{n-m}(s_{n-m+1}) \\ &-V_{n-m}(t) + \sum_{j=n-m+1}^{n} (A_{S_j}(s_j, t) - c_j(s_{j+1} - s_j)) \Bigg\} - \sum_{j=n-m+1}^{n-1} V_j(t). \end{aligned}$$

Substituting (3.7) for $V_{n-m}(s_{n-m+1})$, and arranging the terms yields (3.47).

Taking $m = n - 1$ in (3.47), we obtain

$$V_n(t) = X_n(t) - \sum_{m=1}^{n-1} V_m(t),$$

so that $V_n(t) = X_n(t) - X_{n-1}(t)$. □

3.G Proof of Lemma 3.7.2

Lemma *For any $\theta_1, \ldots, \theta_n$, with $e := \min_{m=1,\ldots,n}\{c_m - \sum_{j=1}^{m} \theta_j\}$,*

$$X_n(t) \geq W_i^e(t) - \sum_{j=1}^{n} U_{S_j}^{\theta_j}(t).$$

Proof Using the fact that $\sum_{j=1}^{n} \theta_j(t - s_j) = \sum_{j=1}^{n} \theta^{(j)}(s_{j+1} - s_j)$ with $\theta^{(j)} := \sum_{m=1}^{j} \theta_m$, we write

$$X_n(t) = \sup_{0 \leq s_1 \leq \cdots \leq s_{n+1} = t} \Bigg\{ A_i(s_1, t) - \sum_{j=1}^{n} (c_j - \theta^{(j)})(s_{j+1} - s_j)$$

$$+\sum_{j=1}^{n}\left(A_{S_j}(s_j,t)-\theta_j(t-s_j)\right)\Bigg\}.$$

Because of (3.10) the right-hand side is lower bounded by

$$\sup_{0\le s_1\le\ldots\le s_{n+1}=t}\left\{A_i(s_1,t)-\sum_{j=1}^{n}(c_j-\theta^{(j)})(s_{j+1}-s_j)\right\}$$
$$-\sum_{j=1}^{n}\sup_{0\le s_j\le t}\left\{\theta_j(t-s_j)-A_{S_j}(s_j,t)\right\}.$$

Using (3.8) and (3.46) for the first supremum, the proof is completed. □

3.H Proof of Lemma 3.7.3

Lemma *For any $\xi_1,\ldots,\xi_n$, with $d:=\min_{m=1,\ldots,n}\{c_m-\sum_{j=1}^{m}\xi_j\}$,*

$$X_n(t)\le W_i^d(t)+\sum_{j=1}^{n}W_{S_j}^{\xi_j}(t).$$

Proof The proof is similar to that of the lower bound. First adding $\sum_{j=1}^{n}\xi_j(t-s_j)=\sum_{j=1}^{n}\xi^{(j)}(s_{j+1}-s_j)$ with $\xi^{(j)}:=\sum_{m=1}^{j}\xi_m$ to $X_n(t)$, then subtracting it again and using (3.46) yields

$$X_n(t)\le\sup_{0\le s\le t}\{A_i(s,t)-(c_k-\xi^{(k)})(t-s)\}+\sum_{j=1}^{n}\sup_{0\le s_j\le t}\{A_{S_j}(s_j,t)-\xi_j(t-s_j)\},$$

and the proof is completed. □

3.I Proof of Lemma 3.8.1

Lemma *For $2\le n\le p$,*

$$V_n^p(t)=X_n^p(t)-X_{n-1}^p(t).$$

Proof First we prove, using induction on r, that for each $r\in\{0,\ldots,n-1\}$,

$$\begin{aligned}V_n^p(t) &= \sup_{0\le s_{n-r}\le\ldots\le s_{n+1}=t}\Bigg\{A_{i,n-r}(s_{n-r},t)+\hat{A}^p_{n-r-1,n-r}(s_{n-r},t)\\ &\quad+\sum_{k=n-r}^{n}\left(\hat{A}_k^p(s_k,t)-c_k^p(s_k,s_{k+1})\right)\Bigg\}-\sum_{k=n-r}^{n-1}V_k^p(t). \qquad (3.48)\end{aligned}$$

For $r = 0$, (3.48) reduces to

$$V_n^p(t) = \sup_{0 \le s_n \le t} \left\{ A_{i,n}(s_n, t) + \hat{A}^p_{n,n}(s_n, t) - c_n^p(t - s_n) \right\},$$

which is true by virtue of Lemma 3.2.1.

Assume (3.48) to hold for $r-1$. Substituting the relation in (3.6) for $A_{i,n-r+1}(s_{n-r+1}, t) + \hat{A}^p_{n-r,n-r+1}(s_{n-r+1}, t)$ yields,

$$\begin{aligned} V_n^p(t) &= \sup_{0 \le s_{n-r+1} \le \cdots \le s_{n+1} = t} \left\{ A_{i,n-r}(s_{n-r+1}, t) + \hat{A}^p_{n-r,n-r}(s_{n-r+1}, t) \right. \\ &\quad \left. + V^p_{n-r}(s_{n-r+1}) + \sum_{k=n-r+1}^{n} \left(\hat{A}_k^p(s_k, t) - c_k^p(s_k, s_{k+1}) \right) \right\} \\ &\quad - \sum_{k=n-r+1}^{n-1} V_k^p(t) - V_{i,n-r}(t) - \sum_{j=1}^{n-r} \sum_{m=p}^{N} V_{S_j^m, n-r}(t). \end{aligned}$$

Using Lemma 3.2.1 to substitute for $V^p_{n-r}(s_{n-r+1})$, i.e.,

$$\begin{aligned} V^p_{n-r}(s_{n-r+1}) &= \sup_{0 \le s_{n-r} \le s_{n-r+1}} \left\{ \hat{A}^p_{n-r,n-r}(s_{n-r}, s_{n-r+1}) \right. \\ &\quad \left. + A_{i,n-r}(s_{n-r}, s_{n-r+1}) - c^p_{n-r}(s_{n-r}, s_{n-r+1}) \right\}, \end{aligned}$$

and rewriting the supremum, we obtain (3.48).

Now taking $r = n-1$ in (3.48) yields $\sum_{k=1}^{n} V_k^p(t) = X_n^p(t)$ for all $n \le p$, and thus we obtain the desired result. □

3.J Proof of Lemma 3.8.2

Lemma *For $n \le p$ and $\gamma_q \le \tilde{\phi}_q$, with $(e_n^p)^\star := \min_{k=1,\ldots,n}\{e_k^p\}$,*

$$X_n^p(t) \ge W_i^{(e_n^p)^\star}(t) - \sum_{k=1}^{n} \sum_{m=p}^{N} U_{S_k^m}^{\zeta_{S_k^m}}(t) - \sum_{k=1}^{n} \sum_{j=1}^{k} \sum_{m=k}^{p-1} \sum_{q \in S_j^m} \left(U_q^{\zeta_q}(t) + W_q^{\gamma_q}(t) \right).$$

Proof Using the lower bound for $B_{q,k}(s_k, s_{k+1})$ as given in Lemma 3.6.1 and using (3.36) and (3.37),

$$\begin{aligned} X_n^p(t) \ge \sup_{0 \le s_1 \le \ldots \le s_{n+1} = t} & \left\{ A_i(s_1, t) + \sum_{k=1}^{n} \left(\sum_{m=p}^{N} A_{S_k^m}(s_k, t) \right. \right. \\ & - c_k(s_{k+1} - s_k) + \sum_{j=1}^{k} \sum_{m=k}^{p-1} \gamma_{S_j^m}(s_{k+1} - s_k) \\ & \left. \left. - \sum_{j=1}^{k} \sum_{m=k}^{p-1} \sum_{q \in S_j^m} \sup_{s_k \le s_m \le s_{k+1}} \{\gamma_q(s_m - s_k) - A_q(s_k, s_m)\} \right) \right\}. \end{aligned}$$

Observe that

$$\zeta_{S_k^m}(t-s_k)=\sum_{j=k}^{n}\zeta_{S_k^m}(s_{j+1}-s_j),$$

which implies that

$$\sum_{k=1}^{n}\sum_{m=p}^{N}\zeta_{S_k^m}(t-s_k)=\sum_{k=1}^{n}\sum_{j=k}^{n}\sum_{m=p}^{N}\zeta_{S_k^m}(s_{j+1}-s_j).$$

First changing the order of summation and then interchanging the indices j and k, the latter term can be written as

$$\sum_{j=1}^{n}\sum_{k=1}^{j}\sum_{m=p}^{N}\zeta_{S_k^m}(s_{j+1}-s_j)=\sum_{k=1}^{n}\sum_{j=1}^{k}\sum_{m=p}^{N}\zeta_{S_j^m}(s_{k+1}-s_k). \tag{3.49}$$

Hence, adding and subtracting $\sum_{k=1}^{n}\sum_{m=p}^{N}\zeta_{S_k^m}(t-s_k)$ yields,

$$\begin{aligned} X_n^p(t) \;\geq\; & \sup_{0\leq s_1\leq\ldots\leq s_{n+1}=t}\Bigg\{A_i(s_1,t)+\sum_{k=1}^{n}\sum_{m=p}^{N}\Big(A_{S_k^m}(s_k,t)-\zeta_{S_k^m}(t-s_k)\Big)\\ & -\sum_{k=1}^{n}\Big(c_k-\sum_{j=1}^{k}\sum_{m=k}^{p-1}\gamma_{S_j^m}-\sum_{j=1}^{k}\sum_{m=p}^{N}\zeta_{S_j^m}\Big)(s_{k+1}-s_k)\\ & -\sum_{k=1}^{n}\sum_{j=1}^{k}\sum_{m=k}^{p-1}\sum_{q\in S_j^m}\sup_{s_k\leq s_m\leq s_{k+1}}\{\gamma_q(s_m-s_k)-A_q(s_k,s_m)\}\Bigg\}. \end{aligned}$$

The inner supremum is upper bounded by

$$\begin{aligned} & \sup_{s_k\leq s_m\leq s_{k+1}}\{\zeta_q(t-s_k)-A_q(s_k,t)\}\\ + & \sup_{s_k\leq s_m\leq s_{k+1}}\{A_q(s_m,t)-\gamma_q(t-s_m)\}+(\gamma_q-\zeta_q)(t-s_k). \end{aligned} \tag{3.50}$$

Because

$$(\gamma_{S_j^m}-\zeta_{S_j^m})(t-s_k)=\sum_{f=k}^{n}(\gamma_{S_j^m}-\zeta_{S_j^m})(s_{f+1}-s_f),$$

we can follow the derivation of (3.49) to obtain

$$\sum_{k=1}^{n}\sum_{j=1}^{k}\sum_{m=k}^{p-1}(\gamma_{S_j^m}-\zeta_{S_j^m})(t-s_k)=\sum_{k=1}^{n}\sum_{f=1}^{k}\sum_{j=1}^{f}\sum_{m=f}^{p-1}(\gamma_{S_j^m}-\zeta_{S_j^m})(s_{k+1}-s_k). \tag{3.51}$$

Then using (3.50) and (3.51), we obtain for the lower bound,

$$\sup_{0\leq s_1\leq\ldots\leq s_{n+1}=t}\Bigg\{A_i(s_1,t)-\sum_{k=1}^{n}e_k(s_{k+1}-s_k)\Bigg\}$$

$$-\sum_{k=1}^{n}\sum_{m=p}^{N}\sup_{0\le s_k\le t}\{\zeta_{S_k^m}(t-s_k)-A_{S_k^m}(s_k,t)\}$$
$$-\sum_{k=1}^{n}\sum_{j=1}^{k}\sum_{m=k}^{p-1}\sum_{q\in S_j^m}\sup_{0\le s_k\le t}\{\zeta_q(t-s_k)-A_q(s_k,t)\}$$
$$-\sum_{k=1}^{n}\sum_{j=1}^{k}\sum_{m=k}^{p-1}\sum_{q\in S_j^m}\sup_{0\le s_m\le t}\{A_q(s_m,t)-\gamma_q(t-s_m)\}\,.$$

Now using (3.8) and (3.46) for the first supremum, the proof is completed. □

3.K Proof of Lemma 3.8.3

Lemma *For $n \le p$ and $\gamma_q \le \tilde{\phi}_q$, with $(d_n^p)^\star := \min_{k=1,\ldots,n}\{d_k^p\}$,*

$$X_n^p(t) \le W_i^{(d_n^p)^\star}(t) + \sum_{k=1}^{n}\sum_{m=p}^{N} W_{S_k^m}^{\theta_{S_k^m}}(t) + \sum_{k=1}^{n}\sum_{j=1}^{k}\sum_{m=k}^{p-1}\sum_{q\in S_j^m}\left(W_q^{\theta_q}(t)+U_q^{\psi_q}(t)\right).$$

Proof Using the upper bound for $B_{q,k}(s_k,s_{k+1})$ as given in Lemma 3.6.4 and using (3.36) and (3.37) yields,

$$\begin{aligned} X_n^p(t) \le & \sup_{0\le s_1\le\ldots\le s_{n+1}=t}\Bigg\{A_i(s_1,t)+\sum_{k=1}^{n}\Bigg(\sum_{m=p}^{N}A_{S_k^m}(s_k,t) \\ & -c_k(s_{k+1}-s_k)+\sum_{j=1}^{k}\sum_{m=k}^{p-1}\gamma_{S_j^m}(s_{k+1}-s_k) \\ & +\sum_{j=1}^{k}\sum_{m=k}^{p-1}\sum_{q\in S_j^m}\sup_{0\le s_m\le s_{k+1}}\{A_q(s_m,s_{k+1})-\gamma_q(s_{k+1}-s_m)\}\Bigg)\Bigg\}. \end{aligned}$$

Analogously to the proof of Lemma 3.8.2, we obtain

$$\sum_{k=1}^{n}\sum_{m=p}^{N}\gamma_{S_k^m}(t-s_k)=\sum_{k=1}^{n}\sum_{j=1}^{k}\sum_{m=p}^{N}\gamma_{S_j^m}(s_{k+1}-s_k).$$

Hence, adding and subtracting $\sum_{k=1}^{n}\sum_{m=p}^{N}\gamma_{S_k^m}(t-s_k)$ yields,

$$\begin{aligned} X_n^p(t) \le & \sup_{0\le s_1\le\ldots\le s_{n+1}=t}\Bigg\{A_i(s_1,t)+\sum_{k=1}^{n}\sum_{m=p}^{N}\left(A_{S_k^m}(s_k,t)-\gamma_{S_k^m}(t-s_k)\right) \\ & -\sum_{k=1}^{n}\left(c_k-\sum_{j=1}^{k}\sum_{m=k}^{p-1}\gamma_{S_j^m}-\sum_{j=1}^{k}\sum_{m=p}^{N}\gamma_{S_j^m}\right)(s_{k+1}-s_k) \end{aligned}$$

$$+\sum_{k=1}^{n}\sum_{j=1}^{k}\sum_{m=k}^{p-1}\sum_{q\in S_j^m}\sup_{0\le s_m\le s_{k+1}}\{A_q(s_m,s_{k+1})-\gamma_q(s_{k+1}-s_m)\}\Bigg\}.$$

The inner supremum is upper bounded by

$$\begin{aligned}&\sup_{0\le s_m\le t}\{A_q(s_m,t)-\gamma_q(t-s_m)\}\\+\quad&\sup_{0\le s_{k+1}\le t}\{\psi_q(t-s_{k+1})-A_q(s_{k+1},t)\}+(\gamma_q-\psi_q)(t-s_{k+1}).\end{aligned}$$

Combining

$$(\gamma_{S_j^m}-\psi_{S_j^m})(t-s_{k+1})=\sum_{f=k+1}^{n}(\theta_{S_j^m}-\psi_{S_j^m})(s_{f+1}-s_f)$$

with the reasoning in the proof of Lemma 3.8.2, it is easily seen that

$$\sum_{k=1}^{n}\sum_{j=1}^{k}\sum_{m=k}^{p-1}(\gamma_{S_j^m}-\psi_{S_j^m})(t-s_{k+1})=\sum_{k=1}^{n}\sum_{f=1}^{k-1}\sum_{j=1}^{f}\sum_{m=f}^{p-1}(\gamma_{S_j^m}-\psi_{S_j^m})(s_{k+1}-s_k).$$

Using this in the upper bound yields,

$$\begin{aligned}&\sup_{0\le s_1\le\ldots\le s_{n+1}=t}\left\{A_i(s_1,t)-\sum_{k=1}^{n}d_k(s_{k+1}-s_k)\right\}\\&+\sum_{k=1}^{n}\sum_{m=p}^{N}\sup_{0\le s_k\le t}\{A_{S_k^m}(s_k,t)-\gamma_{S_k^m}(t-s_k)\}\\&+\sum_{k=1}^{n}\sum_{j=1}^{k}\sum_{m=k}^{p-1}\sum_{q\in S_j^m}\sup_{0\le s_m\le t}\{A_q(s_m,t)-\gamma_q(t-s_m)\}\\&+\sum_{k=1}^{n}\sum_{j=1}^{k}\sum_{m=k}^{p-1}\sum_{q\in S_j^m}\sup_{0\le s_{k+1}\le t}\{\psi_q(t-s_{k+1})-A_q(s_{k+1},t)\}.\end{aligned}$$

Then using (3.8) and (3.46) for the first supremum, the proof is completed. □

Chapter 4

GPS with heterogeneous traffic

In the previous chapter we showed that GPS is capable of protecting a flow with heavy-tailed characteristics against other flows. The results were proved for the case where the heavy-tailed flow travels through a feed-forward network, coming across flows with possibly heavier-tailed characteristics. In this chapter we initiate the investigation of the workload behavior of *light-tailed* traffic when served together with heavy-tailed traffic in a single node according to GPS. Note that this chapter presents the results in [28]. Specific parts of this work have appeared in [26, 27, 29].

Previous work on light-tailed traffic when mixed with heavy-tailed traffic concentrated mainly on the FIFO discipline (see [31, 55, 152]). An exception is [21], where the asymptotic behavior of the workload distribution of an individual flow is derived under the assumption that its GPS weight is smaller than its average rate. The result holds for all types of input for the particular flow, for instance if the flow has light-tailed characteristics. If the other flows have heavier-tailed characteristics, then the workload behavior of the light-tailed flow will inherit these properties. Unfortunately, the result does not indicate to what extent light-tailed flows are affected by heavy-tailed flows in the more plausible situation when their GPS weight is larger than their traffic intensity.

In the present chapter we consider a GPS queue fed by a mixture of light-tailed and heavy-tailed traffic. More specifically, we assume the light-tailed flow to be Markov-modulated fluid. That is, the flow can be in a number of states, in each of which it generates traffic at a certain rate for an exponentially distributed period of time. For the heavy-tailed flow we consider two cases: (i) it generates instantaneous traffic bursts whose sizes have a heavy-tailed distribution, or (ii) it behaves according to an on-off process with heavy-tailed on periods. In principle, it would be possible to extend the analysis to general heavy-tailed fluid models or combinations of instantaneous and fluid heavy-tailed input. However, we will not pursue these extensions in the present chapter since they would blur the results without providing substantial additional insight. We derive the asymptotic workload behavior of the light-tailed flow under the assumption that its GPS weight is larger than its traffic intensity. We examine how the performance experienced by the light-tailed flow is affected by possibly badly behaving heavy-tailed input.

In the analysis, we reduce the space of all possible sample paths to overflow to a single 'most-likely' path which occurs with overwhelming probability. This procedure yields valuable insight in the typical overflow scenario. The most-likely path can be found using the following intuition regarding rare events for light-tailed and heavy-tailed traffic, cf. [31]. From large deviations we know that in the *light-tailed* case the flow behaves such that it remains longer than usual in states with a high traffic rate, and shorter than usual in states with a low traffic rate, thus sending at a higher intensity. Asymptotic results for *heavy-tailed* traffic suggest that the workload caused by the heavy-tailed flow will grow large due to a single large on period or a single large burst. We also mention [33], where a similar combination of light-tailed large deviations and heavy-tailed large deviations is used.

The remainder of this chapter is organized as follows. In Section 4.1, we present a description of the model and state some important preliminary results. In Section 4.2, we provide an overview of the main results, which characterize the exact asymptotic behavior of the workload distribution of the light-tailed traffic flow. The subsequent sections are devoted to the detailed proofs. We start in Section 4.3 with deriving lower and upper bounds for the workload distribution of the light-tailed flow. In Section 4.4, we proceed to prove some auxiliary results for the light-tailed flow in isolation. Although the bounds seem quite crude by themselves, we show in Section 4.5 that they asymptotically coincide, yielding the exact asymptotic behavior. One of the asymptotic terms involves the probability that the heavy-tailed flow is backlogged long enough for overflow to occur. In order to determine the distribution of the backlog period, we first establish in Section 4.6 some preliminary results for the heavy-tailed flow in isolation. We then compute the specific form of the distribution for various traffic processes in Sections 4.7, 4.8 and 4.9.

4.1 Model description

We consider two traffic flows sharing a link of unit rate, which is equipped with the GPS mechanism. We follow the notation as introduced in Section 1.3, unless it deviates from previously introduced notation. For $t > 0$, we denote by $a_i(t)$ the amount of traffic generated by flow i during the time interval $(0, t]$, $i = 1, 2$. For $t \leq 0$, $a_i(t)$ denotes the negative counterpart of the amount of traffic generated by flow i in $(t, 0]$. Assuming that $a_i(\cdot)$ is reversible and has stationary increments, we define $A_i(s,t) := a_i(t) - a_i(s)$ to be the amount of traffic generated in $(s, t]$, $s < t$, $s, t \in \mathbb{R}$. Define $B_i(s,t)$ as the amount of service received by flow i during $(s, t]$. As mentioned in the introduction, a relation similar to (1.2) holds for $V_i^c(t) := \sup_{s \leq t}\{A_i(s,t) - c(t-s)\}$, the workload at time t in a queue where flow i is served in isolation at rate c. With $B_i^c(s,t)$ denoting the amount of service received by flow i in this queue in $(s, t]$, we have for all $s < t$,

$$V_i^c(t) = V_i^c(s) + A_i(s,t) - B_i^c(s,t). \tag{4.1}$$

For any two real functions $f(\cdot)$ and $g(\cdot)$, we use the following notation:

$$f(x) \lesssim g(x) \quad \text{denotes} \quad \limsup_{x\to\infty} \frac{f(x)}{g(x)} \leq 1; \tag{4.2}$$

$$f(x) \gtrsim g(x) \quad \text{denotes} \quad \liminf_{x\to\infty} \frac{f(x)}{g(x)} \geq 1. \tag{4.3}$$

4.1.1 Traffic model flow 1

We assume that flow 1 is light-tailed. Specifically, we make the assumption that the input process $A_1(s,t)$ is a *Markov-modulated fluid.* Such a process can be described as follows. There is an irreducible Markov chain with a finite state space $\{1, 2, \ldots, d\}$. The corresponding transition rate matrix is denoted by $\Lambda := (\lambda_{ij})_{i,j=1,\ldots,d}$, where we follow the convention that $\lambda_{ii} := -\sum_{j\neq i} \lambda_{ij}$. Since the Markov chain is irreducible, there is a unique stationary distribution, which we denote by the (column) vector π. When the source is in state i, traffic is generated (as fluid) at constant rate $R_i < \infty$. Let R be the diagonal matrix with the coefficients R_i on the diagonal. Denote the mean rate by $\rho_1 := \sum_{i=1}^{d} \pi_i R_i$. Denote the peak rate by $R_P := \max_{i=1,\ldots,d} R_i$. We assume that $R_P > \phi_1$ in order to avoid the trivial case where the workload of flow 1 is always zero. It is important to observe that the class of Markov fluid input is closed under superposition, i.e., the superposition of Markov fluid sources can again be modeled as a Markov fluid source. As observed in the Introduction, Markov-modulated fluid models have gained strong ground as a flexible modeling tool for describing the statistical behavior of light-tailed traffic sources. Although we adopt the Markovian structure to ensure mathematical tractability, we expect that qualitatively similar results hold for general light-tailed input, see Remark 4.5.1 below.

The following standard result for Markov-modulated fluid sources follows directly from [59], [72], and [76].

Proposition 4.1.1 *Let $c \in (\rho_1, R_P)$.*

- *The moment generating function of traffic generated in an interval of length t in steady state is given by, in matrix notation,*

  $$\mathbb{E}e^{\theta A_1(0,t)} = \pi^T e^{(\Lambda+\theta R)t}\mathbf{1}, \qquad \theta \in \mathbb{R},$$

  with $\mathbf{1}$ the all one (column) vector of dimension d.

- *There exists a limiting moment generating function:*

  $$\lim_{t\to\infty} \frac{1}{t} \log \mathbb{E}e^{\theta(A_1(0,t)-ct)} = M_c(\theta).$$

  This function is continuous and differentiable. It also holds that there is a finite, positive C such that

  $$\mathbb{E}e^{\theta(A_1(0,t)-ct)} \leq Ce^{M_c(\theta)t}.$$

- *The large-buffer asymptotics of a queue with Markov-modulated fluid input are given by*

$$\liminf_{x\to\infty} \frac{1}{x} \log \mathbb{P}\left(V_1^c > x\right) = -\theta^\star(c),$$

with $\theta^\star(c)$ the unique positive root of $M_c(\theta) = 0$. Moreover, $M_c'(\theta^\star(c)) > 0$.

In the remainder of this chapter we will mainly use $c = \phi_1$. For ease of notation we therefore write $M(\cdot)$ instead of $M_{\phi_1}(\cdot)$, and $\theta^\star$ instead of $\theta^\star(\phi_1)$, when it is clear that we take $c = \phi_1$.

4.1.2 Traffic model flow 2

We assume that flow 2 is heavy-tailed. That is, the input process $A_2(s,t)$ is either instantaneous or on-off, with heavy-tailed burst sizes or on periods, respectively. The traffic models are similar to that in the previous chapter, only the choice of the heavy-tailed distribution differs. Because the notation is different from that introduced in Sections 3.1.1 and 3.1.2, we introduce both the notation and the two preliminary results again.

Instantaneous input

Here, flow 2 generates instantaneous traffic bursts according to a renewal process. The interarrival times between bursts have distribution function $U_2(\cdot)$ with mean $1/\lambda_2$. The burst sizes B_2 have distribution function $B_2(\cdot)$ with mean $\beta_2 < \infty$. Thus, the traffic intensity is $\rho_2 := \lambda_2\beta_2$. We assume that $B_2(\cdot) \in \mathcal{R}_{-\nu_2}$, i.e., $\bar{B}_2(x) := 1 - B_2(x)$ is regularly varying of index $-\nu_2$, for some $\nu_2 > 1$. The next result which is due to Pakes [106] (see also Theorem 3.1.1) then yields the tail behavior of the workload distribution of flow 2 in isolation.

Theorem 4.1.1 *If $B_2^{\mathrm{r}}(\cdot) \in \mathcal{S}$, and $\rho_2 < c$, then*

$$\mathbb{P}\left(V_2^c > x\right) \sim \frac{\rho_2}{c - \rho_2} \mathbb{P}\left(B_2^{\mathrm{r}} > x\right).$$

Recall that for a random variable X with distribution function $F(\cdot)$, we use the notation X^{r} and $F^{\mathrm{r}}(\cdot)$ for its residual lifetime.

Fluid input

Here, flow 2 generates traffic according to an on-off process, alternating between on and off periods. The off periods U_2 have distribution function $U_2(\cdot)$ with mean $1/\lambda_2$. The on periods A_2 have distribution function $A_2(\cdot)$ with mean $\alpha_2 < \infty$. While on, flow i produces traffic at constant rate r_2, so the mean burst size, i.e., the mean amount of traffic sent during an on period, is $\alpha_2 r_2$. The fraction of time that flow 2 is off is

$$p_2 := \frac{1/\lambda_2}{1/\lambda_2 + \alpha_2} = \frac{1}{1 + \lambda_2\alpha_2}.$$

The traffic intensity is

$$\rho_2 := (1 - p_2) r_2 = \frac{\lambda_2 \alpha_2 r_2}{1 + \lambda_2 \alpha_2}.$$

We assume that $A_2(\cdot) \in \mathcal{R}_{-\nu_2}$, i.e., $\bar{A}_2(x) := 1 - A_2(x)$ is regularly varying of index $-\nu_2$, for some $\nu_2 > 1$. The next result as presented in [67] (see also Theorem 3.1.4) then yields the tail behavior of the workload distribution of flow 2 in isolation.

Theorem 4.1.2 *If $A_2^{\mathrm{r}}(\cdot) \in \mathcal{S}$, and $\rho_2 < c < r_2$, then*

$$\mathbb{P}\left(V_2^c > x\right) \sim p_2 \frac{\rho_2}{c - \rho_2} \mathbb{P}\left(A_2^{\mathrm{r}} > \frac{x}{r_2 - c}\right).$$

4.2 Overview of the results

Throughout this chapter, we assume $\rho_i < \phi_i$, $i = 1, 2$, which ensures stability of both flows. We first briefly discuss in Section 4.2.1 what happens if this condition fails to hold. In addition, we make the assumption that $r_2 > \phi_2$ in case of fluid input of flow 2. Otherwise, the workload of flow 2 would be zero, so the workload of flow 1 would be equal to the total workload. The asymptotic distribution of the total workload has been obtained in [31]. In Section 4.2.2 we provide a heuristic explanation of the main results of this chapter. The main result (without proof) is then given in Section 4.2.3, where we also present an example.

4.2.1 The case $\rho_1 > \phi_1$

To put things in perspective, we first briefly review the case that $\rho_1 > \phi_1$, while $\rho_1 + \rho_2 < 1$. If either (i) $B_2^{\mathrm{r}}(\cdot) \in \mathcal{IR}$ (instantaneous input of flow 2), or (ii) $A_2^{\mathrm{r}}(\cdot) \in \mathcal{IR}$, $r_2 > \phi_2$ (fluid input), then from [17],

$$\mathbb{P}\left(V_1 > x\right) \sim \frac{\phi_2 - \rho_2}{\phi_2} \frac{\rho_2}{1 - \rho_1 - \rho_2} \mathbb{P}\left(P_2^{\mathrm{r}} > \frac{x}{\rho_1 - \phi_1}\right),$$

with P_2 a random variable whose distribution is the busy-period distribution in a queue with constant service rate ϕ_2 fed by flow 2.

The above result may be interpreted as follows. Large-deviations arguments suggest that the most-likely way for flow 1 to build a large queue is that flow 2 generates a large burst, or experiences a long on period, while flow 1 itself shows roughly average behavior. Note that when flow 2 produces a large amount of traffic, so it becomes backlogged for a long period of time, it receives service at rate ϕ_2. Thus it will experience a busy period as if it were served at constant rate ϕ_2. During that period, flow 1 receives service at rate ϕ_1, while it generates traffic roughly at rate ρ_1, so its queue will grow approximately at rate $\rho_1 - \phi_1$. When flow 2 is not backlogged, the corresponding queue will drain approximately at rate $1 - \rho_1$.

Thus, the backlog of flow 1 behaves as that in a queue with constant service rate $1-\rho_1$ fed by an on-off source with peak rate ϕ_2. The on and the off periods correspond to the busy and idle periods of flow 2 when served at constant rate ϕ_2, respectively. Recall that the workload asymptotics of such an on-off source are given by Theorem 4.1.2. Setting $c = 1 - \rho_1$, $r_2 = \phi_2$, $p_2 = 1 - \rho_2/\phi_2$, and identifying A_2^{r} with P_2^{r}, we obtain the above result for the workload asymptotics of flow 1.

4.2.2 Heuristic explanation of the main results

We now focus on the case $\rho_1 < \phi_1$. Before presenting the main result, we first provide a heuristic derivation of the asymptotic behavior of $\mathbb{P}(V_1 > x)$ based on large-deviations arguments. The overflow scenario described above for the case $\rho_1 > \phi_1$ cannot occur, and now flow 1 too must deviate from its 'normal' behavior in order for the queue to grow. Specifically, large-deviations results suggest that flow 1 must behave as if its traffic intensity is temporarily increased from ρ_1 to some larger value $\hat{\rho}_1 > \phi_1$ (as will be specified below). During that time period, flow 2 is continuously backlogged, consuming service rate ϕ_2, thus leaving service rate ϕ_1 for flow 1. (Notice that if flow 2 were not permanently backlogged, then flow 1 would have to show even greater anomalous activity in order for a given backlog level to occur.) The intuitive argument may be summarized as follows (see Figure 4.1). A large backlog of level x of flow 1 occurs as a consequence of two rare events: *(i)* Flow 1 shows 'abnormal' behavior similar to what typically causes overflow when served in isolation, thus behaving as if its traffic intensity is increased from ρ_1 to $\hat{\rho}_1$ for at least a period of time $x/(\hat{\rho}_1 - \phi_1)$. *(ii)* During that time period, flow 2 is constantly backlogged, demanding service rate ϕ_2, with service rate ϕ_1 remaining for flow 1. As we will see later, the persistent backlog is most likely caused by flow 2 generating a large burst or initiating a long on period prior to that period.

These considerations lead to the following characterization of the asymptotic

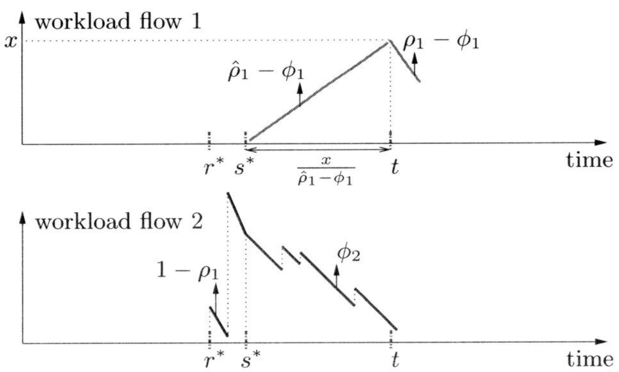

Figure 4.1: Overflow scenario - instantaneous input flow 2.

behavior of $\mathbb{P}(V_1 > x)$:

$$\mathbb{P}(V_1 > x) \sim \mathbb{P}\left(V_1^{\phi_1} > x\right) \mathbb{P}\left(T_2^{1-\rho_1} > \frac{x}{\hat{\rho}_1 - \phi_1}\right). \tag{4.4}$$

The second term represents the probability that flow 2 is continuously backlogged during a period of time $x/(\hat{\rho}_1 - \phi_1)$, receiving a service rate ϕ_2 starting from some time t on, and having received a service rate $1 - \rho_1$ prior to time t. Here $T_2^{1-\rho_1}$ is a random variable whose distribution is the limit distribution of $T_2^{1-\rho_1,t}$ for $t \to \infty$, with

$$T_2^{c,t} := \inf\{u \geq 0 : V_2^c(t) + A_2(t, t+u) - \phi_2 u = 0\}$$

representing the drain time in a queue with service rate ϕ_2 fed by flow 2 with initial workload $V_2^c(t)$. The service rate $1 - \rho_1$ reflects the fact that flow 1 has shown normal behavior prior to time t.

Thus, the workload distribution is asymptotically equivalent to that in an isolated system, multiplied with a certain post-factor. The isolated system consists of flow 1 served in isolation at constant rate ϕ_1. The post-factor represents the probability that flow 2 is backlogged long enough for flow 1 to reach overflow.

Note that the general product form of (4.4) holds irrespective of the detailed traffic characteristics of the two flows. However, the specific form of the two individual terms in (4.4) *does* depend on the detailed properties of the traffic processes. In particular, we need to distinguish whether flow 2 generates instantaneous or fluid input. In the latter case, it also depends on whether the peak rate r_2 exceeds $1 - \rho_1$ or not.

4.2.3 Main results

We now state the main theorem of this chapter.

Theorem 4.2.1 *Defining* $\hat{\rho}_1 := M'_{\phi_1}(\theta^\star(\phi_1)) + \phi_1$,

$$\mathbb{P}(V_1 > x) \sim \mathbb{P}\left(V_1^{\phi_1} > x\right) \mathbb{P}\left(T_2^{1-\rho_1} > \frac{x}{\hat{\rho}_1 - \phi_1}\right).$$

Case I: (instantaneous input)

$$\mathbb{P}\left(T_2^{1-\rho_1} > x\right) \sim \frac{\rho_2}{1 - \rho_1 - \rho_2} \mathbb{P}\left(B_2^{\mathrm{r}} > x(\phi_2 - \rho_2)\right). \tag{4.5}$$

Case II-A: (fluid input with $r_2 < 1 - \rho_1$)

$$\mathbb{P}\left(T_2^{1-\rho_1} > x\right) \sim (1 - p_2)\mathbb{P}\left(A_2^{\mathrm{r}} > \frac{x(\phi_2 - \rho_2)}{r_2 - \rho_2}\right). \tag{4.6}$$

Case II-B: (fluid input with $r_2 > 1 - \rho_1$)

$$\mathbb{P}\left(T_2^{1-\rho_1} > x\right) \sim p_2 \frac{\rho_2}{1 - \rho_1 - \rho_2} \mathbb{P}\left(A_2^{\mathrm{r}} > \frac{x(\phi_2 - \rho_2)}{r_2 - \rho_2}\right). \tag{4.7}$$

Noting that $p_2\rho_2 = (1-p_2)(r_2-\rho_2)$, it is easily seen that in the limiting regime $r_2 \to 1-\rho_1$, cases II-A and II-B coincide. Also, case I can be seen as the limiting case of II-B if we use $r_2 A_2 = B_2$ and let $r_2 \to \infty$ so that $p_2 \downarrow 1$. In Chapter 5 a qualitatively similar result as in case I is derived for a system with two coupled queues, one having heavy-tailed input, the other one exhibiting light-tailed properties.

Before proceeding to the formal proof of Theorem 4.2.1, we first give an example. Assume flow 1 to behave according to an on-off process with exponentially distributed on and off periods with means $1/\mu_1$ and $1/\mu_2$, respectively. When the flow is in the on-state, it generates traffic at rate R_1. We assume flow 2 to generate instantaneous input with regularly varying burst sizes of index $-\nu_2$, i.e.,

$$\mathbb{P}(B_2 > x) \sim C_2 x^{-\nu_2} l_2(x),$$

with $l_2(\cdot)$ some slowly varying function. First we determine the deviant traffic intensity $\hat{\rho}_1$ using [88],

$$\hat{\rho}_1 = \left(\frac{R_1\phi_1^2}{\mu_2}\right) \Big/ \left(\frac{\phi_1^2}{\mu_2} + \frac{(R_1-\phi_1)^2}{\mu_1}\right).$$

Then using [58], we obtain for flow 1,

$$\mathbb{P}\left(V_1^{\phi_1} > x\right) \sim \frac{R_1}{\phi_1}\frac{\mu_2}{\mu_1+\mu_2}\exp\left\{-\left(\frac{\mu_1}{R_1-\phi_1}+\frac{\mu_2}{\phi_1}\right)x\right\}.$$

For flow 2, from (4.5),

$$\mathbb{P}\left(T_2^{1-\rho_1} > x\right) \sim \frac{\rho_2}{1-\rho_1-\rho_2}\frac{C_2}{\beta_2(\nu_2-1)}(x(\phi_2-\rho_2))^{1-\nu_2} l_2(x(\phi_2-\rho_2)).$$

This provides all the ingredients for $\mathbb{P}(V_1 > x)$ as required in Theorem 4.2.1.

4.3 Bounds

In this section we derive lower and upper bounds for the workload distribution of flow 1. The bounds will be instrumental in obtaining the asymptotic behavior of $\mathbb{P}(V_1 > x)$ as given in Theorem 4.2.1.

4.3.1 Lower bound

We start with a lower bound for the workload distribution of flow 1. The main idea (see Figure 4.2) is that the following scenario is sufficient for the event $V_1(t) > x$ to occur (in fact, is the only plausible one, as we will see later). Flow 1 starts to build up at some time $s^\star$, and hence is constantly backlogged throughout the time interval $[s^\star, t]$. Flow 2 is also continuously backlogged during $[s^\star, t]$. Thus, during that time period, flows 1 and 2 both receive service at rates ϕ_1 and ϕ_2, respectively. Flow 2 already becomes backlogged at time $r^\star \leq s^\star$, and receives service approximately at rate $1-\rho_1$ during $[r^\star, s^\star]$, while flow 1 then shows roughly average behavior.

Lemma 4.3.1 *For any $\epsilon > 0$, suppose $r^\star, s^\star$ and y exist such that $r^\star \le s^\star \le t$ and*

$$A_1(s^\star, t) - \phi_1(t - s^\star) > x,$$
$$A_1(r^\star, s^\star) - (\rho_1 - \epsilon)(s^\star - r^\star) \ge -y,$$

and

$$\inf_{s^\star \le u \le t} \{A_2(r^\star, u) - (1 - \rho_1 + \epsilon)(s^\star - r^\star) - \phi_2(u - s^\star)\} \ge y,$$

then $V_1(t) > x$.

Proof From (1.2), we have for all $s \le t$,

$$V_1(t) = V_1(s) + A_1(s,t) - B_1(s,t), \tag{4.8}$$

and by definition,

$$B_1(s,t) + B_2(s,t) \le t - s. \tag{4.9}$$

Because of the GPS discipline, it is easily seen that

$$B_2(s,t) \ge \min\left\{\phi_2(t-s), V_2(s) + \inf_{s \le u \le t}\{A_2(s,u) + \phi_2(t-u)\}\right\}. \tag{4.10}$$

Substituting (4.9) and (4.10) into (4.8) yields

$$\begin{aligned} V_1(t) \quad \ge \quad & A_1(s,t) - \phi_1(t-s) \\ & + \min\left\{0, V_1(s) + V_2(s) + \inf_{s \le u \le t}\{A_2(s,u) - \phi_2(u-s)\}\right\} \end{aligned} \tag{4.11}$$

From (1.2), for all $r \le s$,

$$V_1(s) + V_2(s) = V_1(r) + V_2(r) + A_1(r,s) + A_2(r,s) - B_1(r,s) - B_2(r,s).$$

Together with $B_1(r,s) + B_2(r,s) \le s - r$, this yields

$$V_1(s) + V_2(s) \ge A_1(r,s) + A_2(r,s) - (s - r).$$

Substituting the previous inequality into (4.11), it holds for all $r \le s \le t$ that

$$V_1(t) \quad \ge \quad A_1(s,t) - \phi_1(t-s) + \min\left\{0, A_1(r,s) + A_2(r,s) - (s-r) + \right.$$

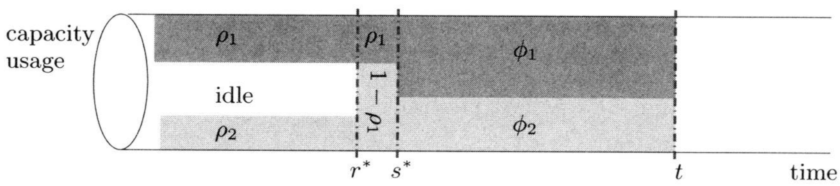

Figure 4.2: Intuitive idea lower bound.

$$\inf_{s \leq u \leq t} \{A_2(s,u) - \phi_2(u-s)\}\Big\}$$

$$= \quad A_1(s,t) - \phi_1(t-s) + \min\Big\{0, A_1(r,s) - (\rho_1 - \epsilon)(s-r) +$$

$$\inf_{s \leq u \leq t} \{A_2(r,u) - (1-\rho_1+\epsilon)(s-r) - \phi_2(u-s)\}\Big\},$$

which completes the proof. □

We now translate the above sample-path result into a probabilistic lower bound. We first introduce some additional notation. For any c and $w \geq 0$, define

$$V_i^c[w] := \sup_{0 \leq s \leq w} \{A_i(-s,0) - cs\}.$$

Note that, for $c > \rho_i$, $V_i^c[\infty] \stackrel{d}{=} V_i^c$ as defined earlier (recall that we use $\stackrel{d}{=}$ to denote equality in distribution). For any c, $v \geq 0$, and y, define

$$T_2^c(v,y) := \inf\Big\{u \geq 0 : \sup_{0 \leq r \leq v} \{A_2(-r,0) - cr\} + A_2(0,u) - \phi_2 u \leq y\Big\}.$$

Thus, $T_2^c(v,y)$ represents the drain time in a queue of service rate ϕ_2 fed by flow 2 with initial workload $\sup_{0 \leq r \leq v}\{A_2(-r,0) - cr\} - y$. Note that, for $c > \rho_2$,

$$T_2^c(y) := T_2^c(\infty, y) = \inf\{u \geq 0 : V_2^c(0) + A_2(0,u) - \phi_2 u \leq y\},$$

and that $T_2^c(0) \stackrel{d}{=} T_2^c$ as defined earlier. Also, define

$$T_2(y) := T_2^c(0,y) = \inf\{u \geq 0 : A_2(0,u) - \phi_2 u \leq y\}.$$

(note that the latter quantity does not depend on the value of c), and denote $T_2 := T_2(0)$. We define $P^{\rho_1-\epsilon}(s^\star, v, x, y) :=$

$$\mathbb{P}\left(\sup_{s^\star - v \leq r \leq s^\star} \{(\rho_1 - \epsilon)(s^\star - r) - A_1(r, s^\star)\} \leq y \mid A_1(s^\star, 0) + \phi_1 s^\star > x\right).$$

Corollary 4.3.1 *For any $v \geq 0$ and y, a lower bound for $\mathbb{P}(V_1 > x)$ is given by*

$$\mathbb{P}\left(V_1^{\phi_1}\left[\frac{(1+\alpha)x}{\hat{\rho}_1 - \phi_1}\right] > x\right)\mathbb{P}\left(T_2^{1-\rho_1+\epsilon}(v,y) > \frac{(1+\alpha)x}{\hat{\rho}_1 - \phi_1}\right)P^{\rho_1-\epsilon}(s^\star, v, x, y).$$

Proof Using Lemma 4.3.1, the independence of $A_1(s,t)$ and $A_2(s,t)$, and the fact that $A_1(s,t)$ and $A_2(s,t)$ have stationary increments, for all $v, w \geq 0$, and y,

$$\mathbb{P}(V_1(t) > x)$$

$$\geq \quad \mathbb{P}\Big(\exists s^\star \in [t-w,t], \exists r^\star \in [s^\star - v, s^\star] : A_1(s^\star, t) - \phi_1(t - s^\star) > x,$$

$$
\begin{aligned}
& A_1(r^\star, s^\star) - (\rho_1 - \epsilon)(s^\star - r^\star) \geq -y, \\
& \inf_{s^\star \leq u \leq t} \{A_2(r^\star, u) - (1 - \rho_1 + \epsilon)(s^\star - r^\star) - \phi_2(u - s^\star)\} \geq y \Big) \\
= \; & \mathbb{P}\Big(\exists s^\star \in [-w, 0], \exists r^\star \in [s^\star - v, s^\star] : A_1(s^\star, 0) + \phi_1 s^\star > x, \\
& A_1(r^\star, s^\star) - (\rho_1 - \epsilon)(s^\star - r^\star) \geq -y, \\
& \inf_{s^\star \leq u \leq 0} \{A_2(r^\star, u) - (1 - \rho_1 + \epsilon)(s^\star - r^\star) - \phi_2(u - s^\star)\} \geq y \Big) \\
\geq \; & \mathbb{P}\Big(\exists s^\star \in [-w, 0], \exists r^\star \in [s^\star - v, s^\star] : A_1(s^\star, 0) + \phi_1 s^\star > x, \\
& \inf_{s^\star - v \leq r \leq s^\star} \{A_1(r, s^\star) - (\rho_1 - \epsilon)(s^\star - r)\} \geq -y, \\
& \inf_{s^\star \leq u \leq s^\star + w} \{A_2(r^\star, u) - (1 - \rho_1 + \epsilon)(s^\star - r^\star) - \phi_2(u - s^\star)\} \geq y \Big) \\
= \; & \mathbb{P}\Big(\exists s^\star \in [-w, 0] : A_1(s^\star, 0) + \phi_1 s^\star > x, \\
& \inf_{s^\star - v \leq r \leq s^\star} \{A_1(r, s^\star) - (\rho_1 - \epsilon)(s^\star - r)\} \geq -y, \\
& \exists r^\star \in [s^\star - v, s^\star] : \\
& \inf_{s^\star \leq u \leq s^\star + w} \{A_2(r^\star, u) - (1 - \rho_1 + \epsilon)(s^\star - r^\star) - \phi_2(u - s^\star)\} \geq y \Big) \\
\geq \; & \mathbb{P}\left(\exists s^\star \in [-w, 0] : A_1(s^\star, 0) + \phi_1 s^\star > x \right) \\
& \times \mathbb{P}\Big(\inf_{s^\star - v \leq r \leq s^\star} A_1(r, s^\star) - (\rho_1 - \epsilon)(s^\star - r) \geq -y \mid A_1(s^\star, 0) + \phi_1 s^\star > x \Big) \\
& \times \mathbb{P}\Big(\exists r^\star \in [s^\star - v, s^\star] : \\
& \inf_{s^\star \leq u \leq s^\star + w} \{A_2(r^\star, u) - (1 - \rho_1 + \epsilon)(s^\star - r^\star) - \phi_2(u - s^\star)\} \geq y \Big) \\
= \; & \mathbb{P}\left(\exists s \in [0, w] : A_1(-s, 0) - \phi_1 s > x \right) \\
& \times \mathbb{P}\Big(\sup_{s^\star - v \leq r \leq s^\star} \{(\rho_1 - \epsilon)(s^\star - r) - A_1(r, s^\star)\} \leq y \mid A_1(s^\star, 0) + \phi_1 s^\star > x \Big) \\
& \times \mathbb{P}\Big(\exists r \in [0, v] : \inf_{0 \leq u \leq w} \{A_2(-r, u) - (1 - \rho_1 + \epsilon) r - \phi_2 u\} \geq y \Big) \\
= \; & \mathbb{P}\Big(\sup_{0 \leq s \leq w} \{A_1(-s, 0) - \phi_1 s\} > x \Big) P^{\rho_1 - \epsilon}(s^\star, v, x, y) \\
& \times \mathbb{P}\Big(\sup_{0 \leq r \leq v} \inf_{0 \leq u \leq w} \{A_2(-r, u) - (1 - \rho_1 + \epsilon) r - \phi_2 u\} \geq y \Big) \\
= \; & \mathbb{P}\left(V_1^{\phi_1}[w] > x \right) P^{\rho_1 - \epsilon}(s^\star, v, x, y)
\end{aligned}
$$

$$\times \mathbb{P}\left(\inf_{0\le u\le w}\left\{\sup_{0\le r\le v}\{A_2(-r,0)-(1-\rho_1+\epsilon)r\}+A_2(0,u)-\phi_2 u\right\}\ge y\right)$$
$$= \mathbb{P}\left(V_1^{\phi_1}[w]>x\right)\mathbb{P}\left(T_2^{1-\rho_1+\epsilon}(v,y)>w\right)P^{\rho_1-\epsilon}(s^\star,v,x,y).$$

Taking $w=((1+\alpha)x)/(\hat\rho_1-\phi_1)$ completes the proof. □

4.3.2 Upper bound

We proceed to derive an upper bound for the workload distribution of flow 1. The idea is that the lower-bound scenario described above is basically also necessary for the event $V_1(t)>x$ to occur.

Lemma 4.3.2 *Suppose that $V_1(t)>x$. Then for any $\epsilon>0$ and for all y there exist $r^\star, s^\star$ such that $r^\star\le s^\star\le t$ and*

$$A_1(s^\star,t)-\phi_1(t-s^\star)>x, \tag{4.12}$$

and at least one of the three following events occurs, either

$$A_1(r^\star,s^\star)-(\rho_1+\epsilon)(s^\star-r^\star)>y, \tag{4.13}$$

or

$$V_1^{\phi_1}(t)>x+y, \tag{4.14}$$

or

$$\inf_{s^\star\le u\le t}\{A_2(r^\star,u)-(1-\rho_1-\epsilon)(s^\star-r^\star)-\phi_2(u-s^\star)\}>-2y. \tag{4.15}$$

Proof First we show that (4.12) is implied by $V_1(t)>x$. Because of the GPS discipline,

$$V_1(t)\le V_1^{\phi_1}(t)=\sup_{s\le t}\{A_1(s,t)-\phi_1(t-s)\}.$$

Hence, there exists an $s\le t$ such that $A_1(s,t)-\phi_1(t-s)>x$. Defining

$$s^\star:=\inf\{s:A_1(u,t)-\phi_1(t-u)\le x\,\forall u>s\}=\sup\{s:A_1(s,t)-\phi_1(t-s)>x\},$$

we first show by contradiction, that $V_1(t)>x$ implies that flow 1 is continuously backlogged during the interval $[s^\star,t]$. Suppose flow 1 is *not* continuously backlogged, then $u\in[s^\star,t]$ exists such that $V_1(u)=0$. Defining $u^\star:=\sup\{u\in[s^\star,t]:V_1(u)=0\}$, it holds that

$$V_1(t)=A_1(u^\star,t)-B_1(u^\star,t),$$

which combined with $B_1(u^\star,t)\ge\phi_1(t-u^\star)$ yields

$$V_1(t)\le A_1(u^\star,t)-\phi_1(t-u^\star).$$

In view of $V_1(t)>x$, we have $A_1(u^\star,t)-\phi_1(t-u^\star)>x$, which contradicts the definition of $s^\star$. Hence, flow 1 must be continuously backlogged during $[s^\star,t^\star]$.

We now show in two steps, that $V_1(t) > x$ implies that at least one of the events (4.13), (4.14) or (4.15) must occur: (i) first we prove that $V_1(t) > x$ implies either (4.14) or

$$\forall u \in [s^\star, t] : B_2(s^\star, u) - \phi_2(u - s^\star) > -y,$$

(ii) next, we prove that the latter event implies either (4.13) or (4.15).

Ad (i). Equivalently, we prove that the events

$$\exists u^\star \in [s^\star, t] : B_2(s^\star, u^\star) - \phi_2(u^\star - s^\star) \leq -y \tag{4.16}$$

and

$$\forall q \leq s^\star \leq t : A_1(q, t) - \phi_1(t - q) \leq x + y \tag{4.17}$$

imply $V_1(t) \leq x$, where (4.17) is the complement of (4.14). Since flow 1 is continuously backlogged during $[s^\star, t]$,

$$V_1(t) = V_1(s^\star) + A_1(s^\star, t) - (t - s^\star) + B_2(s^\star, u^\star) + B_2(u^\star, t)$$

and

$$B_2(u^\star, t) \leq \phi_2(t - u^\star).$$

Because of the GPS discipline,

$$V_1(s^\star) \leq \sup_{r \leq s^\star} \{A_1(r, s^\star) - \phi_1(s^\star - r)\}.$$

Hence, using first (4.16) and then (4.17),

$$\begin{aligned} V_1(t) &\leq \sup_{r \leq s^\star} \{A_1(r, s^\star) - \phi_1(s^\star - r)\} + A_1(s^\star, t) - (t - s^\star) + \phi_2(t - s^\star) - y \\ &= \sup_{r \leq s^\star} \{A_1(r, t) - \phi_1(t - r)\} - y \leq x + y - y = x. \end{aligned}$$

Ad (ii). By definition, using (1.4) for the total workload,

$$\begin{aligned} B_2(s^\star, u) &\leq V_2(s^\star) + A_2(s^\star, u) \leq V_1(s^\star) + V_2(s^\star) + A_2(s^\star, u) \\ &= \sup_{r \leq s^\star} \{A_1(r, s^\star) + A_2(r, s^\star) - (s^\star - r)\} + A_2(s^\star, u). \end{aligned}$$

Hence,

$$\begin{aligned} &\inf_{s^\star \leq u \leq t} \{B_2(s^\star, u) - \phi_2(u - s^\star)\} \\ \leq\ &\inf_{s^\star \leq u \leq t} \left\{ \sup_{r \leq s^\star} \{A_1(r, s^\star) + A_2(r, s^\star) - (s^\star - r)\} + A_2(s^\star, u) - \phi_2(u - s^\star) \right\} \\ =\ &\sup_{r \leq s^\star} \inf_{s^\star \leq u \leq t} \{A_1(r, s^\star) + A_2(r, u) - (s^\star - r) - \phi_2(u - s^\star)\} \\ \leq\ &\sup_{r \leq s^\star} \inf_{s^\star \leq u \leq t} \{A_2(r, u) - (1 - \rho_1 - \epsilon)(s^\star - r) - \phi_2(u - s^\star)\} \\ &+ \sup_{r \leq s^\star} \{A_1(r, s^\star) - (\rho_1 + \epsilon)(s^\star - r)\}, \end{aligned}$$

and therefore (4.16) implies (4.13) or (4.15), which completes part *(ii)* and the proof. □

We now use the above sample-path relation to obtain a probabilistic upper bound. Denote $Q^{\rho_1+\epsilon}(s^\star, x, y) :=$

$$\mathbb{P}\left(\sup_{r\leq s^\star}\{A_1(r, s^\star) - (\rho_1+\epsilon)(s^\star - r)\} > y \mid A_1(s^\star, 0) + \phi_1 s^\star > x\right).$$

Corollary 4.3.2 *For any y, an upper bound for $\mathbb{P}(V_1 > x)$ is given by*

$$\mathbb{P}\left(V_1^{\phi_1} > x\right)\mathbb{P}\left(T_2^{1-\rho_1-\epsilon}(-2y) > \frac{(1-\alpha)x}{\hat{\rho}_1 - \phi_1}\right) + \mathbb{P}\left(V_1^{\phi_1} > x + y\right)$$
$$+\mathbb{P}\left(V_1^{\phi_1}\left[\frac{(1-\alpha)x}{\hat{\rho}_1 - \phi_1}\right] > x\right) + \mathbb{P}\left(V_1^{\phi_1} > x\right)Q^{\rho_1+\epsilon}(s^\star, x, y).$$

Proof Using Lemma 4.3.2, the independence of $A_1(s,t)$ and $A_2(s,t)$, and the fact that $A_1(s,t)$ and $A_2(s,t)$ have stationary increments, for all y and $w \geq 0$,

$$\begin{aligned}
&\mathbb{P}(V_1(t) > x)\\
\leq\ & \mathbb{P}\Big(\{A_1(s^\star,t) - \phi_1(t-s^\star) > x\} \cap \Big\{\{A_1(r^\star,s^\star) - (\rho_1+\epsilon)(s^\star - r^\star) > y\}\\
&\quad \cup \left\{V_1^{\phi_1}(t) > x+y\right\}\\
&\quad \cup \Big\{\inf_{s^\star\leq u\leq t}\{A_2(r^\star,u) - (1-\rho_1-\epsilon)(s^\star - r^\star) - \phi_2(u-s^\star)\} > -2y\Big\}\Big\}\Big)\\
\leq\ & \mathbb{P}(A_1(s^\star,t) - \phi_1(t-s^\star) > x, A_1(r^\star,s^\star) - (\rho_1+\epsilon)(s^\star - r^\star) > y)\\
&+\mathbb{P}\left(A_1(s^\star,t) - \phi_1(t-s^\star) > x, V_1^{\phi_1}(t) > x+y\right)\\
&+\mathbb{P}\Big(A_1(s^\star,t) - \phi_1(t-s^\star) > x,\\
&\qquad \inf_{s^\star\leq u\leq t}\{A_2(r^\star,u) - (1-\rho_1-\epsilon)(s^\star - r^\star) - \phi_2(u-s^\star)\} > -2y\Big)\\
\leq\ & \mathbb{P}(A_1(s^\star,t) - \phi_1(t-s^\star) > x, A_1(r^\star,s^\star) - (\rho_1+\epsilon)(s^\star - r^\star) > y)\\
&+\mathbb{P}\left(V_1^{\phi_1}(t) > x+y\right)\\
&+\mathbb{P}\Big(A_1(s^\star,t) - \phi_1(t-s^\star) > x,\\
&\qquad \inf_{s^\star\leq u\leq t}\{A_2(r^\star,u) - (1-\rho_1-\epsilon)(s^\star - r^\star) - \phi_2(u-s^\star)\} > -2y\Big)\\
=\ & \mathbb{P}(\exists r^\star \leq s^\star \leq t : A_1(s^\star,t) - \phi_1(t-s^\star) > x,\\
&\quad A_1(r^\star,s^\star) - (\rho_1+\epsilon)(s^\star - r^\star) > y) + \mathbb{P}\left(V_1^{\phi_1}(t) > x+y\right)\\
&+\mathbb{P}\Big(\exists r^\star \leq s^\star \leq t : A_1(s^\star,t) - \phi_1(t-s^\star) > x,\\
&\qquad \inf_{s^\star\leq u\leq t}\{A_2(r^\star,u) - (1-\rho_1-\epsilon)(s^\star - r^\star) - \phi_2(u-s^\star)\} > -2y\Big)
\end{aligned}$$

$$
\begin{aligned}
&= \mathbb{P}(\exists r^\star \le s^\star \le 0 : A_1(s^\star,0) + \phi_1 s^\star > x, \\
&\qquad A_1(r^\star,s^\star) - (\rho_1+\epsilon)(s^\star - r^\star) > y) + \mathbb{P}\left(V_1^{\phi_1}(0) > x+y\right) \\
&\quad + \mathbb{P}\Bigg(\exists r^\star \le s^\star \le 0 : A_1(s^\star,0) + \phi_1 s^\star > x, \\
&\qquad \inf_{s^\star \le u \le 0} \{A_2(r^\star,u) - (1-\rho_1-\epsilon)(s^\star - r^\star) - \phi_2(u - s^\star)\} > -2y\Bigg) \\
&= \mathbb{P}\Bigg(\exists s^\star \le 0 : A_1(s^\star,0) + \phi_1 s^\star > x, \\
&\qquad \sup_{r \le s^\star} \{A_1(r,s^\star) - (\rho_1+\epsilon)(s^\star - r)\} > y\Bigg) + \mathbb{P}\left(V_1^{\phi_1}(0) > x+y\right) \\
&\quad + \mathbb{P}\Bigg(\exists s^\star \le 0 : A_1(s^\star,0) + \phi_1 s^\star > x, \\
&\qquad \sup_{r \le s^\star} \inf_{s^\star \le u \le 0} \{A_2(r,u) - (1-\rho_1-\epsilon)(s^\star - r) - \phi_2(u - s^\star)\} > -2y\Bigg) \\
&\le \mathbb{P}\left(\exists s^\star \le 0 : A_1(s^\star,0) + \phi_1 s^\star > x\right) \\
&\quad \times \mathbb{P}\left(\sup_{r \le s^\star} \{A_1(r,s^\star) - (\rho_1+\epsilon)(s^\star - r)\} > y \mid A_1(s^\star,0) + \phi_1 s^\star > x\right) \\
&\quad + \mathbb{P}\left(V_1^{\phi_1}(0) > x+y\right) + \mathbb{P}\left(\exists s^\star \in [-w,0] : A_1(s^\star,0) + \phi_1 s^\star > x\right) \\
&\quad + \mathbb{P}\Bigg(\exists s^\star \le -w : A_1(s^\star,0) + \phi_1 s^\star > x, \\
&\qquad \sup_{r \le s^\star} \inf_{s^\star \le u \le s^\star + w} \{A_2(r,u) - (1-\rho_1-\epsilon)(s^\star - r) - \phi_2(u - s^\star)\} > -2y\Bigg) \\
&= \mathbb{P}\left(\exists s \ge 0 : A_1(-s,0) - \phi_1 s > x\right) Q^{\rho_1+\epsilon}(s^\star,x,y) \\
&\quad + \mathbb{P}\left(V_1^{\phi_1}(0) > x+y\right) + \mathbb{P}\left(\exists s \in [0,w] : A_1(-s,0) - \phi_1 s > x\right) \\
&\quad + \mathbb{P}\left(\sup_{r \le s^\star} \inf_{s^\star \le u \le s^\star + w} \{A_2(r,u) - (1-\rho_1-\epsilon)(s^\star - r) - \phi_2(u - s^\star)\} > -2y\right) \\
&\quad \times \mathbb{P}\left(\exists s^\star \le -w : A_1(s^\star,0) + \phi_1 s^\star > x\right) \\
&\le \mathbb{P}\left(\sup_{s \ge 0} \{A_1(-s,0) - \phi_1 s\} > x\right) Q^{\rho_1+\epsilon}(s^\star,x,y) \\
&\quad + \mathbb{P}\left(V_1^{\phi_1}(0) > x+y\right) + \mathbb{P}\left(\sup_{0 \le s \le w} \{A_1(-s,0) - \phi_1 s\} > x\right) \\
&\quad + \mathbb{P}\left(\exists s \ge 0 : A_1(-s,0) - \phi_1 s > x\right) \\
&\quad \times \mathbb{P}\left(\sup_{r \ge 0} \inf_{0 \le u \le w} \{A_2(-r,u) - (1-\rho_1-\epsilon) r - \phi_2 u\} > -2y\right) \\
&= \mathbb{P}\left(V_1^{\phi_1} > x\right) Q^{\rho_1+\epsilon}(s^\star,x,y) + \mathbb{P}\left(V_1^{\phi_1}(0) > x+y\right) + \mathbb{P}\left(V_1^{\phi_1}[w] > x\right)
\end{aligned}
$$

$$+\mathbb{P}\left(\sup_{s\geq 0}\{A_1(-s,0)-\phi_1 s\}>x\right)$$
$$\times\mathbb{P}\left(\inf_{0\leq u\leq w}\{\sup_{r\geq 0}\{A_2(-r,0)-(1-\rho_1-\epsilon)r\}+A_2(0,u)-\phi_2 u\}>-2y\right)$$
$$= \mathbb{P}\left(V_1^{\phi_1}>x\right)\mathbb{P}\left(T_2^{1-\rho_1-\epsilon}(-2y)>w\right)$$
$$+\mathbb{P}\left(V_1^{\phi_1}(0)>x+y\right)+\mathbb{P}\left(V_1^{\phi_1}[w]>x\right)+\mathbb{P}\left(V_1^{\phi_1}>x\right)Q^{\rho_1+\epsilon}(s^\star,x,y).$$

Taking $w=((1-\alpha)x)/(\hat{\rho}_1-\phi_1)$ completes the proof. □

4.4 Preliminary results for the light-tailed flow

In this section we prove some auxiliary results for flow 1 in isolation. The results will be crucial in obtaining the asymptotic behavior of $\mathbb{P}(V_1>x)$ in the GPS model as given in Theorem 4.2.1.

The following result is proven in [31] for a more general class of input processes than just Markov fluid sources. It states that a large workload of flow 1 is due to a temporary increase in the traffic intensity from ρ_1 to $\hat{\rho}_1$. Moreover, the time to reach level x will be at most $(1+\alpha)x/(\hat{\rho}_1-\phi_1)$ with high probability.

Lemma 4.4.1 *For any* $\alpha>0$,

$$\liminf_{x\to\infty}\frac{\mathbb{P}\left(V_1^{\phi_1}\left[\frac{(1+\alpha)x}{\hat{\rho}_1-\phi_1}\right]>x\right)}{\mathbb{P}\left(V_1^{\phi_1}>x\right)}=1, \tag{4.18}$$

with $\hat{\rho}_1=M'_{\phi_1}(\theta^\star)+\phi_1$.

The next lemma indicates that it is very unlikely that flow 1 generates an amount of traffic that is well *below* average during $(r,t^\star]$, given that it generates well above average during $(t^\star,0]$. In the proof we explicitly use the mild dependence of the amount of traffic generated in $(r,t^\star]$ and $(t^\star,0]$. They only depend on each other via the state of flow 1 at time $t^\star$.

Lemma 4.4.2 *For any* $\gamma,\epsilon>0$, $t^\star<0$,

$$\lim_{x\to\infty}\mathbb{P}\left(\sup_{r\leq t^\star}\{(\rho_1-\epsilon)(t^\star-r)-A_1(r,t^\star)\}\leq\gamma x\mid A_1(t^\star,0)+\phi_1 t^\star>x\right)=1.$$

Proof Recall that flow 1 is a Markov fluid source. We condition on the state of the underlying Markov chain at time $t^\star$. Let $E_j(t^\star)$ be the event that the state at time $t^\star$ is j, $j=1,\dots,d$, and $\pi_j(t^\star):=\mathbb{P}(E_j(t^\star)\mid A_1(t^\star,0)+\phi_1 t^\star>x)$. Then,

$$\mathbb{P}\left(\sup_{r\leq t^\star}\{(\rho_1-\epsilon)(t^\star-r)-A_1(r,t^\star)\}\leq\gamma x\mid A_1(t^\star,0)+\phi_1 t^\star>x\right)$$

$$=\sum_{j=1}^{d}\mathbb{P}\Big(\sup_{r\le t^\star}\{(\rho_1-\epsilon)(t^\star-r)-A_1(r,t^\star)\}\le\gamma x\mid A_1(t^\star,0)+\phi_1 t^\star>x,E_j(t^\star)\Big)$$
$$\times\pi_j(t^\star)$$
$$=\sum_{j=1}^{d}\mathbb{P}\left(\sup_{r\le t^\star}\{(\rho_1-\epsilon)(t^\star-r)-A_1(r,t^\star)\}\le\gamma x\mid E_j(t^\star)\right)\pi_j(t^\star).$$

The statement of the lemma then follows by observing that

$$\lim_{x\to\infty}\mathbb{P}\left(\sup_{r\le t^\star}\{(\rho_1-\epsilon)(t^\star-r)-A_1(r,t^\star)\}\le\gamma x\mid E_j(t^\star)\right)=1$$

for all $j=1,\dots,d$, since $\mathbb{E}A_1(-t,0)=\rho_1 t$. □

The next lemma shows that it is very unlikely that flow 1 generates traffic well *above* average in the last busy period before $t^\star$, given that it generates at least $\phi_1 t^\star+x$ in the interval $(t^\star,0]$. When combining the results from Lemma 4.4.2 and 4.4.3, flow 1 most likely produces on average before time $t^\star$ given that its buffer reaches level x at time 0.

Lemma 4.4.3 *For any $\gamma,\epsilon,\mu>0$, $t^\star<0$,*

$$\lim_{x\to\infty}x^{\mu}\mathbb{P}\left(\sup_{r\le t^\star}\{A_1(r,t^\star)-(\rho_1+\epsilon)(t^\star-r)\}>\gamma x\mid A_1(t^\star,0)+\phi_1 t^\star>x\right)=0.$$

Proof As in the proof of Lemma 4.4.2, let $E_j(t^\star)$ be the event that the state at time $t^\star$ is j, $j=1,\dots,d$, and $\pi_j(t^\star):=\mathbb{P}(E_j(t^\star)\mid A_1(t^\star,0)+\phi_1 t^\star>x)$. Then,

$$\mathbb{P}\left(\sup_{r\le t^\star}\{A_1(r,t^\star)-(\rho_1+\epsilon)(t^\star-r)\}>\gamma x\mid A_1(t^\star,0)+\phi_1 t^\star>x\right)$$
$$=\sum_{j=1}^{d}\mathbb{P}\Big(\sup_{r\le t^\star}\{A_1(r,t^\star)-(\rho_1+\epsilon)(t^\star-r)\}>\gamma x\mid A_1(t^\star,0)+\phi_1 t^\star>x,E_j(t^\star)\Big)$$
$$\times\pi_j(t^\star)$$
$$=\sum_{j=1}^{d}\mathbb{P}\left(\sup_{r\le t^\star}\{A_1(r,t^\star)-(\rho_1+\epsilon)(t^\star-r)\}>\gamma x\mid E_j(t^\star)\right)\pi_j(t^\star).$$

The statement of the lemma then follows by observing that there exist constants C, $\theta_0^\star$ (independent of j) such that

$$\mathbb{P}\left(\sup_{r\le t^\star}\{A_1(r,t^\star)-(\rho_1+\epsilon)(t^\star-r)\}>\gamma x\mid E_j(t^\star)\right)\le Ce^{-\theta_0^\star x},$$

where $\theta_0^\star>0$ is the solution of $M_{\rho_1+\epsilon}(\theta)=0$. In [88, Section 4] it is shown that C can be expressed in terms of the dominant eigenvalue of the matrix $\Lambda+\theta_0^\star R$ and the corresponding (component-wise positive) eigenvalue. □

Using the following relation,

$$\frac{\mathbb{P}\left(V_1^{\phi_1} > (1+\gamma)x\right)}{\mathbb{P}\left(V_1^{\phi_1} > x\right)} = \mathbb{P}\left(V_1^{\phi_1} > (1+\gamma)x \mid V_1^{\phi_1} > x\right),$$

it is easily seen that the next lemma can be interpreted as follows. It is very unlikely that the workload of flow 1 reaches a value considerably larger than x, given that it reaches level x.

Lemma 4.4.4 *For any* $\gamma, \mu > 0$,

$$\limsup_{x\to\infty} \frac{x^{\mu}\mathbb{P}\left(V_1^{\phi_1} > (1+\gamma)x\right)}{\mathbb{P}\left(V_1^{\phi_1} > x\right)} = 0.$$

Proof The proof follows immediately from the fact that $\mathbb{P}(V_1^{\phi_1} > x)$ decays exponentially at rate $\theta^\star$, where $\theta^\star > 0$ is the solution of $M(\theta) = 0$ (see Proposition 4.1.1 with $c = \phi_1$). □

The last lemma of this section can be seen as the companion of Lemma 4.4.1. It states that the time to reach a large level x will be at least $(1-\alpha)x/(\hat{\rho}_1 - \phi_1)$ with large probability. Combining this result with Lemma 4.4.1, we obtain that the most-likely epoch of reaching a large workload level x for flow 1, when served at rate ϕ_1 in isolation, is $x/(\hat{\rho}_1 - \phi_1)$.

Lemma 4.4.5 *For any* $\alpha, \mu > 0$,

$$\limsup_{x\to\infty} \frac{x^{\mu}\mathbb{P}\left(V_1^{\phi_1}\left[\frac{(1-\alpha)x}{\hat{\rho}_1-\phi_1}\right] > x\right)}{\mathbb{P}\left(V_1^{\phi_1} > x\right)} = 0.$$

Proof The proof consists of three steps. (i) First we give a sufficient condition for the lemma to hold, explicitly using the fact that the Markov fluid source has a bounded peak rate R_P. (ii) Then we estimate the decay rate of the event that a queue with service rate ϕ_1 fed by a Markov fluid source reaches overflow at time t. (iii) Finally we identify the most-likely epoch of overflow, and show that this implies the required property.

Ad (i). Obviously,

$$\begin{aligned}\mathbb{P}\left(V_1^{\phi_1}\left[\frac{(1-\alpha)x}{\hat{\rho}_1-\phi_1}\right] > x\right) &\leq \mathbb{P}\left(\exists t \leq T_x(\alpha) : A_1(0,t) - \phi_1 t > x\right) \\ &\leq \sum_{t=0}^{T_x(\alpha)} \mathbb{P}\left(A_1(0,t) - \phi_1 t > x - (R_P - \phi_1)\right),\end{aligned}$$

with

$$T_x(\alpha) := \left\lceil \frac{(1-\alpha)x}{\hat{\rho}_1 - \phi_1} \right\rceil.$$

From

$$\begin{aligned}
&\max_{t=0,\dots,T_x(\alpha)} \mathbb{P}\left(A_1(0,t)-\phi_1 t > x-(R_P-\phi_1)\right) \\
\leq\ & \sum_{t=0}^{T_x(\alpha)} \mathbb{P}\left(A_1(0,t)-\phi_1 t > x-(R_P-\phi_1)\right) \\
\leq\ & (T_x(\alpha)+1) \max_{t=0,\dots,T_x(\alpha)} \mathbb{P}\left(A_1(0,t)-\phi_1 t > x-(R_P-\phi_1)\right)
\end{aligned}$$

and

$$\lim_{x\to\infty} \frac{1}{x}\log(T_x(\alpha)+1)=0,$$

we find that

$$\begin{aligned}
&\limsup_{x\to\infty} \frac{1}{x}\log \sum_{t=0}^{T_x(\alpha)} \mathbb{P}\left(A_1(0,t)-\phi_1 t > x-(R_P-\phi_1)\right) \\
=\ & \limsup_{x\to\infty} \frac{1}{x}\log \max_{t=0,\dots,T_x(\alpha)} \mathbb{P}\left(A_1(0,t)-\phi_1 t > x-(R_P-\phi_1)\right) \\
\leq\ & \limsup_{x\to\infty} \frac{1}{x}\log \sup_{t\in[0,T_x(\alpha)]} \mathbb{P}\left(A_1(0,t)-\phi_1 t > x-(R_P-\phi_1)\right) \\
\leq\ & \limsup_{x\to\infty} \frac{1}{x}\log \sup_{t\in[S_x,T_x(\alpha)]} \mathbb{P}\left(A_1(0,t)-\phi_1 t > x-(R_P-\phi_1)\right), \quad (4.19)
\end{aligned}$$

with $S_x := (x-(R_P-\phi_1))/(R_P-\phi_1)$. Notice that we can indeed exclude all t smaller than S_x from the optimization, because in that range no overflow is possible. Clearly, we have proven the result if we show that the latter decay rate is strictly smaller than $\theta^\star$ for all $\alpha>0$ (recall that $\mathbb{P}(V_1^{\phi_1}>x)$ decays at rate $\theta^\star$).

Ad (ii). For x large enough, and all t between S_x and $T_x(\alpha)$, due to Chebychev's inequality, and Proposition 4.1.1 with $c=\phi_1$,

$$\begin{aligned}
\mathbb{P}\left(A_1(0,t)-\phi_1 t > x-(R_P-\phi_1)\right) &\leq \inf_{\theta>0} \frac{\mathbb{E}e^{\theta(A_1(0,t)-\phi_1 t)}}{e^{\theta(x-(R_P-\phi_1))}} \\
&\leq C \inf_{\theta>0} \frac{e^{M(\theta)t}}{e^{\theta(x-(R_P-\phi_1))}}. \quad (4.20)
\end{aligned}$$

Now replace t in (4.19) by

$$t_x(\beta)=\frac{(1-\beta)x}{\hat{\rho}_1-\phi_1},$$

then the supremum over t in (4.19) turns into a supremum over

$$\beta\in\left[\alpha, 1-\frac{(\hat{\rho}_1-\phi_1)x+(\hat{\rho}_1-\phi_1)(\phi_1-R_P)}{(R_P-\phi_1)x}\right].$$

The next step is to calculate the infimum over $\theta > 0$ in (4.20) for fixed β. Differentiating to θ we obtain the first-order condition:

$$M'(\theta) = \frac{(x-(R_P-\phi_1))(\hat{\rho}_1-\phi_1)}{(1-\beta)x}. \tag{4.21}$$

Recall that $\theta^\star$ is such that $M(\theta^\star) = 0$ and that $M'(\theta^\star) = \hat{\rho}_1 - \phi_1 > 0$, see Proposition 4.1.1 with $c = \phi_1$ and the definition of $\hat{\rho}_1$ as given in Theorem 4.2.1. Using a Taylor approximation for $M'(\theta)$ in $\theta^\star$ then yields

$$M'(\theta) = \hat{\rho}_1 - \phi_1 + M''(\theta^\star)(\theta-\theta^\star) + \mathrm{O}\left((\theta-\theta^\star)^2\right).$$

Observe that the right-hand side of (4.21) goes to $(\hat{\rho}_1-\phi_1)/(1-\beta)$ for $x \to \infty$. Let us denote the solution of

$$\hat{\rho}_1 - \phi_1 + M''(\theta^\star)(\theta-\theta^\star) = \frac{\hat{\rho}_1-\phi_1}{1-\beta}$$

to θ by θ_β. Then it is elementary to show that

$$\theta_\beta = \frac{\beta}{1-\beta}\frac{\hat{\rho}_1-\phi_1}{M''(\theta^\star)} + \theta^\star.$$

Using the Taylor approximation for $M(\theta)$ in $M(\theta^\star)$:

$$M(\theta) = M(\theta^\star) + M'(\theta^\star)(\theta-\theta^\star) + \frac{1}{2}M''(\theta^\star)(\theta-\theta^\star)^2 + \mathrm{O}\left((\theta-\theta^\star)^3\right),$$

we obtain for $M(\theta_\beta)$:

$$\begin{aligned} M(\theta_\beta) &= M'(\theta^\star)\frac{\beta}{1-\beta}\frac{\hat{\rho}_1-\phi_1}{M''(\theta^\star)} + \frac{1}{2}M''(\theta^\star)\left(\frac{\beta}{1-\beta}\frac{\hat{\rho}_1-\phi_1}{M''(\theta^\star)}\right)^2 \\ &\quad + \mathrm{O}\left(\left(\frac{\beta}{1-\beta}\frac{\hat{\rho}_1-\phi_1}{M''(\theta^\star)}\right)^3\right). \end{aligned}$$

Hence, after some calculations,

$$\begin{aligned} \lim_{x\to\infty}\frac{1}{x}\log\inf_{\theta>0}\frac{e^{t_x(\beta)M(\theta)}}{e^{\theta(x-(R_P-\phi_1))}} &= \lim_{x\to\infty}\frac{1}{x}\left(t_x(\beta)M(\theta_\beta) - \theta_\beta x\right) \\ &= -\frac{1}{2}\frac{\hat{\rho}_1-\phi_1}{M''(\theta^\star)}\beta^2 - \theta^\star + \mathrm{O}(\beta^3), \end{aligned}$$

where the convexity of $M(\cdot)$ implies that

$$\frac{\hat{\rho}_1-\phi_1}{M''(\theta^\star)} > 0.$$

Ad (iii). Recall that we have to perform the optimization over β. The supremum over β is clearly attained at $\beta = \alpha > 0$. Now the stated follows from the fact that $\mathbb{P}(V_1^{\phi_1} > x)$ decays at rate $\theta^\star$, as explained in the first step of the proof. □

4.5 Asymptotic analysis

We now use the results from the previous section to show that the lower and upper bounds for $\mathbb{P}(V_1 > x)$ of Section 4.3 asymptotically coincide, resulting in the product form of (4.4). For the proof, we need to make certain assumptions on the behavior of the drain time distribution $\mathbb{P}\big(T_2^{1-\rho_1} > x/(\hat{\rho}_1 - \phi_1)\big)$. In later sections, we will determine the specific form of the drain time distribution, and find that flow 2 indeed satisfies these assumptions. For notational convenience, we frequently switch to a variable $\hat{x}$, which should be thought of as playing the role of $x/(\hat{\rho}_1 - \phi_1)$.

Lemma 4.5.1 *Under Assumptions 4.5.1-4.5.3 listed below with* $c = 1 - \rho_1$,

$$\mathbb{P}(V_1 > x) \sim \mathbb{P}\left(V_1^{\phi_1} > x\right)\mathbb{P}\left(T_2^{1-\rho_1} > \frac{x}{\hat{\rho}_1 - \phi_1}\right).$$

Assumption 4.5.1 *For any* $\alpha, \gamma, \epsilon > 0$, *either (a)*

$$\liminf_{\hat{x}\to\infty} \frac{\mathbb{P}\left(T_2^{c+\epsilon}(\gamma\hat{x}) > (1+\alpha)\hat{x}\right)}{\mathbb{P}\left(T_2^c > \hat{x}\right)} = F^c(\alpha, \gamma, \epsilon),$$

with $\lim_{\alpha,\gamma,\epsilon\downarrow 0} F^c(\alpha, \gamma, \epsilon) = 1$, *or (b)*

$$\liminf_{\hat{x}\to\infty} \frac{\mathbb{P}\left(T_2 > (1+\alpha)\hat{x}\right)}{\mathbb{P}\left(T_2^c > \hat{x}\right)} = F(\alpha),$$

with $\lim_{\alpha\downarrow 0} F(\alpha) = 1$.

Assumption 4.5.2 *For any* $\alpha, \gamma, \epsilon > 0$,

$$\limsup_{\hat{x}\to\infty} \frac{\mathbb{P}\left(T_2^{c-\epsilon}(-\gamma\hat{x}) > (1-\alpha)\hat{x}\right)}{\mathbb{P}\left(T_2^c > \hat{x}\right)} = G^c(\alpha, \gamma, \epsilon),$$

with $\lim_{\alpha,\gamma,\epsilon\downarrow 0} G^c(\alpha, \gamma, \epsilon) = 1$.

Assumption 4.5.3 *For some* $\mu > 0$,

$$\liminf_{x\to\infty} \hat{x}^{\mu}\mathbb{P}\left(T_2^c > \hat{x}\right) \geq 1.$$

Proof of Lemma 4.5.1 The proof consists of a lower bound and an upper bound which asymptotically coincide. We start with the lower bound. We distinguish between two cases: Assumption 4.5.1 (a) and Assumption 4.5.1 (b).

(a) Using Corollary 4.3.1 with $v = \infty$, $y = \gamma x/(\hat{\rho}_1 - \phi_1)$, Lemma 4.4.1, and Lemma 4.4.2,

$$\liminf_{x\to\infty} \frac{\mathbb{P}(V_1 > x)}{\mathbb{P}\left(V_1^{\phi_1} > x\right)\mathbb{P}\left(T_2^{1-\rho_1} > \frac{x}{\hat{\rho}_1-\phi_1}\right)}$$

$$\geq \liminf_{x\to\infty} \frac{\mathbb{P}\left(V_1^{\phi_1}\left[\frac{(1+\alpha)x}{\hat{\rho}_1-\phi_1}\right] > x\right)}{\mathbb{P}\left(V_1^{\phi_1} > x\right)} \liminf_{x\to\infty} \frac{\mathbb{P}\left(T_2^{1-\rho_1+\epsilon}\left(\frac{\gamma x}{\hat{\rho}_1-\phi_1}\right) > \frac{(1+\alpha)x}{\hat{\rho}_1-\phi_1}\right)}{\mathbb{P}\left(T_2^{1-\rho_1} > \frac{x}{\hat{\rho}_1-\phi_1}\right)}$$

$$\times \liminf_{x\to\infty} \mathbb{P}\Bigg(\sup_{r\leq s^\star}\{(\rho_1-\epsilon)(s^\star - r) - A_1(r,s^\star)\} \leq \frac{\gamma x}{\hat{\rho}_1-\phi_1} \mid A_1(s^\star,0)+\phi_1 s^\star > x\Bigg)$$

$$= F^{1-\rho_1}(\alpha,\gamma,\epsilon).$$

Letting $\alpha, \gamma, \epsilon \downarrow 0$ completes the proof.

(b) Using Corollary 4.3.1 with $v = 0$, $y = 0$, and Lemma 4.4.1, we obtain, observing that $P^{\rho_1-\epsilon}(s^\star, 0, x, 0) = 1$,

$$\liminf_{x\to\infty} \frac{\mathbb{P}(V_1 > x)}{\mathbb{P}\left(V_1^{\phi_1} > x\right)\mathbb{P}\left(T_2^{1-\rho_1} > \frac{x}{\hat{\rho}_1-\phi_1}\right)}$$

$$\geq \liminf_{x\to\infty} \frac{\mathbb{P}\left(V_1^{\phi_1}\left[\frac{(1+\alpha)x}{\hat{\rho}_1-\phi_1}\right] > x\right)}{\mathbb{P}\left(V_1^{\phi_1} > x\right)} \liminf_{x\to\infty} \frac{\mathbb{P}\left(T_2 > \frac{(1+\alpha)x}{\hat{\rho}_1-\phi_1}\right)}{\mathbb{P}\left(T_2^{1-\rho_1} > \frac{x}{\hat{\rho}_1-\phi_1}\right)}$$

$$= F(\alpha).$$

Then let $\alpha \downarrow 0$.

We now turn to the upper bound. Using Corollary 4.3.2 with $v = \infty$, $y = \gamma x/(2(\hat{\rho}_1 - \phi_1))$, Lemmas 4.4.3-4.4.5, and Assumptions 4.5.2, 4.5.3, for some $\mu > 0$,

$$\limsup_{x\to\infty} \frac{\mathbb{P}(V_1 > x)}{\mathbb{P}\left(V_1^{\phi_1} > x\right)\mathbb{P}\left(T_2^{1-\rho_1} > \frac{x}{\hat{\rho}_1-\phi_1}\right)}$$

$$\leq \limsup_{x\to\infty} \frac{x^\mu \mathbb{P}\left(V_1^{\phi_1}\left[\frac{(1-\alpha)x}{\hat{\rho}_1-\phi_1}\right] > x\right)}{\mathbb{P}\left(V_1^{\phi_1} > x\right)} + \limsup_{x\to\infty} \frac{x^\mu \mathbb{P}\left(V_1^{\phi_1} > \left(1+\frac{\gamma}{2(\hat{\rho}_1-\phi_1)}\right)x\right)}{\mathbb{P}\left(V_1^{\phi_1} > x\right)}$$

$$+ \limsup_{x\to\infty} \frac{\mathbb{P}\left(T_2^{1-\rho_1-\epsilon}\left(\frac{-\gamma x}{\hat{\rho}_1-\phi_1}\right) > \frac{(1-\alpha)x}{\hat{\rho}_1-\phi_1}\right)}{\mathbb{P}\left(T_2^{1-\rho_1} > \frac{x}{\hat{\rho}_1-\phi_1}\right)}$$

$$+ \limsup_{x\to\infty} x^\mu \mathbb{P}\Bigg(\sup_{r\leq s^\star}\{A_1(r,s^\star) - (\rho_1+\epsilon)(s^\star - r)\} > \frac{\gamma x}{2(\hat{\rho}_1-\phi_1)} \mid A_1(s^\star,0)+\phi_1 s^\star > x\Bigg)$$

$$= G^{1-\rho_1}(\alpha,\gamma,\epsilon).$$

Letting $\alpha, \gamma, \epsilon \downarrow 0$ completes the proof. □

In order to complete the proof of Theorem 4.2.1, it remains to be shown that flow 2 satisfies Assumptions 4.5.1-4.5.3 above, with $\mathbb{P}(T_2^{1-\rho_1} > x/(\hat{\rho}_1 - \phi_1))$ as in (4.5)-(4.7). This is done in the next four sections.

Remark 4.5.1 *As the proof shows, Lemma 4.5.1 and thus Theorem 4.2.1 remain true as long as flow 2 satisfies Assumptions 4.5.1-4.5.3 and Lemmas 4.4.1-4.4.5 hold for flow 1. Both seem to be the case under somewhat milder assumptions than made in Sections 4.1.1 and 4.1.2.*

In particular, for the light-tailed flow, the results in [63] suggest that Lemmas 4.4.1-4.4.5 hold for a more general class of arrival processes than just Markov fluid. Upon inspection of the proofs in the previous section, we see that two properties were explicitly exploited. In the first place it was repeatedly used that the source has a bounded peak rate. Secondly, it is required that the dependence between $A_1(r, t^*)$ and $A_1(t^*, 0)$ is rather mild. This leads us to believe that the lemmas still hold if the exponential sojourn times of the Markov fluid source are replaced by other light-tailed random variables. Probably, an essential prerequisite is that the light-tailed arrival process allows application of the Gärtner-Ellis large-deviations theorem. In particular, this requires that the log moment generating function of the amount of traffic generated in an interval of length t grows at most linearly:

$$\lim_{t\to\infty} \frac{1}{t} \log \mathbb{E}\left[\log \exp\{\theta A_1(0, t)\}\right] < \infty$$

for some positive θ. This rules out input processes such as fractional Brownian motion (with Hurst parameter $H \in (\frac{1}{2}, 1)$), or $M/G/\infty$-type inputs with heavy tailed job sizes.

For the heavy-tailed flow, the results may be extended to semi-Markov fluid input or mixtures of fluid input and instantaneous input. It may also be possible to extend the results to a larger class of subexponential distributions, although that would require elaborate refinements in the proofs. In a somewhat related context, [25], [70] provide a sharp demarcation of the distributional conditions for a so-called reduced-load equivalence to hold. We expect that in general there is a complicated trade-off between the assumptions on the light-tailed flow and the conditions imposed on the heavy-tailed flow.

4.6 Preliminary results for the heavy-tailed flow

To determine the behavior of $\mathbb{P}(T_2^{1-\rho_1} > x/(\hat{\rho}_1 - \phi_1))$ as $x \to \infty$, we will reduce the space of all relevant sample paths to a single most-likely scenario, which occurs with overwhelming probability. In this section, we establish some preliminary results which we will use to neglect the contribution of all non-dominant scenarios. All results are derived under the assumptions on flow 2 as given in Section 4.1.

Large-deviations arguments for heavy-tailed distributions suggest that a persistent backlog as associated with the event $T_2^{1-\rho_1} > x/(\hat{\rho}_1 - \phi_1)$, for large x, is most likely due to just a single large burst or long on period. To formalize this idea, we first introduce some additional notation. A burst is called large if the size exceeds $\kappa\hat{x}$, with $\kappa > 0$ some constant, independent of $\hat{x}$. Also, an on period is called long if the length exceeds $\kappa\hat{x}$. In case of instantaneous input, we denote by $\mathcal{N}_{\kappa\hat{x}}[l,r]$ the number of large bursts of flow 2 arriving in the time interval $[l,r]$. In case of an on-off process, we define $\mathcal{N}_{\kappa\hat{x}}[l,r]$ as the number of long on periods overlapping with the time interval $[l,r]$, including the on periods which may be in progress at time l and r.

Depending on the traffic scenario, we denote by $N(t)$ either the number of bursts or the number of on periods of flow 2 in the time interval $[0,t]$. An upper bound for this process is given by

$$N(t) \leq N_U(t) := \sup\left\{n : \sum_{i=1}^{n} U_{2,i} \leq t\right\} + 1,$$

with $U_{2,i}$ i.i.d. random variables representing either interarrival times or off periods of flow 2, depending on the traffic scenario.

We now state a crucial lemma which will allow us to limit the attention to large bursts and long on periods, and replace all remaining traffic activity by its average rate. The lemma is a minor modification of Lemma 3 in [118].

Lemma 4.6.1 *Let $S_n = X_1 + \cdots + X_n$ be a random walk with i.i.d. step sizes such that $\mathbb{E}X_1 < 0$ and $\mathbb{E}[(\max\{X_1, 0\})^p] < \infty$ for some $p > 1$. Then, for any $\mu < \infty$, there exists a $\kappa^\star > 0$ and a function $\phi(\cdot) \in \mathcal{R}_{-\mu}$ such that for all $\kappa \in (0, \kappa^\star]$,*

$$\mathbb{P}\left(S_n > \hat{x} | X_i \leq \kappa\hat{x}, i = 1, \ldots, n\right) \leq \phi(\hat{x})$$

for all n and $\hat{x}$.

Note that if X_i can be represented as the difference of two non-negative independent random variables $X_{i,1}$ and $X_{i,2}$, then the lemma remains valid if the X_i's are replaced by the $X_{i,1}$'s.

We now use the above lemma to show that the traffic process of flow 2 cannot significantly deviate from the normal drift over intervals of the order $\hat{x}$ when there are no large bursts.

Lemma 4.6.2 (Instantaneous input) *For any $\eta, \theta > 0$, a $\kappa^\star > 0$ exists such that for all $\kappa \in (0, \kappa^\star]$, as $\hat{x} \to \infty$,*

$$\mathbb{P}\left(T_2((\theta - (\phi_2 - \rho_2)\eta)\hat{x}) > \eta\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0, \eta\hat{x}] = 0\right) = \mathrm{o}(\mathbb{P}\left(B_2^{\mathrm{r}} > \hat{x}(\phi_2 - \rho_2)\right)).$$

Proof The event $T_2((\theta - (\phi_2 - \rho_2)\eta)\hat{x}) > \eta\hat{x}$ means that

$$\inf_{0 \leq u \leq \eta\hat{x}} \{A_2(0, u) - \phi_2 u\} > (\theta - (\phi_2 - \rho_2)\eta)\hat{x},$$

which in particular implies that

$$A_2(0, \eta\hat{x}) - \phi_2\eta\hat{x} > (\theta - (\phi_2 - \rho_2)\eta)\hat{x},$$

or equivalently,

$$A_2(0, \eta\hat{x}) - (\rho_2 + \theta/2\eta)\eta\hat{x} > \theta\hat{x}/2,$$

so that also

$$\sup_{0\le u\le \eta\hat{x}} \{A_2(0,u) - (\rho_2 + \theta/2\eta)u\} > \theta\hat{x}/2.$$

Now let $S_n := X_1 + \cdots + X_n$ be a random walk where the step sizes are denoted by $X_i := B_{2,i} - (\rho_2 + \theta/2\eta)\, U_{2,i}$, with $U_{2,i}$ and $B_{2,i}$ i.i.d. random variables representing the interarrival times and burst sizes of flow 2, respectively. Note that X_i represents the net increase in the workload in a queue with service rate $\rho_2 + \theta/2\eta$ between two consecutive bursts, and that $\mathbb{E}X_i < 0$. Because of the saw-tooth nature of the process $\{A_2(0,u) - (\rho_2 + \theta/2\eta)u\}$, we have

$$\sup_{0\le u\le t} \{A_2(0,u) - (\rho_2 + \theta/2\eta)u\} \le B_{2,0} + \sup_{1\le n\le N(t)} S_n.$$

Thus,

$$\begin{aligned}
&\mathbb{P}\left(T_2((\theta - (\phi_2 - \rho_2)\eta)\hat{x}) > \eta\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0,\eta\hat{x}] = 0\right)\\
\le\ & \mathbb{P}\left(B_{2,0} + \sup_{1\le n\le N(\eta\hat{x})} S_n \ge \theta\hat{x}/2, \mathcal{N}_{\kappa\hat{x}}[0,\eta\hat{x}] = 0\right)\\
\le\ & \mathbb{P}\left(B_{2,0} + \sup_{1\le n\le N(\eta\hat{x})} S_n \ge \theta\hat{x}/2 \mid \mathcal{N}_{\kappa\hat{x}}[0,\eta\hat{x}] = 0\right)\\
\le\ & \mathbb{P}\left(B_{2,0} + \sup_{1\le n\le N(\eta\hat{x})} S_n \ge \theta\hat{x}/2 \mid B_{2,i} \le \kappa\hat{x}, i \ge 0\right)\\
\le\ & \mathbb{P}\left(\sup_{1\le n\le N(\eta\hat{x})} S_n \ge (\theta/2 - \kappa)\hat{x} \mid B_{2,i} \le \kappa\hat{x}, i \ge 1\right)\\
\le\ & \mathbb{P}\left(\sup_{1\le n\le (\lambda_2+\epsilon)\eta\hat{x}} S_n \ge (\theta/2 - \kappa)\hat{x} \mid B_{2,i} \le \kappa\hat{x}, i \ge 1\right)\\
& +\mathbb{P}\left(N(\eta\hat{x}) > (\lambda_2 + \epsilon)\eta\hat{x}\right)\\
\le\ & \sum_{i=1}^{(\lambda_2+\epsilon)\eta\hat{x}} \mathbb{P}\left(S_n \ge (\theta/2 - \kappa)\hat{x} \mid B_{2,i} \le \kappa\hat{x}, i = 1,\dots,n\right)\\
& +\mathbb{P}\left(N(\eta\hat{x}) > (\lambda_2 + \epsilon)\eta\hat{x}\right).
\end{aligned}$$

Observe that the second term decays exponentially fast as $\hat{x} \to \infty$. According to Lemma 4.6.1, there exists a $\kappa^\star > 0$ and a function $\phi(\cdot) \in \mathcal{R}_{-\mu}$, $\mu > \nu_2$, such that for all $\kappa \in (0, \kappa^\star]$, each of the probabilities in the first term is upper bounded by $\phi(\hat{x})$. The statement then follows. □

We now formulate the counterpart of the above lemma for on-off processes, meaning that the traffic process of flow 2 closely follows the drift over intervals of the order $\hat{x}$ when there are no long on periods.

Lemma 4.6.3 (Fluid input) *For any $\eta, \theta > 0$, a $\kappa^\star > 0$ exists such that for all $\kappa \in (0, \kappa^\star]$, as $\hat{x} \to \infty$,*

$$\mathbb{P}\left(T_2((\theta - (\phi_2 - \rho_2)\eta)\hat{x}) > \eta\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0,\eta\hat{x}] = 0\right) = o\left(\mathbb{P}\left(A_2^{\mathrm{r}} > \frac{\hat{x}(\phi_2 - \rho_2)}{r_2 - \rho_2}\right)\right).$$

Proof Let $S_n := X_1 + \cdots + X_n$ be a random walk where the step sizes are denoted by $X_i := (r_2 - \rho_2 - \theta/2\eta)\ A_{2,i}\ -(\rho_2 + \theta/2\eta)\ U_{2,i}$, with $A_{2,i}$ and $U_{2,i}$ i.i.d. random variables representing the on periods and the off periods of flow 2, respectively. Note that X_i represents the net increase in the workload in a queue with service rate $\rho_2 + \theta/2\eta$ during an off period and a consecutive on period, and that $\mathbb{E}X_i < 0$. Because of the saw-tooth nature of the process $\{A_2(0,u) - (\rho_2 + \theta/2\eta)u\}$, we have

$$\sup_{0\le u\le t} \{A_2(0,u) - (\rho_2 + \theta/2\eta)u\} \le (r_2 - \rho_2)A_{2,0} + \sup_{1\le n\le N(t)} S_n \le$$

$$(r_2 - \rho_2)A_{2,0} + \sup_{1\le n\le N_U(t)} S_n.$$

The remainder of the proof is similar to that of Lemma 4.6.2. □

We now prove that it is relatively unlikely for flow 2 to generate two or more large bursts in an interval of order $\hat{x}$.

Lemma 4.6.4 (Instantaneous input) *For any $\alpha < 1$, $\kappa > 0$, as $\hat{x} \to \infty$,*

$$\mathbb{P}\left(\mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha)\hat{x}] \ge 2\right) = \mathrm{o}(\mathbb{P}\left(B_2^{\mathrm{r}} > \hat{x}(\phi_2 - \rho_2)\right)).$$

Proof By definition,

$$\mathbb{P}(\mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha)\hat{x}] \ge 2) = \mathbb{P}(\#\{j \in \{1,\ldots,N((1-\alpha)\hat{x})\} : B_{2,j} \ge \kappa\hat{x}\} \ge 2)$$
$$\le \mathbb{P}(\#\{j \in \{1,\ldots,N_U((1-\alpha)\hat{x})\} : B_{2,j} \ge \kappa\hat{x}\} \ge 2).$$

Now condition on $N_U((1-\alpha)\hat{x})$. This yields the following upper bound

$$\mathbb{E}\left[(N_U((1-\alpha)\hat{x}))^2\right]\mathbb{P}(B_2 \ge \kappa\hat{x})^2.$$

Finally, observe that $\mathbb{E}[(N_U((1-\alpha)\hat{x}))^2]$ is quadratic in $\hat{x}$ for $\hat{x} \to \infty$. □

We now state the counterpart of the above lemma for on-off processes, meaning that the probability that flow 2 experiences at least two long on periods during an interval of order $\hat{x}$ is negligibly small.

Lemma 4.6.5 (Fluid input) *For any $\alpha < 1$, $\kappa > 0$, as $\hat{x} \to \infty$,*

$$\mathbb{P}\left(\mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha)\hat{x}] \ge 2\right) = \mathrm{o}\left(\mathbb{P}\left(A_2^{\mathrm{r}} > \frac{\hat{x}(\phi_2 - \rho_2)}{r_2 - \rho_2}\right)\right).$$

Proof This lemma is a variant of a technical result in [152]. Note that

$$\begin{aligned} &\mathbb{P}\left(\mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha)\hat{x}] \ge 2\right) \\ \le\ & (1-p_2)\mathbb{P}\left(A_2^{\mathrm{r}} \ge \kappa\hat{x}\right)\mathbb{P}\left(\#\{j \in \{1,\ldots,N_U((1-\alpha)\hat{x})\} : A_{2,j} \ge \kappa\hat{x}\} \ge 1\right) \\ &+\mathbb{P}\left(\mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha)\hat{x}] \ge 2, \text{flow 2 is off at time 0}\right). \end{aligned}$$

By conditioning upon $N_U((1-\alpha)\hat{x})$, one can bound the second probability in the first term by $\mathbb{E}[N_U((1-\alpha)\hat{x})]\mathbb{P}(A_2 \geq \kappa\hat{x})$. The first factor is linear in $\hat{x}$ for $\hat{x} \to \infty$, whereas the second is in $\mathcal{R}_{-\nu_2}$. Hence, the first term is in $\mathcal{R}_{2(1-\nu_2)}$. To bound the second term, condition (again) on $N_U((1-\alpha)\hat{x})$. This yields

$$\begin{aligned} & \mathbb{P}\left(\mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha)\hat{x}] \geq 2, \text{flow 2 is off at time } 0\right) \\ \leq \; & \mathbb{E}[(N_U((1-\alpha)\hat{x}))^2]\mathbb{P}\left(A_2 \geq \kappa\hat{x}\right)^2 . \end{aligned}$$

Finally, note that, as in Lemma 4.6.4, $\mathbb{E}[(N_U((1-\alpha)\hat{x}))^2]$ is quadratic in $\hat{x}$ for $\hat{x} \to \infty$. □

We now prove that the amount of traffic generated by flow 2 after turning off is not below average by any significant margin.

Lemma 4.6.6 (Fluid input) *Suppose that flow 2 turns off at time v. Then for any $\delta, \theta > 0$,*

$$\lim_{\hat{x}\to\infty} \mathbb{P}\left(\sup_{u \geq v}\{(\rho_2-\delta)(u-v) - A_2(v,u)\} \leq \theta\hat{x}\right) = 1.$$

Proof Let $S_n := X_1+\cdots+X_n$ be a random walk with step sizes $X_i := (\rho_2-\delta-r_2)A_{2,i}+(\rho_2-\delta)U_{2,i}$, with $A_{2,i}$ and $U_{2,i}$ i.i.d. random variables representing the on periods and the off periods of flow 2, respectively. Note that X_i represents the net decrease in the workload in a queue with service rate $\rho_2-\delta$ fed by flow 2 during an on period and consecutive off period, and that $\mathbb{E}X_i < 0$. Now observe that

$$\sup_{u\geq v}\{(\rho_2-\delta)(u-v) - A_2(v,u)\} \leq (\rho_2-\delta)U_{2,0} + \sup_{n\geq 1} S_n,$$

so that

$$\begin{aligned} & \mathbb{P}\left(\sup_{u\geq v}\{(\rho_2-\delta)(u-v) - A_2(v,u)\} \leq \theta\hat{x}\right) \\ = \; & 1 - \mathbb{P}\left(\sup_{u\geq v}\{(\rho_2-\delta)(u-v) - A_2(v,u)\} > \theta\hat{x}\right) \\ \geq \; & 1 - \mathbb{P}\left((\rho_2-\delta)U_{2,0} + \sup_{n\geq 1} S_n > \theta\hat{x}\right). \end{aligned}$$

The probability in the last term goes to 0 as $\hat{x} \to \infty$ for any $\theta > 0$, since the maximum of a random walk with negative drift is finite with probability 1. □

The following lemma shows that it is not likely for flow 2 to have a workload of at least order $\hat{x}$ at time 0 and to generate at the same time one large burst in an interval of order $\hat{x}$.

Lemma 4.6.7 (Instantaneous input) *For any $0 < \xi < 1-\alpha$, $\zeta, \kappa > 0$, as $\hat{x} \to \infty$,*

$$\mathbb{P}\left(\mathcal{N}_{\kappa\hat{x}}[\xi\hat{x},(1-\alpha)\hat{x}] = 1, V_2^c(0) > \zeta\hat{x}\right) = \mathrm{o}(\mathbb{P}\left(B_2^{\mathrm{r}} > \hat{x}(\phi_2-\rho_2)\right)).$$

Proof Because of independence, the probability in the left-hand side equals

$$\mathbb{P}\left(\mathcal{N}_{\kappa\hat{x}}[\xi\hat{x},(1-\alpha)\hat{x}]=1\right)\mathbb{P}\left(V_2^c(0)>\zeta\hat{x}\right).$$

By conditioning upon $N_U((1-\alpha-\xi)\hat{x})$, we have

$$\mathbb{P}\left(\mathcal{N}_{\kappa\hat{x}}[\xi\hat{x},(1-\alpha)\hat{x}]=1\right)\leq\mathbb{E}\left[N_U((1-\alpha-\xi)\hat{x})\right]\mathbb{P}\left(B_2>\kappa\hat{x}\right).$$

As before the first term is linear in $\hat{x}$ for $\hat{x}\to\infty$. The statement then follows from the fact that $B_2(\cdot)\in\mathcal{R}_{-\nu_2}$ in combination with Theorem 4.1.1. □

The next lemma is the counterpart for on-off processes of the above lemma.

Lemma 4.6.8 (Fluid input) *For any* $c\in(\rho_2,r_2)$, $0<\xi<1-\alpha$, $\zeta,\kappa>0$, *as* $\hat{x}\to\infty$,

$$\begin{aligned}&\mathbb{P}\left(\mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha)\hat{x}]=1,\mathcal{N}_{\kappa\hat{x}}[0,\xi\hat{x}]=0,V_2^c(0)>\zeta\hat{x}\right)\\=\;&o\left(\mathbb{P}\left(A_2^{\mathrm{r}}>\frac{\hat{x}(\phi_2-\rho_2)}{r_2-\rho_2}\right)\right).\end{aligned}$$

Proof The event $\mathcal{N}_{\kappa\hat{x}}[0,\xi\hat{x}]=0$ in conjunction with $\mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha)\hat{x}]=1$ implies that flow 2 has started a long on period at some time t in the interval $[\xi\hat{x},(1-\alpha)\hat{x}]$. Therefore, an upper bound is given by

$$\begin{aligned}&\mathbb{P}\left(\mathcal{N}_{\kappa\hat{x}}[\xi\hat{x},(1-\alpha)\hat{x}]=1,\text{long on period started after time }\xi\hat{x},V_2^c(0)\geq\zeta\hat{x}\right)\\&=\mathbb{P}\left(\#\{j\in\{1,\ldots,N_U((1-\alpha-\xi)\hat{x})\}:A_{2,j}>\kappa\hat{x}\}=1\right)\mathbb{P}\left(V_2^c(0)\geq\zeta\hat{x}\right).\end{aligned}$$

By conditioning upon $N_U((1-\alpha-\xi)\hat{x})$, the first term can be bounded by

$$\mathbb{E}\left[N_U((1-\alpha-\xi)\hat{x})\right]\mathbb{P}\left(A_2>\kappa\hat{x}\right).$$

Combining the fact that $A_2(\cdot)\in\mathcal{R}_{-\nu_2}$ with Theorem 4.1.2 then completes the proof. □

4.7 Case I: instantaneous heavy-tailed input

In this section we consider the case where flow 2 generates instantaneous traffic bursts of regularly varying size. The next theorem shows that then (4.5) in Theorem 4.2.1 holds and that flow 2 satisfies Assumptions 4.5.1-4.5.3. Recall the notation (4.2) and (4.3) as introduced in Section 4.1.

Theorem 4.7.1 *For any* $c>\rho_2$, *and* $\alpha,\gamma>0$,

$$\mathbb{P}\left(T_2^c(\gamma\hat{x})>(1+\alpha)\hat{x}\right)\gtrsim\frac{\rho_2}{c-\rho_2}\mathbb{P}\left(B_2^{\mathrm{r}}>((\phi_2-\rho_2)(1+\alpha)+\gamma)\hat{x}\right),\tag{4.22}$$

$$\mathbb{P}\left(T_2^c(-\gamma\hat{x})>(1-\alpha)\hat{x}\right)\lesssim\frac{\rho_2}{c-\rho_2}\mathbb{P}\left(B_2^{\mathrm{r}}>\psi_I\hat{x}\right),\tag{4.23}$$

with

$$\psi_I := \left((\phi_2 - \rho_2)(1 - \alpha) - \gamma \frac{c + \phi_2 - 2\rho_2}{\phi_2 - \rho_2} \right),$$

and

$$\mathbb{P}\left(T_2^c > \hat{x}\right) \sim \frac{\rho_2}{c - \rho_2} \mathbb{P}\left(B_2^{\mathrm{r}} > \hat{x}(\phi_2 - \rho_2)\right). \tag{4.24}$$

Before giving the formal proof of the above theorem, we first provide an intuitive argument. Consider a queue with service rate ϕ_2 fed by the arrival process of flow 2. In order for the event $T_2^c > \hat{x}$ to occur, the workload must remain positive throughout the interval $[0, \hat{x}]$, given that the initial workload is $V_2^c(0)$. Note that the normal drift in the workload is $\rho_2 - \phi_2 < 0$. Thus, there is a 'deficit' $(\phi_2 - \rho_2)\hat{x}$, which must be compensated for by the initial workload $V_2^c(0)$ plus possibly flow 2 showing above-average activity during the interval $[0, \hat{x}]$.

We claim that the most-likely way for the gap to be filled is by a large initial workload only, i.e., $V_2^c(0) > (\phi_2 - \rho_2)\hat{x}$. This in turn is most probably due to an extremely large burst of flow 2 at some point before time 0, which is consistent with the usual situation for heavy-tailed distributions that a large deviation is caused by just a single exceptional event. Using Theorem 4.1.1, we see that the probability of this event is indeed exactly the right-hand side of (4.24).

Note that it is unlikely for the gap to be filled by flow 2 producing extra traffic during the interval $[0, \hat{x}]$, because this would require a large burst arriving almost immediately after time 0. The probability of this event is negligibly small compared to that of $V_2^c(0) > (\phi_2 - \rho_2)\hat{x}$. A combination of both is even less likely, since this would amount to two rare events occurring simultaneously.

The above arguments will be formalized in the proof below. We first prove that the event $V_2^c(0) > (\phi_2 - \rho_2)\hat{x}$ indeed implies that $T_2^c > \hat{x}$ for large $\hat{x}$, thus obtaining a lower bound for the probability of the latter event. Next we show that for large $\hat{x}$ the event $V_2^c(0) > (\phi_2 - \rho_2)\hat{x}$ is also necessary for $T_2^c > \hat{x}$ to occur, by proving that the probability of all other possible scenarios is negligibly small.

Proof of Theorem 4.7.1 We start with the proof of (4.22). We first prove that for any $\alpha, \gamma, \delta, \theta > 0$, the event

$$T_2^c(\gamma\hat{x}) > (1 + \alpha)\hat{x} \tag{4.25}$$

is implied by the events (i):

$$V_2^c(z^\star) > ((\phi_2 - \rho_2 + \delta)(1 + \alpha) + \gamma + \theta)\hat{x} + cz^\star,$$

where $z^\star$ is the last time before 0 that a burst arrived, and (ii):

$$\sup_{0 \le u \le (1+\alpha)\hat{x}} \{(\rho_2 - \delta)u - A_2(0, u)\} \le \theta\hat{x}.$$

The event (ii) means that for all $u \in [0, (1 + \alpha)\hat{x}]$: $A_2(0, u) \ge (\rho_2 - \delta)u - \theta\hat{x}$. Because of this, for all $u \in [0, (1 + \alpha)\hat{x}]$,

$$V_2^c(0) + A_2(0, u) - \phi_2 u$$

$$\begin{aligned}
&= V_2^c(z^\star) + A_2(z^\star, 0) - B_2^c(z^\star, 0) + A_2(0, u) - \phi_2 u \\
&\geq V_2^c(z^\star) - cz^\star + A_2(0, u) - \phi_2 u \\
&> ((\phi_2 - \rho_2 + \delta)(1+\alpha) + \gamma + \theta)\hat{x} + (\rho_2 - \delta)u - \theta\hat{x} - \phi_2 u \\
&= (\phi_2 - \rho_2 + \delta)((1+\alpha)\hat{x} - u) + \gamma\hat{x} \\
&\geq \gamma\hat{x},
\end{aligned}$$

where the first equality is obtained using (4.1), and the first inequality using the fact that $B_2^c(z^\star, 0) \leq cz^\star$ (and the trivial inequality $A_2(z^\star, 0) \geq 0$). Hence,

$$\inf\{u \geq 0 : V_2^c(0) + A_2(0, u) - \phi_2 u \leq \gamma\hat{x}\} > (1+\alpha)\hat{x},$$

which gives (4.25).

Using independence of $V_2^c(z^\star)$ and $A_2(0, u)$,

$$\begin{aligned}
&\mathbb{P}\left(T_2^c(\gamma\hat{x}) > (1+\alpha)\hat{x}\right) \\
\geq\ &\mathbb{P}\left(V_2^c(z^\star) > ((\phi_2 - \rho_2 + \delta)(1+\alpha) + \gamma + \theta)\hat{x} + cz^\star\right) \\
&\times \mathbb{P}\left(\sup_{0 \leq u \leq (1+\alpha)\hat{x}} \{(\rho_2 - \delta)u - A_2(0, u)\} \leq \theta\hat{x}\right).
\end{aligned}$$

Now observe that $V_2^c(0) > 0$ implies $V_2^c(0) = V_2^c(z^\star) - cz^\star$, so that

$$\begin{aligned}
&\mathbb{P}\left(V_2^c(z^\star) > ((\phi_2 - \rho_2 + \delta)(1+\alpha) + \gamma + \theta)\hat{x} + cz^\star\right) \\
\geq\ &\mathbb{P}\left(V_2^c(0) > ((\phi_2 - \rho_2 + \delta)(1+\alpha) + \gamma + \theta)\hat{x}\right) \\
\sim\ &\frac{\rho_2}{c - \rho_2}\mathbb{P}\left(B_2^{\mathrm{r}} > ((\phi_2 - \rho_2 + \delta)(1+\alpha) + \gamma + \theta)\hat{x})\right),
\end{aligned}$$

where the last term is due to Theorem 4.1.1. Also, for all $\alpha, \delta, \theta > 0$,

$$\begin{aligned}
&\mathbb{P}\left(\sup_{0 \leq u \leq (1+\alpha)\hat{x}} \{(\rho_2 - \delta)u - A_2(0, u)\} \leq \theta\hat{x}\right) \\
\geq\ &\mathbb{P}\left(\sup_{u \geq 0}\{(\rho_2 - \delta)u - A_2(0, u)\} \leq \theta\hat{x}\right) \to 1,
\end{aligned}$$

as $\hat{x} \to \infty$, since $\mathbb{E}A_2(0, u) = \rho_2 u$. Thus, for all $\alpha, \gamma, \delta, \theta > 0$,

$$\mathbb{P}\left(T_2^c(\gamma\hat{x}) > (1+\alpha)\hat{x}\right) \gtrsim \frac{\rho_2}{c - \rho_2}\mathbb{P}\left(B_2^{\mathrm{r}} > ((\phi_2 - \rho_2 + \delta)(1+\alpha) + \gamma + \theta)\hat{x}\right).$$

Letting $\delta \downarrow 0$, $\theta \downarrow 0$, using the fact that $B_2^{\mathrm{r}}(\cdot) \in \mathcal{IR}$, (4.22) follows.

We now turn to the proof of (4.23). By partitioning, we obtain for any $\alpha, \gamma, \zeta, \theta, \kappa > 0$, $w \geq 0$,

$$\begin{aligned}
&\mathbb{P}\left(T_2^c(-\gamma\hat{x}) > (1-\alpha)\hat{x}\right) \\
&= \mathbb{P}\left(T_2^c(-\gamma\hat{x}) > (1-\alpha)\hat{x}, V_2^c(w) > ((\phi_2 - \rho_2)(1-\alpha) - \gamma - \theta)\hat{x} - (c - \rho_2)w\right) \\
&+ \mathbb{P}\left(T_2^c(-\gamma\hat{x}) > (1-\alpha)\hat{x}, V_2^c(w) \leq ((\phi_2 - \rho_2)(1-\alpha) - \gamma - \theta)\hat{x} - (c - \rho_2)w,\right.
\end{aligned}$$

$$\begin{aligned}
&\mathcal{N}_{\kappa\hat{x}}[0,w] \leq 1, \mathcal{N}_{\kappa\hat{x}}[w,(1-\alpha)\hat{x}] = 0) \\
&+\mathbb{P}\left(T_2^c(-\gamma\hat{x}) > (1-\alpha)\hat{x}, V_2^c(w) \leq ((\phi_2-\rho_2)(1-\alpha)-\gamma-\theta)\hat{x}-(c-\rho_2)w,\right. \\
&\quad \mathcal{N}_{\kappa\hat{x}}[0,w] = 0, \mathcal{N}_{\kappa\hat{x}}[w,(1-\alpha)\hat{x}] = 1, V_2^c(0) \leq \zeta\hat{x}) \\
&+\mathbb{P}\left(T_2^c(-\gamma\hat{x}) > (1-\alpha)\hat{x}, V_2^c(w) \leq ((\phi_2-\rho_2)(1-\alpha)-\gamma-\theta)\hat{x}-(c-\rho_2)w,\right. \\
&\quad \mathcal{N}_{\kappa\hat{x}}[0,w] = 0, \mathcal{N}_{\kappa\hat{x}}[w,(1-\alpha)\hat{x}] = 1, V_2^c(0) > \zeta\hat{x}) \\
&+\mathbb{P}\left(T_2^c(-\gamma\hat{x}) > (1-\alpha)\hat{x}, V_2^c(w) \leq ((\phi_2-\rho_2)(1-\alpha)-\gamma-\theta)\hat{x}-(c-\rho_2)w,\right. \\
&\quad \left.\mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha)\hat{x}] \geq 2\right),
\end{aligned}$$

which is obviously upper bounded by $I_a + I_b + I_c + I_d + I_e$, where

$$\begin{aligned}
I_a &:= \mathbb{P}\left(V_2^c(w) > ((\phi_2-\rho_2)(1-\alpha)-\gamma-\theta)\hat{x}-(c-\rho_2)w\right), \\
I_b &:= \mathbb{P}\left(T_2^c(-\gamma\hat{x}) > (1-\alpha)\hat{x}, \mathcal{N}_{\kappa\hat{x}}[w,(1-\alpha)\hat{x}] = 0,\right. \\
&\qquad \left. V_2^c(w) \leq ((\phi_2-\rho_2)(1-\alpha)-\gamma-\theta)\hat{x}-(c-\rho_2)w\right), \\
I_c &:= \mathbb{P}\left(T_2^c(-\gamma\hat{x}) > (1-\alpha)\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0,w] = 0, V_2^c(0) \leq \zeta\hat{x}\right), \\
I_d &:= \mathbb{P}\left(\mathcal{N}_{\kappa\hat{x}}[w,(1-\alpha)\hat{x}] = 1, V_2^c(0) > \zeta\hat{x}\right), \\
I_e &:= \mathbb{P}\left(\mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha)\hat{x}] \geq 2\right).
\end{aligned}$$

Take $w = \xi\hat{x}$, with

$$\xi := \frac{\gamma+\zeta+\theta}{\phi_2-\rho_2} < 1-\alpha.$$

We first concentrate on the event $T_2^c(-\gamma\hat{x}) > (1-\alpha)\hat{x}$. It is equivalent to

$$\inf_{0\leq u\leq(1-\alpha)\hat{x}} \{V_2^c(0) + A_2(0,u) - \phi_2 u\} > -\gamma\hat{x}.$$

Observe that

$$\inf_{0\leq u\leq(1-\alpha)\hat{x}} \{V_2^c(0) + A_2(0,u) - \phi_2 u\} \leq V_2^c(0) + \inf_{0\leq u\leq w} \{A_2(0,u) - \phi_2 u\} \quad (4.26)$$

and

$$\begin{aligned}
&\inf_{0\leq u\leq(1-\alpha)\hat{x}} \{V_2^c(0) + A_2(0,u) - \phi_2 u\} \\
\leq\ & V_2^c(0) + \inf_{w\leq u\leq(1-\alpha)\hat{x}} \{A_2(0,u) - \phi_2 u\} \\
=\ & V_2^c(0) + \inf_{w\leq u\leq(1-\alpha)\hat{x}} \{A_2(0,w) - \phi_2 w + A_2(w,u) - \phi_2(u-w)\} \\
=\ & V_2^c(0) + A_2(0,w) - \phi_2 w + \inf_{w\leq u\leq(1-\alpha)\hat{x}} \{A_2(w,u) - \phi_2(u-w)\} \\
\leq\ & V_2^c(w) + (c-\phi_2)w + \inf_{w\leq u\leq(1-\alpha)\hat{x}} \{A_2(w,u) - \phi_2(u-w)\}. \qquad (4.27)
\end{aligned}$$

Consider term I_a. Using Theorem 4.1.1, I_a is equal to

$$\begin{aligned}
&\mathbb{P}\left(V_2^c > ((\phi_2-\rho_2)(1-\alpha)-\gamma-\theta-(c-\rho_2)\xi)\hat{x}\right) \\
\sim\ & \frac{\rho_2}{c-\rho_2}\mathbb{P}\left(B_2^{\mathrm{r}} > ((\phi_2-\rho_2)(1-\alpha)-\gamma-\theta-(c-\rho_2)\xi)\hat{x}\right)
\end{aligned}$$

$$= \frac{\rho_2}{c-\rho_2}\mathbb{P}\left(B_2^{\mathrm{r}} > \left((\phi_2-\rho_2)(1-\alpha)-\gamma-\theta-\frac{(c-\rho_2)(\gamma+\zeta+\theta)}{\phi_2-\rho_2}\right)\hat{x}\right).$$

Next, consider term I_b. Using (4.27), the event $T_2^c(-\gamma\hat{x}) > (1-\alpha)\hat{x}$ implies

$$\inf_{w\leq u\leq(1-\alpha)\hat{x}}\{A_2(w,u)-\phi_2(u-w)\} > -V_2^c(w)-(c-\phi_2)w-\gamma\hat{x},$$

so that an upper bound for I_b is given by

$$\mathbb{P}\Big(\inf_{w\leq u\leq(1-\alpha)\hat{x}}\{A_2(w,u)-\phi_2(u-w)\} > -V_2^c(w)-(c-\phi_2)w-\gamma\hat{x},$$
$$V_2^c(w)\leq((\phi_2-\rho_2)(1-\alpha)-\gamma-\theta)\hat{x}-(c-\rho_2)w, \mathcal{N}_{\kappa\hat{x}}[w,(1-\alpha)\hat{x}]=0\Big)$$
$$\leq \mathbb{P}\Big(\inf_{w\leq u\leq(1-\alpha)\hat{x}}\{A_2(w,u)-\phi_2(u-w)\} > \theta\hat{x}-(\phi_2-\rho_2)((1-\alpha)\hat{x}-w),$$
$$\mathcal{N}_{\kappa\hat{x}}[w,(1-\alpha)\hat{x}]=0\Big)$$
$$= \mathbb{P}\Big(\inf_{0\leq u\leq(1-\alpha)\hat{x}-w}\{A_2(0,u)-\phi_2 u\} > \theta\hat{x}-(\phi_2-\rho_2)((1-\alpha)\hat{x}-w),$$
$$\mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha)\hat{x}-w]=0\Big)$$
$$= \mathbb{P}\Big(\inf_{0\leq u\leq(1-\alpha-\xi)\hat{x}}\{A_2(0,u)-\phi_2 u\} > (\theta-(\phi_2-\rho_2)(1-\alpha-\xi))\hat{x},$$
$$\mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha-\xi)\hat{x}]=0\Big)$$
$$= \mathbb{P}\left(T_2(\theta-(\phi_2-\rho_2)(1-\alpha-\xi))\hat{x}) > (1-\alpha-\xi)\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha-\xi)\hat{x}]=0\right).$$

Finally, consider term I_c. Using (4.26), the event $T_2^c(-\gamma\hat{x}) > (1-\alpha)\hat{x}$ implies (recall that $w < (1-\alpha)\hat{x}$)

$$\inf_{0\leq u\leq w}\{A_2(0,u)-\phi_2 u\} > -V_2^c(0)-\gamma\hat{x},$$

so that an upper bound for I_c is given by

$$\mathbb{P}\left(\inf_{0\leq u\leq w}\{A_2(0,u)-\phi_2 u\} > -V_2^c(0)-\gamma\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0,w]=0, V_2^c(0)\leq\zeta\hat{x}\right)$$
$$\leq \mathbb{P}\left(\inf_{0\leq u\leq w}\{A_2(0,u)-\phi_2 u\} > -(\gamma+\zeta)\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0,w]=0\right)$$
$$= \mathbb{P}\left(\inf_{0\leq u\leq \xi\hat{x}}\{A_2(0,u)-\phi_2 u\} > (\theta-(\phi_2-\rho_2)\xi)\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0,\xi\hat{x}]=0\right)$$
$$= \mathbb{P}\left(T_2((\theta-(\phi_2-\rho_2)\xi)\hat{x}) > \xi\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0,\xi\hat{x}]=0\right).$$

Now taking $\eta = 1-\alpha-\xi$ in Lemma 4.6.2 for I_b, taking $\eta=\xi$ in Lemma 4.6.2 for I_c, using Lemma 4.6.7 for I_d, and using Lemma 4.6.4 for I_e, we obtain

$$\mathbb{P}\left(T_2^c(-\gamma\hat{x}) > (1-\alpha)\hat{x}\right) \lesssim$$

$$\frac{\rho_2}{c-\rho_2}\mathbb{P}\left(B_2^{\mathrm{r}} > ((\phi_2-\rho_2)(1-\alpha)-\gamma-\theta-\frac{(c-\rho_2)(\gamma+\zeta+\theta)}{\phi_2-\rho_2})\hat{x}\right).$$

Letting $\zeta \downarrow 0$, $\theta \downarrow 0$, using the fact that $B_2^{\mathrm{r}}(\cdot) \in \mathcal{IR}$, (4.23) follows.

Finally, note that (4.24) follows from (4.22) and (4.23) by letting $\alpha \downarrow 0$, $\gamma \downarrow 0$, and using again the fact that $B_2^{\mathrm{r}}(\cdot) \in \mathcal{IR}$. □

4.8 Case II-A: fluid heavy-tailed input with $r_2 < 1 - \rho_1$

We now consider the case where flow 2 generates traffic according to an on-off process with peak rate $r_2 < 1-\rho_1$. The next theorem shows that flow 2 satisfies Assumptions 4.5.1-4.5.3 and that (4.6) in Theorem 4.2.1 holds.

Theorem 4.8.1 *For any $\alpha, \gamma > 0$,*

$$\mathbb{P}(T_2 > (1+\alpha)\hat{x}) \gtrsim (1-p_2)\mathbb{P}\left(A_2^{\mathrm{r}} > \frac{\phi_2-\rho_2}{r_2-\rho_2}(1+\alpha)\hat{x}\right), \tag{4.28}$$

$$\mathbb{P}(T_2(-\gamma\hat{x}) > (1-\alpha)\hat{x}) \lesssim (1-p_2)\mathbb{P}(A_2^{\mathrm{r}} > \psi_{II-A}\hat{x}), \tag{4.29}$$

with

$$\psi_{II-A} := \left(\frac{(\phi_2-\rho_2)(1-\alpha)-\gamma}{r_2-\rho_2} - \frac{\gamma}{\phi_2-\rho_2}\right),$$

and

$$\mathbb{P}(T_2 > \hat{x}) \sim (1-p_2)\mathbb{P}\left(A_2^{\mathrm{r}} > \frac{\phi_2-\rho_2}{r_2-\rho_2}\hat{x}\right). \tag{4.30}$$

Before giving the formal proof of the above theorem, we first provide an intuitive argument. Consider a queue with service rate ϕ_2 fed by the arrival process of flow 2. In order for the event $T_2 > \hat{x}$ to occur, the workload must remain positive throughout the interval $[0, \hat{x}]$, given that the initial workload is 0. Note that the normal drift in the workload is $\rho_2 - \phi_2 < 0$. Thus, there is a 'deficit' $(\phi_2-\rho_2)\hat{x}$, which must be made up for by flow 2 showing above-average activity during the interval $[0, \hat{x}]$.

We claim that the most-likely way for the gap to be filled is by a single long on period of flow 2. This long on period covers the entire interval $[0, v]$, with $v := ((\phi_2-\rho_2)\hat{x})/(r_2-\rho_2)$. When on, flow 2 generates above-average traffic at rate $r_2 > \rho_2$, so this event (call it $E(\hat{x})$) makes up for the entire deficit. This is consistent with the usual situation for heavy-tailed distributions that a large deviation is caused by just a single exceptional event. Observe that the probability of this event is indeed exactly the right-hand side of (4.30). Note that it is unlikely for the gap to be filled by several long on periods, since the probability of this happening is an order of magnitude smaller.

The above arguments will be formalized in the proof below. We first prove that the event $E(\hat{x})$ combined with normal behavior during the interval $(v, \hat{x}]$, indeed implies that $T_2 > \hat{x}$ for large $\hat{x}$, thus obtaining a lower bound for the probability of the latter event. Next we show that for large $\hat{x}$ the event $E(\hat{x})$ is also necessary for $T_2 > \hat{x}$ to occur, by proving that the probability of all other possible scenarios is negligibly small.

Proof of Theorem 4.8.1 We first prove that for any $\alpha, \delta, \theta > 0$, the event

$$T_2 > (1+\alpha)\hat{x} \tag{4.31}$$

is implied by the event $E(\hat{x})$ that flow 2 is on at time 0 and turns off again at time $\tilde{v} > \tilde{\tau}\hat{x}$, with

$$\tilde{\tau} := \frac{(\phi_2 - \rho_2 + \delta)(1+\alpha) + \theta}{r_2 - \rho_2 + \delta},$$

combined with the event

$$\sup_{\tilde{v} \leq u \leq (1+\alpha)\hat{x}} \{(\rho_2 - \delta)(u - \tilde{v}) - A_2(\tilde{v}, u)\} \leq \theta\hat{x}.$$

Note that $\tilde{v}$ differs slightly from time v as defined in the intuition above. The second event means that for all $u \in [\tilde{v}, (1+\alpha)\hat{x}]$: $A_2(\tilde{v}, u) \geq (\rho_2 - \delta)(u - \tilde{v}) - \theta\hat{x}$. We distinguish between two cases. (i) $0 \leq u \leq \tilde{v}$. Then

$$A_2(0,u) - \phi_2 u = r_2 u - \phi_2 u \geq 0.$$

(ii) $\tilde{v} \leq u \leq (1+\alpha)\hat{x}$. Then

$$\begin{aligned} A_2(0,u) - \phi_2 u &= A_2(0,\tilde{v}) + A_2(\tilde{v}, u) - \phi_2 u \\ &\geq r_2\tilde{v} + (\rho_2 - \delta)(u - \tilde{v}) - \theta\hat{x} - \phi_2 u \\ &= (r_2 - \rho_2 + \delta)\tilde{v} - (\phi_2 - \rho_2 + \delta)u - \theta\hat{x} \\ &> ((\phi_2 - \rho_2 + \delta)(1+\alpha) + \theta)\hat{x} - (\phi_2 - \rho_2 + \delta)u - \theta\hat{x} \\ &\geq (\phi_2 - \rho_2 + \delta)(1+\alpha)\hat{x} - (\phi_2 - \rho_2 + \delta)(1+\alpha)\hat{x} \\ &= 0. \end{aligned}$$

So, $\inf\{u \geq 0 : A_2(0,u) - \phi_2 u \leq 0\} > (1+\alpha)\hat{x}$, which gives (4.31).

Hence, due to the independence of the two events above, using Lemma 4.6.6, for any $\alpha, \delta, \theta > 0$,

$$\begin{aligned} &\mathbb{P}(T_2 > (1+\alpha)\hat{x}) \\ \geq\ &\mathbb{P}(E(\hat{x}))\,\mathbb{P}\left(\sup_{\tilde{v} \leq u \leq (1+\alpha)\hat{x}} \{(\rho_2 - \delta)(u - \tilde{v}) - A_2(\tilde{v}, u)\} \leq \theta\hat{x}\right) \\ \gtrsim\ &\mathbb{P}(E(\hat{x})) \\ =\ &(1 - p_2)\mathbb{P}\left(A_2^{\mathrm{r}} > \frac{(\phi_2 - \rho_2 + \delta)(1+\alpha) + \theta}{r_2 - \rho_2 + \delta}\hat{x}\right). \end{aligned}$$

Letting $\delta \downarrow 0$, $\theta \downarrow 0$, using the fact that $A_2^{\mathrm{r}}(\cdot) \in \mathcal{IR}$, (4.28) follows.

We now turn to the proof of (4.29). Note that we use $\hat{v}$ which is again somewhat different from v as given in the intuition above. By partitioning, we obtain for any $\alpha, \gamma, \theta, \kappa > 0$, $\hat{v} \geq w \geq 0$,

$$\mathbb{P}(T_2(-\gamma\hat{x}) > (1-\alpha)\hat{x})$$

$$
\begin{aligned}
&= \mathbb{P}\left(T_2(-\gamma\hat{x}) > (1-\alpha)\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha)\hat{x}] \leq 1, \mathcal{N}_{\kappa\hat{x}}[0,w] \geq 1,\right.\\
&\qquad \left.\mathcal{N}_{\kappa\hat{x}}[\hat{v},(1-\alpha)\hat{x}] \geq 1\right)\\
&\quad + \mathbb{P}\left(T_2(-\gamma\hat{x}) > (1-\alpha)\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha)\hat{x}] \leq 1, \mathcal{N}_{\kappa\hat{x}}[0,w] \geq 1,\right.\\
&\qquad \left.\mathcal{N}_{\kappa\hat{x}}[\hat{v},(1-\alpha)\hat{x}] = 0\right)\\
&\quad + \mathbb{P}\left(T_2(-\gamma\hat{x}) > (1-\alpha)\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha)\hat{x}] \leq 1, \mathcal{N}_{\kappa\hat{x}}[0,w] = 0\right)\\
&\quad + \mathbb{P}\left(T_2(-\gamma\hat{x}) > (1-\alpha)\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha)\hat{x}] \geq 2\right),
\end{aligned}
$$

which is clearly upper bounded by $II\text{–}A_a + II\text{–}A_b + II\text{–}A_c + II\text{–}A_d$, with

$$
\begin{aligned}
II\text{–}A_a &:= \mathbb{P}\left(\mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha)\hat{x}] \leq 1, \mathcal{N}_{\kappa\hat{x}}[0,w] \geq 1, \mathcal{N}_{\kappa\hat{x}}[\hat{v},(1-\alpha)\hat{x}] \geq 1\right),\\
II\text{–}A_b &:= \mathbb{P}\left(T_2(-\gamma\hat{x}) > (1-\alpha)\hat{x}, \mathcal{N}_{\kappa\hat{x}}[\hat{v},(1-\alpha)\hat{x}] = 0\right),\\
II\text{–}A_c &:= \mathbb{P}\left(T_2(-\gamma\hat{x}) > (1-\alpha)\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0,w] = 0\right),\\
II\text{–}A_d &:= \mathbb{P}\left(\mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha)\hat{x}] \geq 2\right).
\end{aligned}
$$

Take $\hat{v} = \hat{\tau}\hat{x}$ and $w = \xi\hat{x}$, where

$$\hat{\tau} := \frac{(\phi_2 - \rho_2)(1-\alpha) - \gamma - \theta}{r_2 - \rho_2} < 1 - \alpha,$$

and ξ is redefined as

$$\xi := \frac{\gamma + \theta}{\phi_2 - \rho_2} < \hat{\tau}.$$

We remark that $\hat{\tau}$ plays a similar role as $\tilde{\tau}$ in the first part of this proof (proof of (4.28)) and that ξ plays a similar role as ξ defined in the proof of Theorem 4.7.1. Now consider term $II\text{–}A_a$. For the relevant events to occur, flow 2 must be on during the entire interval $[w, \hat{v}]$, so that

$$
\begin{aligned}
II\text{–}A_a &\leq (1-p_2)\mathbb{P}\left(A_2^{\mathrm{r}} > \hat{v} - w\right)\\
&= (1-p_2)\mathbb{P}\left(A_2^{\mathrm{r}} > \left(\frac{(\phi_2 - \rho_2)(1-\alpha) - \gamma - \theta}{r_2 - \rho_2} - \frac{\gamma + \theta}{\phi_2 - \rho_2}\right)\hat{x}\right).
\end{aligned}
$$

Next, consider term $II\text{–}A_b$. The event $T_2(-\gamma\hat{x}) > (1-\alpha)\hat{x}$ means that

$$\inf_{0 \leq u \leq (1-\alpha)\hat{x}} \{A_2(0,u) - \phi_2 u\} > -\gamma\hat{x}.$$

Now observe that

$$
\begin{aligned}
&\inf_{0 \leq u \leq (1-\alpha)\hat{x}} \{A_2(0,u) - \phi_2 u\}\\
\leq\ &\inf_{\hat{v} \leq u \leq (1-\alpha)\hat{x}} \{A_2(0,u) - \phi_2 u\}\\
=\ &\inf_{\hat{v} \leq u \leq (1-\alpha)\hat{x}} \{A_2(0,\hat{v}) - \phi_2\hat{v} + A_2(\hat{v},u) - \phi_2(u - \hat{v})\}\\
=\ &A_2(0,\hat{v}) - \phi_2\hat{v} + \inf_{\hat{v} \leq u \leq (1-\alpha)\hat{x}} \{A_2(\hat{v},u) - \phi_2(u - \hat{v})\}\\
\leq\ &(r_2 - \phi_2)\hat{v} + \inf_{\hat{v} \leq u \leq (1-\alpha)\hat{x}} \{A_2(\hat{v},u) - \phi_2(u - \hat{v})\}.
\end{aligned}
$$

Thus, the event $T_2(-\gamma\hat{x}) > (1-\alpha)\hat{x}$ implies

$$\inf_{\hat{v}\le u\le(1-\alpha)\hat{x}} \{A_2(\hat{v},u) - \phi_2(u-\hat{v})\} > -(r_2-\phi_2)\hat{v} - \gamma\hat{x},$$

so that an upper bound for II–A_b is given by

$$\begin{aligned}
&\mathbb{P}\Big(\inf_{\hat{v}\le u\le(1-\alpha)\hat{x}} \{A_2(\hat{v},u) - \phi_2(u-\hat{v})\} > -(r_2-\phi_2)\hat{v} - \gamma\hat{x},\\
&\qquad \mathcal{N}_{\kappa\hat{x}}[\hat{v},(1-\alpha)\hat{x}] = 0\Big)\\
= \;& \mathbb{P}\Big(\inf_{0\le u\le(1-\alpha)\hat{x}-\hat{v}} \{A_2(0,u) - \phi_2 u\} > -(r_2-\phi_2)\hat{v} - \gamma\hat{x},\\
&\qquad \mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha)\hat{x}-\hat{v}] = 0\Big)\\
= \;& \mathbb{P}\Big(\inf_{0\le u\le(1-\alpha-\hat{\tau})\hat{x}} \{A_2(0,u) - \phi_2 u\} > (\theta-(\phi_2-\rho_2)(1-\alpha-\hat{\tau}))\hat{x},\\
&\qquad \mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha-\hat{\tau})\hat{x}] = 0\Big)\\
= \;& \mathbb{P}\Big(T_2((\theta-(\phi_2-\rho_2)(1-\alpha-\hat{\tau}))\hat{x}) > (1-\alpha-\hat{\tau})\hat{x},\\
&\qquad \mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha-\hat{\tau})\hat{x}] = 0\Big),
\end{aligned}$$

where the second equality follows from the fact that $-(r_2-\phi_2)\hat{v} = -(r_2-\rho_2+\rho_2-\phi_2)\hat{\tau}\hat{x}$ and, using the definition of $\hat{\tau}$, $(r_2-\rho_2)\hat{\tau} = (\phi_2-\rho_2)(1-\alpha)-\gamma-\theta$. Finally, consider term II–A_c, i.e.,

$$\begin{aligned}
II\text{–}A_c &\le \mathbb{P}\left(T_2(-\gamma\hat{x}) > w, \mathcal{N}_{\kappa\hat{x}}[0,w] = 0\right)\\
&= \mathbb{P}\left(T_2((\theta-(\phi_2-\rho_2)\xi)\hat{x}) > \xi\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0,\xi\hat{x}] = 0\right).
\end{aligned}$$

Thus, taking $\eta = 1-\alpha-\hat{\tau}$ in Lemma 4.6.3 for II–A_b, taking $\eta = \xi$ in Lemma 4.6.3 for II–A_c, and using Lemma 4.6.5 for II–A_d, we obtain

$$\begin{aligned}
&\mathbb{P}\left(T_2(-\gamma\hat{x}) > (1-\alpha)\hat{x}\right)\\
\lesssim\;& (1-p_2)\mathbb{P}\left(A_2^{\mathrm{r}} > \Big(\frac{(\phi_2-\rho_2)(1-\alpha)-\gamma-\theta}{r_2-\rho_2} - \frac{\gamma+\theta}{\phi_2-\rho_2}\Big)\hat{x}\right).
\end{aligned}$$

Letting $\theta \downarrow 0$, using the fact that $A_2^{\mathrm{r}}(\cdot) \in \mathcal{IR}$, (4.29) follows.

Letting $\alpha \downarrow 0$, $\gamma \downarrow 0$ and using that $A_2^{\mathrm{r}}(\cdot) \in \mathcal{IR}$, it is easily seen that (4.30) follows from (4.28) and (4.29). □

4.9 Case II-B: fluid heavy-tailed input with $r_2 > 1-\rho_1$

We now consider the case where flow 2 generates traffic according to an on-off process with peak rate $r_2 > 1-\rho_1$. The next theorem shows that flow 2 satisfies Assumptions 4.5.1-4.5.3 and that (4.7) in Theorem 4.2.1 holds.

Theorem 4.9.1 *For any $c \in (\rho_2, r_2)$, $\alpha, \gamma > 0$,*

$$\mathbb{P}\left(T_2^c(\gamma\hat{x}) > (1-\alpha)\hat{x}\right) \gtrsim p_2 \frac{\rho_2}{c-\rho_2} \mathbb{P}\left(A_2^{\mathrm{r}} > \psi_{II\text{-}B1}\hat{x}\right), \tag{4.32}$$

with

$$\psi_{II\text{-}B1} := \left(\frac{(\phi_2-\rho_2)(1+\alpha)}{r_2-\rho_2} + \frac{\gamma}{r_2-c}\right),$$

$$\mathbb{P}\left(T_2^c(-\gamma\hat{x}) > (1-\alpha)\hat{x}\right) \lesssim p_2 \frac{\rho_2}{c-\rho_2} \mathbb{P}\left(A_2^{\mathrm{r}} > \psi_{II\text{-}B2}\hat{x}\right), \tag{4.33}$$

with

$$\psi_{II\text{-}B2} := \left(\frac{(\phi_2-\rho_2)(1-\alpha)}{r_2-\rho_2} + \frac{\gamma(2r_2-\phi_2-c)}{(r_2-\rho_2)(\phi_2-\rho_2)}\right),$$

and

$$\mathbb{P}\left(T_2^c > \hat{x}\right) \sim p_2 \frac{\rho_2}{c-\rho_2} \mathbb{P}\left(A_2^{\mathrm{r}} > \frac{\phi_2-\rho_2}{r_2-\rho_2}\hat{x}\right). \tag{4.34}$$

Before giving the formal proof of the above theorem, we first provide an intuitive argument. Consider a queue with service rate ϕ_2 fed by the arrival process of flow 2. In order for the event $T_2^c > \hat{x}$ to occur, the workload must remain positive throughout the interval $[0, \hat{x}]$, given that the initial workload is $V_2^c(0)$. Note that the normal drift in the workload is $\rho_2 - \phi_2 < 0$. Thus, there is a 'deficit' $(\phi_2 - \rho_2)\hat{x}$, which must be compensated for by the initial workload $V_2^c(0)$ plus possibly flow 2 showing above-average activity during the interval $[0, \hat{x}]$.

As before, we claim that the most-likely way for the gap to be filled is by an extremely long on period of flow 2 which started at some point before time 0. Unfortunately, it is harder to pin down exactly how long that on period must last, since it depends on when it started. No matter when the on period started however, it turns out that we must always have $V_2^c(v) > (r_2 - c)v$, with $v := ((\phi_2 - \rho_2)\hat{x})/(r_2 - \rho_2)$. Using Theorem 4.1.2, we see that the probability of this event is indeed exactly the right-hand side of (4.34).

The above arguments will be formalized in the proof below. We first prove that a sufficiently long on period, combined with normal behavior of flow 2 in the interval $[v, \hat{x}]$, implies that $T_2^c > \hat{x}$ for large $\hat{x}$, thus obtaining a lower bound for the probability of the latter event. Next we show that for large $\hat{x}$ the event $V_2^c(v) > (r_2 - c)v$ is indeed necessary for $T_2^c > \hat{x}$ to occur, by proving that the probability of all other possible scenarios is negligibly small.

Proof of Theorem 4.9.1 We start with the proof of (4.32). Consider the following two events. (i) Flow 2 turns on at time $-y$ with $y \in (\gamma\hat{x}/(r_2 - c), \infty)$ and turns off at time $\tilde{v} > -y$ such that

$$(y+\tilde{v})(r_2-\rho_2+\delta) - y(c-\rho_2+\delta) > ((\phi_2-\rho_2+\delta)(1+\alpha)+\gamma+\theta)\hat{x}, \tag{4.35}$$

where $y + \tilde{v}$ is the length of the on period that started at $-y$. Note that $\tilde{v}$ differs slightly from v as defined above (and plays a similar role as $\tilde{v}$ in the proof of

Theorem 4.8.1. (ii) The event

$$\sup_{\tilde{v}\le u\le(1+\alpha)\hat{x}} \{(\rho_2-\delta)(u-\tilde{v})-A_2(\tilde{v},u)\}\le\theta\hat{x},$$

which implies that $\forall u\in[\tilde{v},(1+\alpha)\hat{x}]: A_2(\tilde{v},u)\ge(\rho_2-\delta)(u-\tilde{v})-\theta\hat{x}$.

We now prove that for any $\alpha,\gamma,\delta,\theta>0$, the above events (i) and (ii) imply

$$T_2^c(\gamma\hat{x})>(1+\alpha)\hat{x}. \tag{4.36}$$

Applying (4.1) to $V_2^c(0)$ and observing that $B_2^c(s,t)\le c(t-s)$ for all $s<t$, we have

$$V_2^c(0)+A_2(0,u)-\phi_2 u\ge A_2(-y,u)-cy-\phi_2 u.$$

We split the interval $[0,(1-\alpha)\hat{x}]$ and consider two regimes for u. For both regimes we show that

$$V_2^c(0)+A_2(0,u)-\phi_2 u\ge\gamma\hat{x}.$$

First consider $0\le u\le\tilde{v}$. Because of (i), flow 2 is still on at time u, and thus,

$$\begin{aligned}V_2^c(0)+A_2(0,u)-\phi_2 u &\ge r_2(y+u)-cy-\phi_2 u\\ &\ge (r_2-c)y\\ &> \gamma\hat{x},\end{aligned}$$

where the strict inequality is due to (i), where $y\in(\gamma\hat{x}/(r_2-c),\infty)$. The second regime for u is then $\tilde{v}\le u\le(1+\alpha)\hat{x}$. Because of (i), flow 2 remains on until time $\tilde{v}$, and thus,

$$\begin{aligned}&V_2^c(0)+A_2(0,u)-\phi_2 u\\ \ge\;& r_2(y+\tilde{v})+A_2(\tilde{v},u)-cy-\phi_2 u\\ \ge\;& r_2(y+\tilde{v})+(\rho_2-\delta)(u-\tilde{v})-\theta\hat{x}-cy-\phi_2 u\\ >\;& ((\phi_2-\rho_2+\delta)(1+\alpha)+\theta+\gamma)\hat{x}-u(\phi_2-\rho_2+\delta)-\theta\hat{x}\\ \ge\;& \gamma\hat{x},\end{aligned}$$

where the second inequality follows from (ii), the strict inequality from (4.35) in (i), and the last inequality from $u\le(1+\alpha)\hat{x}$. So, combining the two regimes for u we obtain:

$$\inf\{u\ge0: V_2^c(0)+A_2(0,u)-\phi_2 u\le\gamma\hat{x}\}>(1+\alpha)\hat{x},$$

which gives (4.36).

We need the following additional notation. We define Z_2 to be a random variable whose distribution is the distribution of the length of the on period starting at time $-y$. We use (4.35) to define z_y as a realization of Z_2, which depends on y, i.e.,

$$z_y:=\frac{((\phi_2-\rho_2+\delta)(1+\alpha)+\gamma+\theta)\hat{x}+y(c-\rho_2+\delta)}{r_2-\rho_2+\delta}.$$

Because the events *(I)* and *(II)* are independent, we have for any $\alpha, \gamma, \delta, \theta > 0$,

$$\begin{aligned}
& \mathbb{P}\left(T_2^c(\gamma\hat{x}) \geq (1+\alpha)\hat{x}\right) \\
\geq\ & \mathbb{P}\left(\text{turns on at } -y, \text{ turns off at } \tilde{v}, (4.35), y \in \left(\frac{\gamma\hat{x}}{r_2 - c}, \infty\right)\right) \\
& \times\mathbb{P}\left(\sup_{\tilde{v} \leq u \leq (1+\alpha)\hat{x}} \{(\rho_2 - \delta)(u - \tilde{v}) - A_2(\tilde{v}, u)\} \leq \theta\hat{x}\right) \\
\sim\ & \mathbb{P}\left(\text{turns on at } -y, y \in \left(\frac{\gamma\hat{x}}{r_2 - c}, \infty\right), Z_2 > z_y\right) \qquad (4.37) \\
=\ & \frac{1}{\alpha_2 + \frac{1}{\lambda_2}} \int_{\frac{\gamma\hat{x}}{r_2 - c}}^{\infty} \mathbb{P}\left(A_2 > z_y\right) \mathrm{d}y, \qquad (4.38)
\end{aligned}$$

where (4.37) is due to Lemma 4.6.6. Defining

$$\beta := \frac{\gamma\hat{x}}{r_2 - c} + \frac{((\phi_2 - \rho_2 + \delta)(1+\alpha) + \theta)\hat{x}}{r_2 - \rho_2 + \delta},$$

and changing the integration variable, (4.38) can be evaluated as

$$\begin{aligned}
& \frac{r_2 - \rho_2 + \delta}{c - \rho_2 + \delta} \frac{1}{\alpha_2 + \frac{1}{\lambda_2}} \int_{\beta}^{\infty} \mathbb{P}\left(A_2 > z_y\right) \mathrm{d}z_y, \\
=\ & \frac{r_2 - \rho_2 + \delta}{c - \rho_2 + \delta} \frac{\alpha_2}{\alpha_2 + \frac{1}{\lambda_2}} \mathbb{P}\left(A_2^{\mathrm{r}} > \beta\right) \\
=\ & \frac{\delta + \rho_2 p_2 - \delta p_2}{c - \rho_2 + \delta} \mathbb{P}\left(A_2^{\mathrm{r}} > \beta\right),
\end{aligned}$$

where the last equality follows from the definition of p_2. Letting $\delta \downarrow 0$, $\theta \downarrow 0$, using the fact that $A_2^{\mathrm{r}}(\cdot) \in \mathcal{IR}$, (4.32) follows.

We now turn to the proof of (4.33). We use time $\hat{v}$, which is related to v as defined above, but differs slightly, and plays a similar role as $\tilde{v}$ defined in the first part of this proof. By partitioning, we obtain for all $\alpha, \gamma, \theta, \kappa > 0$, $\hat{v} \geq w \geq 0$,

$$\begin{aligned}
& \mathbb{P}\left(T_2^c(-\gamma\hat{x}) > (1-\alpha)\hat{x}\right) \\
=\ & \mathbb{P}\left(T_2^c(-\gamma\hat{x}) > (1-\alpha)\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0, (1-\alpha)\hat{x}] \leq 1, V_2^c(\hat{v}) \geq (r_2 - c)(\hat{v} - w)\right) \\
& +\mathbb{P}\Big(T_2^c(-\gamma\hat{x}) > (1-\alpha)\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0, (1-\alpha)\hat{x}] \leq 1, \\
& \qquad V_2^c(\hat{v}) < (r_2 - c)(\hat{v} - w), \mathcal{N}_{\kappa\hat{x}}[\hat{v}, (1-\alpha)\hat{x}] = 0\Big) \\
& +\mathbb{P}\Big(T_2^c(-\gamma\hat{x}) > (1-\alpha)\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0, (1-\alpha)\hat{x}] \leq 1, \\
& \qquad V_2^c(\hat{v}) < (r_2 - c)(\hat{v} - w), \mathcal{N}_{\kappa\hat{x}}[\hat{v}, (1-\alpha)\hat{x}] = 1\Big)
\end{aligned}$$

$$+\mathbb{P}\left(T_2^c(-\gamma\hat{x}) > (1-\alpha)\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha)\hat{x}] \geq 2\right).$$

Now consider the third term. Suppose that $\mathcal{N}_{\kappa\hat{x}}[0,w] \geq 1$, i.e., there is a long on period in the interval $[0,w]$. Since it holds that $\mathcal{N}_{\kappa\hat{x}}[\hat{v},(1-\alpha)\hat{x}] = 1$ and $\mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha)\hat{x}] \leq 1$, this long on period must then last till at least time $\hat{v}$. However, this contradicts the fact that $V_2^c(\hat{v}) < (r_2-c)(\hat{v}-w)$. Hence, the third term may be rewritten as

$$\begin{aligned}
&\mathbb{P}\left(T_2^c(-\gamma\hat{x}) > (1-\alpha)\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha)\hat{x}] = 1,\right.\\
&\qquad \left. V_2^c(\hat{v}) < (r_2-c)(\hat{v}-w), \mathcal{N}_{\kappa\hat{x}}[\hat{v},(1-\alpha)\hat{x}] = 1, \mathcal{N}_{\kappa\hat{x}}[0,w] = 0\right)\\
=\; &\mathbb{P}\left(T_2^c(-\gamma\hat{x}) > (1-\alpha)\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha)\hat{x}] = 1, V_2^c(0) \leq \zeta\hat{x},\right.\\
&\qquad \left. V_2^c(\hat{v}) < (r_2-c)(\hat{v}-w), \mathcal{N}_{\kappa\hat{x}}[\hat{v},(1-\alpha)\hat{x}] = 1, \mathcal{N}_{\kappa\hat{x}}[0,w] = 0\right)\\
&+\mathbb{P}\left(T_2^c(-\gamma\hat{x}) > (1-\alpha)\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha)\hat{x}] = 1, V_2^c(0) > \zeta\hat{x},\right.\\
&\qquad \left. V_2^c(\hat{v}) < (r_2-c)(\hat{v}-w), \mathcal{N}_{\kappa\hat{x}}[\hat{v},(1-\alpha)\hat{x}] = 1, \mathcal{N}_{\kappa\hat{x}}[0,w] = 0\right).
\end{aligned}$$

We thus arrive at the upper bound, for all $\zeta > 0$,

$$\mathbb{P}\left(T_2^c(-\gamma\hat{x}) > (1-\alpha)\hat{x}\right) \leq II\text{–}B_a + II\text{–}B_b + II\text{–}B_c + II\text{–}B_d + II\text{–}B_e,$$

with,

$$\begin{aligned}
II\text{–}B_a &:= \mathbb{P}\left(V_2^c(\hat{v}) \geq (r_2-c)(\hat{v}-w)\right),\\
II\text{–}B_b &:= \mathbb{P}(T_2^c(-\gamma\hat{x}) > (1-\alpha)\hat{x}, V_2^c(\hat{v}) < (r_2-c)(\hat{v}-w),\\
&\qquad \mathcal{N}_{\kappa\hat{x}}[\hat{v},(1-\alpha)\hat{x}] = 0),\\
II\text{–}B_c &:= \mathbb{P}\left(T_2^c(-\gamma\hat{x}) > (1-\alpha)\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0,w] = 0, V_2^c(0) \leq \zeta\hat{x}\right),\\
II\text{–}B_d &:= \mathbb{P}\left(\mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha)\hat{x}] = 1, \mathcal{N}_{\kappa\hat{x}}[0,w] = 0, V_2^c(0) > \zeta\hat{x}\right),\\
II\text{–}B_e &:= \mathbb{P}\left(\mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha)\hat{x}] \geq 2\right).
\end{aligned}$$

Take $\hat{v} = \tau\hat{x}$ and $w = \xi\hat{x}$, where $\tau < 1-\alpha$ is redefined by

$$\tau := \frac{(\phi_2-\rho_2)(1-\alpha)}{r_2-\rho_2} + \frac{(\gamma+\theta)(\phi_2-\rho_2-r_2+c)}{(\phi_2-\rho_2)(\rho_2-r_2)} + \frac{(r_2-c)\zeta}{(\phi_2-\rho_2)(r_2-\rho_2)},$$

and $\xi < \tau$ is redefined by

$$\xi := \frac{\gamma+\zeta+\theta}{\phi_2-\rho_2}.$$

Observe again the similarity with τ and ξ as used here and in the proof of Theorem 4.8.1 Now consider term $II\text{–}B_a$. Using Theorem 4.1.2, it holds that $II\text{–}B_a$ is equal to

$$\begin{aligned}
&\mathbb{P}\left(V_2^c \geq (r_2-c)(\tau-\xi)\hat{x}\right)\\
\sim\; &p_2\frac{\rho_2}{c-\rho_2}\mathbb{P}\left(A_2^{\mathrm{r}} > (\tau-\xi)\hat{x}\right) = p_2\frac{\rho_2}{c-\rho_2}\mathbb{P}\left(A_2^{\mathrm{r}} > \psi_3\hat{x}\right),
\end{aligned}$$

with

$$\psi_3 := \left(\frac{(\phi_2-\rho_2)(1-\alpha)}{r_2-\rho_2} + \frac{(\gamma+\theta)(2r_2-\phi_2-c)+\zeta(\rho_2-c)}{(r_2-\rho_2)(\phi_2-\rho_2)}\right).$$

Next, consider term II–B_b. Similar to the derivation of term I_b in the proof of Theorem 4.7.1, the event $T_2(-\gamma\hat{x}) > (1-\alpha)\hat{x}$ implies

$$\inf_{\hat{v} \leq u \leq (1-\alpha)\hat{x}} \{A_2(\hat{v}, u) - \phi_2(u - \hat{v})\} > -V_2^c(\hat{v}) - (c - \phi_2)\hat{v} - \gamma\hat{x},$$

so that an upper bound for II–B_b is given by

$$\begin{aligned}
& \mathbb{P}\Big(\inf_{\hat{v} \leq u \leq (1-\alpha)\hat{x}} \{A_2(\hat{v}, u) - \phi_2(u - \hat{v})\} > -V_2^c(\hat{v}) - (c - \phi_2)\hat{v} - \gamma\hat{x}, \\
& \quad V_2^c(\hat{v}) < (r_2 - c)(\hat{v} - w), \mathcal{N}_{\kappa\hat{x}}[\hat{v}, (1-\alpha)\hat{x}] = 0 \Big) \\
= \; & \mathbb{P}\Big(\inf_{\hat{v} \leq u \leq (1-\alpha)\hat{x}} \{A_2(\hat{v}, u) - \phi_2(u - \hat{v})\} > -(r_2 - \phi_2)\hat{v} + (r_2 - c)w - \gamma\hat{x}, \\
& \quad \mathcal{N}_{\kappa\hat{x}}[\hat{v}, (1-\alpha)\hat{x}] = 0 \Big) \\
= \; & \mathbb{P}\Big(\inf_{0 \leq u \leq (1-\alpha)\hat{x} - \hat{v}} \{A_2(0, u) - \phi_2 u\} > -(r_2 - \phi_2)\hat{v} + (r_2 - c)w - \gamma\hat{x}, \\
& \quad \mathcal{N}_{\kappa\hat{x}}[0, (1-\alpha)\hat{x} - \hat{v}] = 0 \Big) \\
= \; & \mathbb{P}\Big(\inf_{0 \leq u \leq (1-\alpha-\tau)\hat{x}} \{A_2(0, u) - \phi_2 u\} > (\theta - (\phi_2 - \rho_2)(1 - \alpha - \tau))\hat{x}, \\
& \quad \mathcal{N}_{\kappa\hat{x}}[0, (1-\alpha-\tau)\hat{x}] = 0 \Big) \\
= \; & \mathbb{P}(T_2^c((\theta - (\phi_2 - \rho_2)(1 - \alpha - \tau))\hat{x}) > (1 - \alpha - \tau)\hat{x}, \\
& \quad \mathcal{N}_{\kappa\hat{x}}[0, (1-\alpha-\tau)\hat{x}] = 0).
\end{aligned}$$

Finally, consider term II–B_c. In the proof of Theorem 4.7.1, we showed that the event $T_2(-\gamma\hat{x}) > (1-\alpha)\hat{x}$ implies

$$\inf_{0 \leq u \leq w} \{A_2(0, u) - \phi_2 u\} > -V_2^c(0) - \gamma\hat{x},$$

so that an upper bound for II–B_c is given by

$$\begin{aligned}
& \mathbb{P}\left(\inf_{0 \leq u \leq w} \{A_2(0, u) - \phi_2 u\} > -V_2^c(0) - \gamma\hat{x}, V_2^c(0) \leq \zeta\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0, w] = 0 \right) \\
\leq \; & \mathbb{P}\left(\inf_{0 \leq u \leq w} \{A_2(0, u) - \phi_2 u\} > -(\gamma + \zeta)\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0, w] = 0 \right) \\
= \; & \mathbb{P}\left(\inf_{0 \leq u \leq \xi\hat{x}} \{A_2(0, u) - \phi_2 u\} > (\theta - (\phi_2 - \rho_2)\xi)\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0, \xi\hat{x}] = 0 \right) \\
= \; & \mathbb{P}\left(T_2((\theta - (\phi_2 - \rho_2)\xi)\hat{x}) > \xi\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0, \xi\hat{x}] = 0\right).
\end{aligned}$$

Thus, taking $\eta = 1 - \alpha - \tau$ in Lemma 4.6.3 for II–B_b, taking $\eta = \xi$ in Lemma 4.6.3 for II–B_c, using Lemma 4.6.8 for II–B_d, and using Lemma 4.6.5

for II–B_e, we obtain

$$\mathbb{P}\left(T_2^c(-\gamma\hat{x}) \geq (1-\alpha)\hat{x}\right) \lesssim p_2 \frac{\rho_2}{c-\rho_2} \mathbb{P}\left(A_2^{\mathrm{r}} > \psi_3 \hat{x}\right).$$

Letting $\zeta \downarrow 0$, $\theta \downarrow 0$, using the fact that $A_2^{\mathrm{r}}(\cdot) \in \mathcal{IR}$, (4.33) follows.

Finally, note that (4.34) follows from (4.32) and (4.33) by letting $\alpha \downarrow 0$, $\gamma \downarrow 0$, and using again the fact that $A_2^{\mathrm{r}}(\cdot) \in \mathcal{IR}$. □

4.10 Concluding remarks

We analyzed a GPS queue with two flows, one having light-tailed characteristics, the other one exhibiting heavy-tailed properties. We showed that the workload distribution of the light-tailed flow is asymptotically equivalent to its workload distribution when served in isolation at its minimum guaranteed rate, multiplied with a certain pre-factor. The pre-factor may be interpreted as the probability that the heavy-tailed flow is backlogged long enough for the light-tailed flow to reach overflow. We did not consider the case where the traffic intensity of the heavy-tailed flow exceeds its minimum guaranteed rate. In this case, the pre-factor – representing again the probability that the heavy-tailed flow is continuously backlogged during the period to overflow of the light-tailed flow – is likely to be some constant. Determining the exact value of the constant seems however a rather challenging task.

In the present chapter we have focused on a scenario with two flows. Observe however that the light-tailed flow may be thought of as an aggregate flow, given that the class of Markov-modulated fluid input is closed under superposition of independent processes. In case of instantaneous input, the heavy-tailed flow too may actually represent an aggregate flow, since the superposition of independent Poisson streams with regularly varying bursts produces again a Poisson stream with regularly varying bursts. Unfortunately, the class of on-off sources is clearly not closed under superposition. In fact, the superposition exhibits a fundamentally more complex structure than a single on-off-source, which drastically complicates the analysis of the queueing behavior, see [152].

Despite the above and earlier observations, it would still be interesting to extend the analysis to general scenarios with several light-tailed flows, let's say $N_1 \geq 1$, and $N_2 \geq 1$ heavy-tailed flows.

In case $N_1 = 1$, $N_2 > 1$, we expect that the asymptotic workload distribution of the light-tailed flow is equivalent to its workload distribution when served in isolation at its minimum guaranteed rate, multiplied with a certain pre-factor, exactly as before. In this case however, the pre-factor represents the probability that each of the heavy-tailed flows is constantly backlogged during the period to overflow of the light-tailed flow. Calculating this probability seems a demanding task, since the most likely scenario cannot be easily pinned down due to the complicated interaction of the heavy-tailed flows prior to the overflow period.

In case $N_1 > 1$, $N_2 = 1$, we conjecture that the asymptotic workload distribution of the light-tailed flows is equivalent to that in an isolated GPS queue consisting of the light-tailed flows only, multiplied again with a pre-factor. The pre-factor reflects the probability that the heavy-tailed flow is constantly backlogged during the time to overflow of the light-tailed flows. Unfortunately however, there are only logarithmic asymptotics known for a GPS queue with several light-tailed flows.

Not surprisingly, the two above-described complicating circumstances conspire in scenarios with $N_1 > 1$, $N_2 > 1$.

Chapter 5

Coupled processors

In this chapter we consider a model that is intimately related to the two-class GPS system, the coupled-processors model. As explained in the introductory chapter, i.e., Section 1.7, a two-class GPS system can be seen as a limiting case of the coupled-processors model. The two-class coupled-processors model is a system with two traffic classes, each with their own queue, in which the service capacity is shared as follows. When both classes are backlogged, the two corresponding queues are each served at unit rate. However, the service rate for class 1 increases to r_1 when the queue of class 2 (denoted by Q_2) empties, and that for class 2 increases to r_2 when the queue of class 1 (denoted by Q_1) empties. One of the differences with the two-class GPS system is that for the coupled-processors model we have to assume that $1/r_1 + 1/r_2 \neq 1$, whereas $1/r_1 + 1/r_2 = 1$ corresponds to the two-class GPS system. This chapter presents the results of [23] and [24]. For completeness we also include the results that have already appeared in Section 6 of [20] (see Sections 5.3 and 5.5).

For all models in this chapter we are interested in the asymptotic behavior of the workload at Q_1, under various assumptions for the input traffic of Q_1 and Q_2. We consider both the scenario where the class-1 offered traffic load $\rho_1 < 1$ and the scenario $\rho_1 > 1$. The case $\rho_1 < 1$ will be denoted by *underload* and is divided into a case with $\rho_2 < 1$, i.e., both Q_1 and Q_2 are stable by themselves, and a case with $\rho_2 > 1$, i.e., Q_1 is stable, but Q_2 relies on surplus capacity of Q_1 to be stable. The results show that if $\rho_1 < 1$ then the value of ρ_2 only affects the prefactor of the asymptotic workload distribution of class 1. In case Q_1 has exponential input, then the result will be somewhat more subtle as explained below. Similarly, we denote the scenario $\rho_1 > 1$ by *overload*. In this case we need $\rho_2 < 1$ for the stability of the total system.

Section 5.4. We start this introduction with the model in Section 5.4, because the results are similar to the results obtained in the previous chapter. Recall that in Chapter 4 we were interested in the workload of a light-tailed flow when served in a two-class GPS system together with a heavy-tailed flow. For the heavy-tailed flow we considered two cases in the previous chapter, either it behaves as an on-off process, or it sends instantaneous bursts. In this chapter we can only deal with the scenario where the heavy-tailed flow generates instantaneous bursts. This is because we rely completely on the Laplace Stieltjes transform of the joint workload distribution as derived in [42], which is not available for fluid input. Another difference between the model in Section 5.4 and the previous chapter concerns the light-tailed flow. In the previous chapter we assumed the light-tailed flow to behave according to a Markov modulated fluid process, whereas in Section 5.4 we have to assume that the light-tailed flow sends exponential bursts according to a Poisson process. Apart from these minor differences in the model assumptions, the major difference between the previous chapter and Section 5.4 lies in the method. In Chapter 4 we used a probabilistic sample-path approach, whereas in this chapter we rely completely on Tauberian theorems (see Section 2.2 for introductions on both methods). The asymptotic behavior that we derive in Section 5.4 is as follows. In the underload scenario the asymptotic workload distribution of class 1 is either exponential or semi-exponential, depending on whether ρ_2 exceeds its nominal service rate or not.

Before introducing the other models that are presented in this chapter, we remark that several recent studies have revealed a similar dichotomy in the interplay between exponential and heavy-tailed traffic processes. For example, [36] shows a similar contrast in an M/M/1 queue which alternates between exponentially distributed periods of high service speed and heavy-tailed periods of low service speed. A related phenomenon is observed in [33] for an M/G/2 queue with heterogeneous servers, one having exponential properties, the other one exhibiting heavy-tailed characteristics.

Sections 5.3, 5.5 and 5.6. A special case of the coupled-processors model as presented in Section 5.3 (underload scenario), and in Section 5.5 (overload scenario) has been studied in [19]. It is assumed in [19] that both classes have heavy-tailed characteristics and the value of r_2 is taken to be 1, so that Q_2 is not influenced by Q_1. In the present chapter, we generalize this model, and the results obtained for this model, in various ways. First, we remove the assumption that $r_2 = 1$ in all models that we consider, so that Q_2 is now also influenced by Q_1. Second, in Section 5.3, we extend the results for $\rho_1 < 1$ to the case where class 2 has a general service time distribution. Third, in Section 5.6, we extend the results for $\rho_1 > 1$ to the case where one of the classes has a regularly varying service time distribution and the other has a service time distribution with a lighter general tail. Similar to the results in [19], the results in this chapter show a stark contrast in the asymptotic workload behavior at Q_1, depending on whether the offered traffic at Q_1 exceeds the nominal unit service rate or not. In the underload scenario (Section 5.3), the workload at Q_1 behaves asymptotically as if Q_1 is served in isolation at a *constant* rate, which only depends on the service time distribution at Q_2 through its mean. In the overload scenario (Section 5.5), the asymptotic workload distribution at Q_1 is determined by the *heaviest* of the two service time distributions, so that Q_1 may inherit potentially heavier-tailed characteristics from Q_2. In Section 5.6 we show that the results of Section 5.5 continue to hold.

This chapter is organized in the following way (see Table 5.1, where we use the notation B_i for the service time distribution at Q_i, $i = 1, 2$). In Section 5.1,

		B_1			
		$\mathcal{R}_{-\nu_1}$		exp	$o(B_2)$
B_2	ρ_2	$\rho_1 < 1$	$\rho_1 > 1$	$\rho_1 < 1$	$\rho_1 > 1$
general	$\rho_2 < 1$	Thm 5.3.1			
	$\rho_2 > 1$	Thm 5.3.2	total queue unstable		total queue unstable
$\mathcal{R}_{-\nu_2}$	$\rho_2 < 1$	see Thm 5.3.1	Thm 5.5.1	Thm 5.4.1	Thm 5.6.1 part *(ii)*
	$\rho_2 > 1$	see Thm 5.3.2	total queue unstable	Thm 5.4.2	total queue unstable
$o(B_1)$	$\rho_2 < 1$	see Thm 5.3.1	Thm 5.6.1 part *(i)*		

Table 5.1: Overview of results in Chapter 5.

we present a detailed model description. Section 5.2 contains some preliminary results which provide the basis for the subsequent analysis. In Section 5.3, we obtain the workload asymptotics at Q_1 in case its service time distribution is heavy-tailed, while the service time distribution at Q_2 is allowed to be general. In Section 5.4, we derive the workload asymptotics at Q_1 in case its service time distribution is exponential, while the service time distribution at Q_2 is assumed to be heavy-tailed. In both Sections 5.3 and 5.4 we assume that $\rho_1 < 1$, and consider the case $\rho_2 < 1$ as well as $\rho_2 > 1$. In Section 5.5, we characterize the workload asymptotics at Q_1 for the case where $\rho_1 > 1$ (which forces $\rho_2 < 1$) and the service time distribution at both Q_1 and Q_2 is heavy-tailed. In Section 5.6, we extend the results of Section 5.5 to situations where the service time distribution at one of the queues is heavy-tailed and the service time distribution at the other queue has a lighter, general tail.

5.1 Model description

We consider a system with two heterogeneous traffic classes. Class-i customers arrive as a Poisson process of rate λ_i, and require an amount of service B_i with mean $\beta_i < \infty$ and Laplace-Stieltjes Transform (LST) $\beta_i\{s\} := \mathbb{E}e^{-sB_i}$, Re $s \geq 0$. Define $\rho_i := \lambda_i \beta_i$ as the traffic intensity of class i.

In the next sections we will consider various scenarios for the service time distributions of the two classes. In all scenarios, the service time distribution of at least one of the two classes will be assumed to be regularly varying. We take a regularly varying function that has a different form than the function defined in Definition 2.1.2, and assume for either $i = 1$, $i = 2$ or both,

$$\mathbb{P}(B_i > t) \sim \frac{C_i}{-\Gamma(1-\nu_i)} t^{-\nu_i} l_i(t), \qquad t \to \infty, \tag{5.1}$$

with $l_i(\cdot)$ some slowly varying function (recall that $\lim_{t\to\infty} l_i(\eta t)/l_i(t) = 1$,

$\eta > 1$), C_i a constant, and $\Gamma(\cdot)$ the Gamma function. Throughout the present chapter, we assume that $1 < \nu_i < 2$, which is a particularly interesting case since the variance of B_i is then infinite. With some minor modifications the results may be extended to the case $\nu_i > 2$. Recall that we denote with B_i^{r} the random variable with the distribution of the residual life of B_i. In particular, note that if $\mathbb{P}(B_i > t)$ behaves as in (5.1), then

$$\mathbb{P}(B_i^{\mathrm{r}} > t) \sim \frac{C_i}{\beta_i \Gamma(2-\nu_i)} t^{1-\nu_i} l_i(t), \qquad t \to \infty; \tag{5.2}$$

it is again regularly varying, but of index $1 - \nu_i$.

There are separate queues maintained for each class, Q_1 and Q_2. When both classes are backlogged, each of the queues is served at unit rate. However, the service rate at Q_i increases to $r_i \geq 1$ when the other queue is empty. As explained, for $r_i = 1/\phi_i$, with $\phi_1 + \phi_2 = 1$, the model may equivalently be viewed as a two-class GPS system (with suitably scaled inputs and service). In this chapter we concentrate on the more general case $1/r_1 + 1/r_2 \neq 1$. A careful examination shows however that the case $1/r_1 + 1/r_2 = 1$ does not represent a boundary case, suggesting that the main results remain valid when $1/r_1 + 1/r_2 \to 1$.

Throughout this chapter, we assume that the ergodicity conditions are satisfied. These conditions are discussed in Section III.3.7 of [42]. Here, it suffices to observe that $\max\{\rho_1, \rho_2\} < 1$ is sufficient but not necessary for ergodicity, while $\min\{\rho_1, \rho_2\} < 1$ is necessary.

5.2 Preliminary results

In this section we review some preparatory results which will provide the starting point for the analysis. Denote by V_i a random variable representing the workload at Q_i in steady state. For $c > 0$, denote by V_i^c a random variable representing the steady-state workload at Q_i when served in isolation at a constant rate c. For $\mathrm{Re}\, a_1 \geq 0$, $\mathrm{Re}\, a_2 \geq 0$, let

$$\begin{aligned}
\psi(a_1, a_2) &:= \mathbb{E}e^{-a_1 V_1 - a_2 V_2}, \\
\psi_1(a_2) &:= \mathbb{E}\left[e^{-a_2 V_2} \mathrm{I}_{\{V_1 = 0\}}\right], \\
\psi_2(a_1) &:= \mathbb{E}\left[e^{-a_1 V_1} \mathrm{I}_{\{V_2 = 0\}}\right], \\
\psi_0 &:= \mathbb{P}(V_1 = 0, V_2 = 0),
\end{aligned}$$

with $\mathrm{I}_{\{A\}}$ denoting the indicator function of the event A.

According to Formula (2.16) of Chapter III.3 of [42] (in the remainder of this chapter we omit Chapter III.3 when referring to formulas from [42]), for $\mathrm{Re}\, s \geq 0$,

$$\mathbb{E}e^{-sV_1} = \mathbb{E}e^{-sV_1^1} \left(\frac{\psi_1(0)}{1-\rho_1} + \frac{r_1 - 1}{1-\rho_1}(\psi_0 - \psi_2(s)) \right), \tag{5.3}$$

with, since we have assumed the class-1 service times to be exponentially distributed,

$$\mathbb{E}e^{-sV_1^1} = \frac{(1-\rho_1)(1+\beta_1 s)}{1-\rho_1+\beta_1 s}. \tag{5.4}$$

It should be noted that the denominator has a pole $s = s_1 := \lambda_1 - 1/\beta_1 < 0$. This pole s_1 will play an essential role in the analysis of Section 5.4. Taking $s = 0$ in (5.3), we obtain

$$\frac{\psi_1(0)}{1-\rho_1} + \frac{r_1-1}{1-\rho_1}(\psi_0 - \psi_2(0)) = 1,$$

so that (5.3) may be rewritten as

$$\mathbb{E}e^{-sV_1} = \mathbb{E}e^{-sV_1^1}\left(1 - \frac{r_1-1}{1-\rho_1}(\psi_2(s) - \psi_2(0))\right). \tag{5.5}$$

We now focus on the function $\psi_2(s)$. According to Formulas (6.21), (6.22) and (6.23) of [42],

$$\psi_2(\delta_1(w)) - \psi_0 = \frac{1}{r_1}\frac{\psi_0}{1-\frac{1}{r_1}-\frac{1}{r_2}}\left(1 - e^{-R_1(w)+R_2(w)}\right), \quad \text{Re } w \geq 0, \tag{5.6}$$

and

$$\psi_1(\delta_2(w)) - \psi_0 = \frac{1}{r_2}\frac{\psi_0}{1-\frac{1}{r_1}-\frac{1}{r_2}}\left(1 - e^{P_1(w)-P_2(w)}\right), \quad \text{Re } w \leq 0, \tag{5.7}$$

with

$$\psi_0 = e^{-P_1(0)-R_2(0)}. \tag{5.8}$$

It remains to specify the functions $R_i(w)$, $P_i(w)$, and $\delta_i(w)$, $i = 1, 2$:

$$R_i(w) := \sum_{n=1}^{\infty} \frac{b_i^n}{n}\mathbb{E}\left[e^{-w\sigma_n^{(i)}} \mathrm{I}_{\{\sigma_n^{(i)}>0\}}\right], \quad \text{Re } w \geq 0,$$

and

$$P_i(w) := \sum_{n=1}^{\infty} \frac{b_i^n}{n}\mathbb{E}\left[e^{-w\sigma_n^{(i)}} \mathrm{I}_{\{\sigma_n^{(i)}<0\}}\right], \quad \text{Re } w \leq 0, \tag{5.9}$$

with

$$b_1 := \rho_1\left(1 - \frac{1}{r_2}\right) + \frac{\rho_2}{r_2}, \tag{5.10}$$

$$b_2 := \rho_2\left(1 - \frac{1}{r_1}\right) + \frac{\rho_1}{r_1}, \tag{5.11}$$

and for $i = 1, 2$,

$$\sigma_n^{(i)} := X_{i1} + \cdots + X_{in}, \tag{5.12}$$

with $X_{11}, \ldots, X_{1n}$ i.i.d. and $X_{21}, \ldots, X_{2n}$ i.i.d., and

$$X_{11} = \begin{cases} \hat{P}_1 & w.p.\ \pi_1 := \frac{\rho_1}{b_1}\left(1 - \frac{1}{r_2}\right), \\ -\hat{P}_2 & w.p.\ 1 - \pi_1 = \frac{\rho_2}{b_1 r_2}, \end{cases} \tag{5.13}$$

$$X_{21} = \begin{cases} \hat{P}_1 & w.p.\ \pi_2 := \frac{\rho_1}{b_2 r_1}, \\ -\hat{P}_2 & w.p.\ 1 - \pi_2 = \frac{\rho_2}{b_2}\left(1 - \frac{1}{r_1}\right). \end{cases} \tag{5.14}$$

Here $\hat{P}_i$ is a random variable representing the busy period of class i when served in isolation at unit rate, starting with an *exceptional* first service time B_i^{r} (a residual service time). The functions $\delta_i(w)$, $i = 1, 2$, play a crucial role in the analysis: $\delta_1(w)$ is defined for Re $w \geq 0$ as zero of the function

$$f_1(s, w) := \lambda_1 (1 - \beta_1\{s\}) - s + w, \tag{5.15}$$

which has for Re $w \geq 0$, $w \neq 0$, exactly one zero $s = \delta_1(w)$ in Re $s \geq 0$, and this zero has multiplicity one. Following the three regimes for ρ_1 yields:

- $f_1(s, 0)$ has for $\rho_1 < 1$ exactly one zero $s = \delta_1(0) = 0$ in Re $s \geq 0$, with multiplicity one;
- $f_1(s, 0)$ has for $\rho_1 = 1$ exactly one zero $s = \delta_1(0) = 0$ in Re $s \geq 0$, with multiplicity two;
- $f_1(s, 0)$ has for $\rho_1 > 1$ two zeroes $s = \delta_1(0) > 0$ and $s = \epsilon_1(0) = 0$ in Re $s \geq 0$, each with multiplicity one.

Similarly, $\delta_2(w)$ is defined for Re $w \leq 0$ as zero of the function

$$f_2(s, w) := \lambda_2 (1 - \beta_2\{s\}) - s - w. \tag{5.16}$$

5.3 Q1 heavy-tailed and Q2 general: underload

5.3.1 Assumptions and main results

In this section we consider the case where the service time distribution at Q_1 is regularly varying of index $-\nu_1$, i.e.,

$$\mathbb{P}(B_1 > t) \sim \frac{C_1}{-\Gamma(1 - \nu_1)} t^{-\nu_1} l_1(t), \qquad t \to \infty, \tag{5.17}$$

with $l_1(\cdot)$ some slowly varying function. According to Theorem 2.2.1, the asymptotic relation in (5.17) with $1 < \nu_1 < 2$ is equivalent to (using $\sim$ according to the same notational convention for $s \downarrow 0$ as previously for $t \to \infty$):

$$1 - \frac{1 - \beta_1\{s\}}{\beta_1 s} \sim \frac{C_1}{\beta_1} s^{\nu_1 - 1} l_1\left(\frac{1}{s}\right), \qquad s \downarrow 0. \tag{5.18}$$

The service time distribution at Q_2 is allowed to be general. We assume that $\rho_1 < 1$, so that even the lower service speed at Q_1 is sufficient for stability. The case $\rho_1 > 1$ is treated in Section 5.5. We show that the workload distribution at Q_1 is regularly varying of index $1 - \nu_1$, irrespective of the service time distribution at Q_2. We consider the cases $\rho_2 < 1$ and $\rho_2 > 1$ separately, as their proofs involve some subtle differences. The next theorem gives the asymptotic behavior of the workload distribution for the case where $\rho_2 < 1$ (Theorem 6.1 in [20]).

Theorem 5.3.1 *If $\mathbb{P}(B_1 > t) \in \mathcal{R}_{-\nu_1}$, $1 < \nu_1 < 2$, as given in (5.17), and if $\rho_1 < 1$, $\rho_2 < 1$, then $\mathbb{P}(V_1 > t) \in \mathcal{R}_{1-\nu_1}$ and*

$$\mathbb{P}(V_1 > t) \sim \frac{1}{K_1 - \rho_1} \frac{\lambda_1 C_1}{\Gamma(2-\nu_1)} t^{1-\nu_1} l_1(t) \sim \frac{\rho_1}{K_1 - \rho_1} \mathbb{P}(B_1^{\mathrm{r}} > t), \qquad t \to \infty,$$

with $K_1 := \rho_2 + (1-\rho_2) r_1 \geq 1$ representing the average service rate at Q_1 when continuously backlogged.

In the next theorem the asymptotic behavior of the workload distribution is given for the case $\rho_2 > 1$ (Remark 6.3 in [20]).

Theorem 5.3.2 *If $\mathbb{P}(B_1 > t) \in \mathcal{R}_{-\nu_1}$, $1 < \nu_1 < 2$, as given in (5.17), and if $\rho_1 < 1$, $\rho_2 > 1$, then $\mathbb{P}(V_1 > t) \in \mathcal{R}_{1-\nu_1}$ and*

$$\mathbb{P}(V_1 > t) \sim \frac{1}{1 - \rho_1} \frac{\lambda_1 C_1}{\Gamma(2-\nu_1)} t^{1-\nu_1} l_1(t) \sim \frac{\rho_1}{1 - \rho_1} \mathbb{P}(B_1^{\mathrm{r}} > t), \qquad t \to \infty.$$

The proofs of Theorems 5.3.1 and 5.3.2 are presented in Section 5.3.3. We give the proofs because they are relatively easy compared to the proofs in Section 5.4, so they could serve as a first step towards understanding the more complicated proofs.

5.3.2 Discussion and interpretation

We start with Theorem 5.3.1 for the case $\rho_2 < 1$. Invoking a result in [39], we can rewrite the result as

$$\mathbb{P}(V_1 > t) \sim \mathbb{P}\left(V_1^{K_1} > t\right), \qquad t \to \infty. \tag{5.19}$$

Thus, asymptotically, the workload at Q_1 behaves exactly as if Q_1 is served in isolation at a *constant* rate K_1. This may be intuitively interpreted as follows. Large-deviations arguments suggest that the most likely scenario for the workload at Q_1 to reach a large level is that class 1 generates a large amount of traffic, while class 2 shows average behavior. Specifically, suppose that a class-1 customer arrives with a large service time, so that Q_1 becomes backlogged for a long period of time. During that period, Q_2 will not receive any surplus capacity, and hence be empty only a fraction $1 - \rho_2$ of the time. As a result, the average service rate at Q_1 while backlogged is $K_1 = \rho_2 + (1-\rho_2) r_1$. Thus, Q_1 is effectively served at a constant rate $K_1 \geq 1$, as confirmed by Formula (5.19).

We stress that the result holds regardless of the service time distribution at Q_2 as long as both $\rho_1 < 1$ and $\rho_2 < 1$. In that sense, Q_1 is not significantly affected by the interaction with Q_2. In particular, Q_1 is virtually immune for 'heavier'-tailed service time characteristics at Q_2. That will no longer be the case when $\rho_1 > 1$, as we will see in Section 5.5.

We now turn to the case $\rho_2 > 1$. The overflow scenario in the case $\rho_2 < 1$ as described above suggests that Theorem 5.3.1 should continue to hold, except that the service rate K_1 should be reduced to 1. Specifically, the most likely scenario for the workload at Q_1 to reach a large level is again that a class-1 customer arrives with a large service time, so that Q_1 becomes backlogged for a long period of time. As before, Q_2 will not receive any surplus capacity, and hence soon become persistently backlogged too, since now $\rho_2 > 1$. Thus, Q_1 is only served at rate 1 (rather than $K_1 \geq 1$) while backlogged, as is confirmed by Theorem 5.3.2.

5.3.3 Proofs

Proof of Theorem 5.3.1 The approach may be outlined as follows. Formula (5.5) expresses $\mathbb{E}e^{-sV_1}$ in terms of $\psi_2(s)$. Formula (5.6) relates $\psi_2(s)$, or rather $\psi_2(\delta_1(w))$, to $R_1(w)$ and $R_2(w)$. We use these formulas to derive the behavior of $\mathbb{E}e^{-sV_1}$ for $s \downarrow 0$. Theorem 2.2.1 then yields the behavior of $\mathbb{P}(V_1 > t)$ for $t \to \infty$. Therefore, we now concentrate on the behavior of $R_1(w)$ and $R_2(w)$ and, first, $\delta_1(w)$ for $w \downarrow 0$.

Let P_1 denote a random variable with distribution the steady-state distribution of a busy period at Q_1 in isolation, i.e., an M/G/1 queue with arrival rate λ_1 and service time distribution $B_1(\cdot)$. Comparing (5.15) with the Takács equation for the busy-period LST $\mathbb{E}e^{-wP_1}$, cf. p. 250 in [40], it is seen that

$$\delta_1(w) = w + \lambda_1 \left(1 - \mathbb{E}e^{-wP_1}\right).$$

In [47] it is proven that $\mathbb{P}(P_1 > t)$ is regularly varying of index $-\nu_1$ iff $\mathbb{P}(B_1 > t)$ is regularly varying of index $-\nu_1$, and if either holds then

$$\mathbb{P}(P_1 > t) \sim \frac{1}{1-\rho_1}\mathbb{P}\left(\frac{B_1}{1-\rho_1} > t\right), \qquad t \to \infty. \tag{5.20}$$

Theorem 2.2.1 then gives the behavior of $\mathbb{E}e^{-wP_1} - 1$ for $w \downarrow 0$. We conclude that, if (5.17) holds, then

$$\delta_1(w) - \frac{w}{1-\rho_1} \sim -\frac{\lambda_1 C_1}{1-\rho_1}\left(\frac{w}{1-\rho_1}\right)^{\nu_1} l_1\left(\frac{1}{w}\right), \qquad w \downarrow 0. \tag{5.21}$$

In addition, using (5.15), we have for $\delta_1^{-1}(s) = s - \lambda_1(1 - \beta_1\{s\})$:

$$\delta_1^{-1}(s) - (1-\rho_1)s \sim \lambda_1 C_1 s^{\nu_1} l_1\left(\frac{1}{s}\right), \qquad s \downarrow 0.$$

In the study of $R_i(w)$, a key role is played by the LST of $\hat{P}_1$, a busy period at Q_1 in isolation that starts with a *residual* service time. From (6.4) of [42],

$$\mathbb{E}e^{-w\hat{P}_1} = \frac{1-\beta_1\{\delta_1(w)\}}{\beta_1\delta_1(w)}, \qquad \text{Re } w \geq 0.$$

It is now readily verified that

$$1 - \mathbb{E}e^{-w\hat{P}_1} \sim \frac{C_1}{\beta_1}\left(\frac{w}{1-\rho_1}\right)^{\nu_1-1} l_1\left(\frac{1}{\delta_1(w)}\right), \qquad w \downarrow 0,$$

and hence, using Theorem 2.2.1, $\mathbb{P}\left(\hat{P}_1 > t\right)$ is seen to be regularly varying of index $1-\nu_1$,

$$\mathbb{P}\left(\hat{P}_1 > t\right) \sim \frac{C_1}{\beta_1\Gamma(2-\nu_1)}((1-\rho_1)t)^{1-\nu_1} l_1(t), \qquad t \to \infty. \tag{5.22}$$

The difference with (5.20) is caused by the *residual* service time with which the busy period starts; it is regularly varying of one index higher than an ordinary service time.

We are now ready to study the tail behavior of $R_i(w)$. Observe that $R_i(w)$ is the LST of

$$r_i(t) := \sum_{n=1}^{\infty} \frac{b_i^n}{n} \mathbb{P}\left(0 < X_{i1} + \cdots + X_{in} < t\right), \qquad t > 0.$$

Consider

$$R_i(0) - r_i(t) = \sum_{n=1}^{\infty} \frac{b_i^n}{n} \mathbb{P}\left(X_{i1} + \cdots + X_{in} > t\right), \qquad t > 0. \tag{5.23}$$

Using well-known properties of heavy-tailed random variables (see Proposition 1.1 in [127]), it may be shown as in [19] that, irrespective of the distribution of B_2,

$$\mathbb{P}\left(X_{21} + \cdots + X_{2n} > t\right) \sim n\pi_2 \mathbb{P}\left(\hat{P}_1 > t\right), \qquad t \to \infty. \tag{5.24}$$

We conclude from (5.22), (5.23) and (5.24) that

$$\begin{aligned} R_2(0) - r_2(t) &\sim \frac{b_2\pi_2}{1-b_2}\mathbb{P}\left(\hat{P}_1 > t\right) \\ &\sim \frac{1}{(1-b_2)r_1}\frac{\lambda_1 C_1}{\Gamma(2-\nu_1)}((1-\rho_1)t)^{1-\nu_1} l_1(t), \quad t \to \infty. \end{aligned} \tag{5.25}$$

Again applying Theorem 2.2.1,

$$R_2(w) - R_2(0) \sim -\frac{\lambda_1 C_1}{(1-b_2)r_1}\left(\frac{w}{1-\rho_1}\right)^{\nu_1-1} l_1\left(\frac{1}{w}\right), \qquad w \downarrow 0. \tag{5.26}$$

Similarly, it is seen that

$$\mathbb{P}\left(X_{11}+\cdots+X_{1n}>t\right)\sim n\pi_1\mathbb{P}\left(\hat{P}_1>t\right),\qquad t\to\infty,$$

leading to

$$R_1(w)-R_1(0)\sim-\frac{\lambda_1C_1\left(1-\frac{1}{r_2}\right)}{1-b_1}\left(\frac{w}{1-\rho_1}\right)^{\nu_1-1}l_1\left(\frac{1}{w}\right),\qquad w\downarrow 0.\tag{5.27}$$

It follows from (5.6), (5.26) and (5.27) after a lengthy calculation that

$$\psi_2(\delta_1(w))-\psi_2(0)\sim-\frac{\lambda_1C_1(1-\rho_2)}{(1-b_2)r_1}\left(\frac{w}{1-\rho_1}\right)^{\nu_1-1}l_1\left(\frac{1}{w}\right),\qquad w\downarrow 0.\tag{5.28}$$

Finally, see (5.21),

$$\psi_2(s)-\psi_2(0)\sim-\frac{\lambda_1C_1(1-\rho_2)}{(1-b_2)r_1}s^{\nu_1-1}l_1\left(\frac{1}{s}\right),\qquad s\downarrow 0.\tag{5.29}$$

Using (5.5), (5.18) and (5.29), it follows that for $s\downarrow 0$:

$$\mathbb{E}e^{-sV_1}-1\sim-\left(\frac{1}{1-\rho_1}-\frac{(r_1-1)(1-\rho_2)}{1-\rho_1}\frac{1}{(1-b_2)r_1}\right)\lambda_1C_1s^{\nu_1-1}l_1\left(\frac{1}{s}\right).$$

Using (5.11), we can rewrite this into

$$\mathbb{E}e^{-sV_1}-1\sim-\frac{\lambda_1C_1}{K_1-\rho_1}s^{\nu_1-1}l_1\left(\frac{1}{s}\right),\qquad s\downarrow 0,\tag{5.30}$$

with $K_1=\rho_2+(1-\rho_2)r_1$. Applying Theorem 2.2.1 once more completes the proof of the theorem. □

Proof of Theorem 5.3.2 The proof is largely similar to that of Theorem 5.3.1. The main difference is that the distribution of $\hat{P}_2$ is now defective: $\mathbb{P}(\hat{P}_2<\infty)=1/\rho_2<1$. As in Section IV of [19], it may be shown that the first part of (5.25) changes into

$$R_2(0)-r_2(t)\sim\frac{b_2\pi_2}{1-b_2\left(\pi_2+\frac{1-\pi_2}{\rho_2}\right)}\mathbb{P}\left(\hat{P}_1>t\right),\qquad t\to\infty.\tag{5.31}$$

Similarly,

$$R_1(0)-r_1(t)\sim\frac{b_1\pi_1}{1-b_1\left(\pi_1+\frac{1-\pi_1}{\rho_2}\right)}\mathbb{P}\left(\hat{P}_1>t\right),\qquad t\to\infty.\tag{5.32}$$

From (5.10)-(5.14), it may be checked that the constants in the two right-hand sides both equal $\rho_1/(1-\rho_1)$. Using (5.22) and Theorem 2.2.1, we obtain

$$\begin{aligned}R_1(w)-R_1(0)&\sim R_2(w)-R_2(0)\\&\sim-\frac{\lambda_1C_1}{1-\rho_1}\left(\frac{-w}{1-\rho_1}\right)^{\nu_1-1}l_1\left(\frac{-1}{w}\right),\qquad w\uparrow 0.\end{aligned}\tag{5.33}$$

Substituting into (5.6), we have

$$\psi_2(\delta_1(w)) - \psi_2(0) = o\left(w^{\nu_1 - 1} l_1\left(\frac{1}{w}\right)\right), \quad w \downarrow 0. \tag{5.34}$$

Using (5.5), (5.18) and (5.34) (instead of (5.28)), we obtain

$$\mathbb{E}e^{-sV_1} - 1 \sim -\frac{\lambda_1 C_1}{1-\rho_1} s^{\nu_1 - 1} l_1\left(\frac{1}{s}\right), \quad s \downarrow 0.$$

Applying Theorem 2.2.1 then completes the proof of the theorem. □

5.4 Q1 exponential and Q2 heavy-tailed: underload

5.4.1 Assumptions and main results

In this section we consider the case where the service time distribution at Q_1 is exponential. The service time distribution at Q_2 is assumed to be regularly varying of index $-\nu_2$, $1 < \nu_2 < 2$, i.e.,

$$\mathbb{P}(B_2 > t) \sim \frac{C_2}{-\Gamma(1-\nu_2)} t^{-\nu_2} l_2(t), \quad t \to \infty, \tag{5.35}$$

with $l_2(\cdot)$ some slowly varying function. We assume that $\rho_1 < 1$, so that even the lower service speed at Q_1 is sufficient for stability. The case $\rho_1 > 1$ is covered in Section 5.6. We show that the workload distribution at Q_1 is either semi-exponential (see Theorem 5.4.1), if $\rho_2 < 1$, or exponential (see Theorem 5.4.2), if $\rho_2 > 1$. The next theorem gives the asymptotic behavior of $\mathbb{P}(V_1 > t)$ for $\rho_2 < 1$.

Theorem 5.4.1 *If $\mathbb{P}(B_1 > t)$ is exponential, and $\mathbb{P}(B_2 > t) \in \mathcal{R}_{-\nu_2}$, $1 < \nu_2 < 2$, as given in (5.35), and if $\rho_1 < 1$, $\rho_2 < 1$, then*

$$\begin{aligned}\mathbb{P}(V_1 > t) &\sim \rho_1 e^{(\lambda_1 - 1/\beta_1)t} \frac{1}{K_2 - \rho_2} \frac{\lambda_2 C_2}{\Gamma(2-\nu_2)} \left(\frac{\rho_1(1-\rho_2)}{1-\rho_1} t\right)^{1-\nu_2} l_2(t) \\ &\sim \mathbb{P}(V_1^1 > t) \frac{\rho_2}{K_2 - \rho_2} \mathbb{P}\left(B_2^{\mathrm{r}} > \frac{\rho_1(1-\rho_2)}{1-\rho_1} t\right), \quad t \to \infty,\end{aligned}$$

with $K_2 := \rho_1 + (1-\rho_1) r_2 \geq 1$ representing the average service rate at Q_2 when continuously backlogged.

Recall that a two-class coupled-processors model can be obtained from a two-class GPS model by rescaling the input processes (see Section 1.7). As explained in Remark 4.5.1, with some minor modifications of the proofs, the results in the previous chapter can be shown to hold for the case where flow 1 generates exponentially distributed bursts. We now verify that the result of Theorem 4.2.1 in that case agrees with the result for $\rho_2 < 1$ in Theorem 5.4.1 above, under the assumption $1/r_1 + 1/r_2 = 1$. We denote the bursts of flow i in the

GPS model by $B_{i,G}$, and the bursts of flow i in the coupled-processors model by $B_{i,C}$. Following the reasoning in Section 1.7, we have $B_{i,G} = \phi_i B_{i,C}$, with $\phi_i = 1/r_i$. Similarly we define the amount of traffic generated in the interval $(s,t]$ in the coupled-processors model and the GPS model by $A_{i,C}(s,t)$ and $A_{i,G}(s,t)$ respectively, with $A_{i,G}(s,t) = \phi_i A_{i,C}(s,t)$. We define $V_{1,C}$ to be the workload of flow 1 in the coupled-processors model, $V_{1,G}$ to be the workload of flow 1 in the GPS model, and $V_{1,B}^{c}$ to be the workload of flow 1 with bursts B when it is served in isolation at rate c. Denoting the average rates in the GPS model with $\rho_{i,G}$, we have $\rho_{i,G} = \phi_i \rho_{i,C}$. Now recall the result of Case I as given in Theorem 4.2.1:

$$\begin{aligned} &\mathbb{P}\left(V_{1,G} > x\right) \\ \sim\ &\mathbb{P}\left(V_{1,B_{1,G}}^{\phi_1} > x\right) \frac{\rho_{2,G}}{1-\rho_{1,G}-\rho_{2,G}} \mathbb{P}\left(B_{2,G}^{\mathrm{r}} > \frac{x(\phi_2-\rho_{2,G})}{\hat{\rho}_{1,G}-\phi_1}\right). \end{aligned}$$

We will first rewrite the workload in the coupled-processors model in terms of the workload in the GPS model. Denoting the amount of service available for flow 1 in the coupled-processors model by $c_{1,C}(s,t)$, and defining $c_{1,G}(s,t)$ similarly, we have $\phi_1 c_{1,C}(s,t) = c_{1,G}(s,t)$. Then it is easily seen that the workload in the coupled-processors model can be written in terms of that in the GPS model:

$$\mathbb{P}\left(V_{1,C} > x\right) = \mathbb{P}\left(V_{1,G} > \phi_1 x\right).$$

Hence,

$$\begin{aligned} &\mathbb{P}\left(V_{1,C} > x\right) \\ \sim\ &\mathbb{P}\left(V_{1,B_{1,G}}^{\phi_1} > \phi_1 x\right) \frac{\rho_{2,G}}{1-\rho_{1,G}-\rho_{2,G}} \mathbb{P}\left(B_{2,G}^{\mathrm{r}} > \frac{\phi_1 x(\phi_2-\rho_{2,G})}{\hat{\rho}_{1,G}-\phi_1}\right). \end{aligned}$$

Using that

$$\mathbb{P}\left(V_{1,B_{1,G}}^{\phi_1} > \phi_1 x\right) = \mathbb{P}\left(V_{1,B_{1,C}}^{1} > x\right),$$

and

$$\mathbb{P}\left(B_{2,G}^{\mathrm{r}} > x\right) = \mathbb{P}\left(B_{2,C}^{\mathrm{r}} > \frac{x}{\phi_2}\right),$$

together with the fact that (see [126]) $\hat{\rho}_{1,G} = 1/\rho_{1,G}$, and therefore, $\hat{\rho}_{1,G} = \phi_1/\rho_{1,C}$, we obtain the result in Theorem 5.4.1 after some calculations.

The result for the case $\rho_2 > 1$ is presented in the next theorem.

Theorem 5.4.2 *If $\mathbb{P}(B_1 > t)$ is exponential, and $\mathbb{P}(B_2 > t) \in \mathcal{R}_{-\nu_2}$, $1 < \nu_2 < 2$, as given in (5.35), and if $\rho_1 < 1$, $\rho_2 > 1$, then*

$$\mathbb{P}\left(V_1 > t\right) \sim H_2 \rho_1 e^{(\lambda_1 - 1/\beta_1)t} \sim H_2 \mathbb{P}\left(V_1^1 > t\right), \quad t \to \infty,$$

with

$$H_2 := \frac{\rho_2 - 1}{r_2 - 1} \frac{1}{1-\rho_1} = \frac{\rho_2 - 1}{K_2 - 1}.$$

Remark 5.4.1 *It follows from the discussion on p. 316 of [42] that $K_2 > \rho_2$ if the system is ergodic, as may also be argued from the interpretation of K_2 given earlier, which implies that $H_2 < 1$. This is consistent with the observation that $\mathbb{P}(V_1 > t) \le \mathbb{P}(V_1^1 > t)$ for all $t \ge 0$ if $\rho_1 < 1$, since Q_1 is guaranteed to receive a minimum service rate of 1 when backlogged. Note that the latter observation is also in agreement with Theorems 5.3.1, 5.3.2, and 5.4.1.*

5.4.2 Discussion and interpretation

We start with a heuristic interpretation of the result for $\rho_2 < 1$. Invoking a result in [39], the result in Theorem 5.4.1 may be rewritten as

$$\mathbb{P}(V_1 > t) \sim \mathbb{P}(V_1^1 > t)\,\mathbb{P}\left(V_2^{K_2} > \frac{1-\rho_2}{\hat{\rho}_1 - 1} t\right), \qquad t \to \infty, \tag{5.36}$$

with $\hat{\rho}_1 := 1/\rho_1 > 1$. To understand the above formula, it is useful to draw a comparison with the workload V_1^1 when Q_1 is served in isolation at constant rate 1. Large-deviations results for the M/M/1 queue [126] suggest that the most likely way for V_1^1 to reach a large level x is that class 1 temporarily experiences 'abnormal' traffic activity. Specifically, class 1 essentially behaves as if its traffic intensity were increased from the normal value ρ_1 to the value $\hat{\rho}_1$, causing a positive drift $\hat{\rho}_1 - 1 > 0$ in the workload. In order for the workload to reach a large level x, the deviant behavior must persist for a period of time $x/(\hat{\rho}_1 - 1)$.

Now observe that the above scenario may occur in the shared system as well, provided Q_2 is continuously backlogged while class 1 shows deviant behavior. In fact, given that the workload at Q_1 reaches a large level, Q_2 is constantly backlogged with overwhelming probability, since otherwise class 1 must show even greater anomalous activity. The probability of that happening is negligibly small compared to that of Q_2 being continuously backlogged, because of the highly bursty nature of class-2 traffic and relatively smooth behavior of class-1 traffic.

Note that the normal drift in the workload at Q_2 is $\rho_2 - 1 < 0$. Thus, in order for Q_2 to remain backlogged during the period of deviant behavior of class 1, an additional amount of work of at least $(1-\rho_2)x/(\hat{\rho}_1 - 1)$ must be accounted for. The most likely scenario is that class 2 generates a large amount of traffic prior to the deviant behavior of class 1, so that when that starts, the workload at Q_2 is at least $(1-\rho_2)x/(\hat{\rho}_1 - 1)$. During that prior period, class 1 shows average behavior, and Q_1 does not receive any surplus capacity from Q_2, so that Q_1 is empty a fraction $1-\rho_1$ of the time. Thus, the average service rate at Q_2 during that period is $K_2 = \rho_1 + (1-\rho_1)r_2$. Hence, the probability that Q_2 is sufficiently long backlogged is approximately equal to $\mathbb{P}(V_2^{K_2} > (1-\rho_2)x/(\hat{\rho}_1 - 1))$. Combined, these considerations yield (5.36).

It is instructive to observe the similarities with the intuitive interpretation of Theorem 4.2.1 in Section 4.2.2.

We now turn to the case $\rho_2 > 1$. In this case, it is relatively likely for Q_2 to be continuously backlogged during the period in which class 1 shows deviant behavior. This suggests that the asymptotic behavior of $\mathbb{P}(V_1 > t)$ should not fundamentally change, except that the pre-factor of $\mathbb{P}(V_1^1 > t)$ should be O(1), as is confirmed by Theorem 5.4.2.

5.4.3 Proofs

Proof of Theorem 5.4.1 The approach may be outlined as follows. We are going to apply Theorem 2.2.2 with $f(t) = \mathbb{P}(V_1 > t)$, and hence with Laplace transform $\phi(s) = (1 - \mathbb{E}e^{-sV_1})/s$. From Equations (5.3) and (5.4), we see that $\mathbb{E}e^{-sV_1}$ has a singularity in $s = s_1 = \lambda_1 - 1/\beta_1 < 0$. The fact that the workload of class 1 cannot exceed that when served in isolation implies $\mathbb{P}(V_1 > t) \leq \mathbb{P}(V_1^1 > t)$, so that $\mathbb{E}e^{-sV_1}$ does not have any singularities right of the singularity $s = s_1$ of $\mathbb{E}e^{-sV_1^1}$. In applying Theorem 2.2.2 we thus need to consider

$$(s - s_1)\phi(s) = (s - s_1)\frac{1 - \mathbb{E}e^{-sV_1}}{s}, \tag{5.37}$$

and determine the series expansion for s near s_1.

Substituting (5.5) into (5.37), we obtain

$$\begin{aligned}(s - s_1)\frac{1 - \mathbb{E}e^{-sV_1}}{s} &= \rho_1 + (r_1 - 1)\frac{1 + \beta_1 s}{\beta_1 s}(\psi_2(s) - \psi_2(0)) \\ &= \rho_1 + (r_1 - 1)\frac{1 + \beta_1 s}{\beta_1 s}(\psi_2(s_1) - \psi_2(0)) \\ &\quad + (r_1 - 1)\frac{1 + \beta_1 s}{\beta_1 s}(\psi_2(s) - \psi_2(s_1)).\end{aligned} \tag{5.38}$$

Below we will show that $\psi_2(s_1) - \psi_2(0) = (1 - \rho_1)/(r_1 - 1)$. But first we turn to the most involved part of the proof: determining the series expansion of the term $\psi_2(s) - \psi_2(s_1)$ in (5.38) for s near s_1. Note that we cannot use (5.6), since

$$w = \delta_1^{-1}(s) = s - \frac{\rho_1 s}{1 + \beta_1 s} = \frac{\beta_1 s}{1 + \beta_1 s}(s - s_1) \sim \frac{\rho_1 - 1}{\rho_1}(s - s_1) \uparrow 0 \tag{5.39}$$

for $s \downarrow s_1$, while (5.6) is only valid for Re $w \geq 0$. In the appendix we prove that for $s \downarrow s_1$,

$$\psi_2(s) - \psi_0 \sim \frac{1}{r_1}\frac{\psi_0}{1 - \frac{1}{r_1} - \frac{1}{r_2}}\left(1 - e^{P_1(\delta_1^{-1}(s)) - P_2(\delta_1^{-1}(s))}\frac{r_1}{(r_1 - 1)r_2}\right). \tag{5.40}$$

Using (5.8), (5.39) and (5.40), $\psi_2(s) - \psi_2(s_1)$ is equal to, for $s \downarrow s_1$,

$$\begin{aligned}&\psi_2(s) - \psi_0 - (\psi_2(s_1) - \psi_0) \\ \sim\ & \frac{1}{(r_1 - 1)r_2}\frac{\psi_0}{1 - \frac{1}{r_1} - \frac{1}{r_2}}\left(e^{P_1(0) - P_2(0)} - e^{P_1(\delta_1^{-1}(s)) - P_2(\delta_1^{-1}(s))}\right) \\ =\ & \frac{1}{(r_1 - 1)r_2}\frac{e^{-P_2(0) - R_2(0)}}{1 - \frac{1}{r_1} - \frac{1}{r_2}}\left(1 - e^{P_1(\delta_1^{-1}(s)) - P_1(0) - (P_2(\delta_1^{-1}(s)) - P_2(0))}\right).\end{aligned} \tag{5.41}$$

Note that

$$P_i(0) + R_i(0) = \sum_{n=1}^{\infty} \frac{b_i^n}{n} = -\ln(1 - b_i), \tag{5.42}$$

and $e^x = 1 + x + O\left(x^2\right)$ for x small. Hence, for $s \downarrow s_1$,

$$\begin{aligned}&\psi_2(s) - \psi_2(s_1)\\ &\sim \frac{1}{(r_1-1)r_2} \frac{1-b_2}{1-\frac{1}{r_1}-\frac{1}{r_2}} \left(P_2(\delta_1^{-1}(s)) - P_2(0) - (P_1(\delta_1^{-1}(s)) - P_1(0))\right). \end{aligned} \tag{5.43}$$

We now determine the behavior of $P_2(w) - P_2(0) - (P_1(w) - P_1(0))$ for $w \uparrow 0$, which corresponds to $s \downarrow s_1$. It is easily seen that $P_i(w)$ is the LST of

$$p_i(t) := \sum_{n=1}^{\infty} \frac{b_i^n}{n} \mathbb{P}\left(-t < \sigma_n^{(i)} < 0\right).$$

Using Equations (5.12), (5.13), (5.14), we have

$$P_i(0) - p_i(t) = \sum_{n=1}^{\infty} \frac{b_i^n}{n} \mathbb{P}\left(-\sigma_n^{(i)} > t\right) =$$

$$\sum_{n=1}^{\infty} \frac{b_i^n}{n} \sum_{k=0}^{n} \binom{n}{k} \pi_i^k (1-\pi_i)^{n-k} \mathbb{P}\left(\sum_{j=1}^{n-k} \hat{P}_{2,j} - \sum_{j=1}^{k} \hat{P}_{1,j} > t\right). \tag{5.44}$$

From Formula (6.5) of [42], it follows that

$$\mathbb{E}e^{w\hat{P}_2} = \frac{1 - \beta_2\{\delta_2(w)\}}{\beta_2 \delta_2(w)}, \qquad \mathrm{Re}\, w \le 0. \tag{5.45}$$

Since $\mathbb{P}(B_2 > t)$ is regularly varying of index $-\nu_2$, we have, using Theorem 2.2.1,

$$1 - \frac{1-\beta_2\{s\}}{\beta_2 s} \sim \frac{C_2}{\beta_2} s^{\nu_2 - 1} l_2\left(\frac{1}{s}\right), \qquad s \downarrow 0. \tag{5.46}$$

In view of the symmetry between the two regularly-varying-tail assumptions and between the definitions of $\delta_1(w)$ and $\delta_2(w)$, it is readily seen from (5.21) that

$$\delta_2(w) \sim \frac{-w}{1-\rho_2}\left(1 - \frac{\lambda_2 C_2}{1-\rho_2}\left(\frac{-w}{1-\rho_2}\right)^{\nu_2-1} l_2(\frac{-1}{w})\right), \qquad w \uparrow 0. \tag{5.47}$$

Using Equations (5.45), (5.46), and (5.47),

$$1 - \mathbb{E}e^{w\hat{P}_2} \sim \frac{C_2}{\beta_2}\left(\frac{-w}{1-\rho_2}\right)^{\nu_2-1} l_2\left(\frac{-1}{w}\right), \qquad w \uparrow 0. \tag{5.48}$$

Applying Theorem 2.2.1, we find that Relation (5.48) gives

$$\mathbb{P}\left(\hat{P}_2 > t\right) \sim \frac{C_2}{\beta_2 \Gamma(2-\nu_2)} \left((1-\rho_2)t\right)^{1-\nu_2} l_2(t), \qquad t \to \infty, \tag{5.49}$$

indicating that $\mathbb{P}(\hat{P}_2 > t)$ is regularly varying of index $1-\nu_2$. The latter fact, in combination with $\mathbb{P}(\hat{P}_1 < \infty) = 1$, implies (see Section 2.1 and [127]):

$$\mathbb{P}\left(\sum_{j=1}^{n-k} \hat{P}_{2,j} - \sum_{j=1}^{k} \hat{P}_{1,j} > t\right) \sim \mathbb{P}\left(\sum_{j=1}^{n-k} \hat{P}_{2,j} > t\right) \sim (n-k)\mathbb{P}\left(\hat{P}_2 > t\right), \; t \to \infty.$$

Using the above relation in (5.44), we obtain, for $t \to \infty$,

$$\begin{aligned} P_i(0) - p_i(t) &\sim \sum_{n=1}^{\infty} \frac{b_i^n}{n} \sum_{k=0}^{n} \binom{n}{k} \pi_i^k (1-\pi_i)^{n-k} (n-k) \mathbb{P}\left(\hat{P}_2 > t\right) \\ &= (1-\pi_i) \sum_{n=1}^{\infty} b_i^n \sum_{k=0}^{n-1} \binom{n-1}{k} \pi_i^k (1-\pi_i)^{n-k-1} \mathbb{P}\left(\hat{P}_2 > t\right) \\ &= (1-\pi_i) \sum_{n=1}^{\infty} b_i^n \mathbb{P}\left(\hat{P}_2 > t\right) \\ &= \frac{(1-\pi_i) b_i}{1-b_i} \mathbb{P}\left(\hat{P}_2 > t\right). \end{aligned} \tag{5.50}$$

Applying Theorem 2.2.1 (or rather a minor adaptation, since that theorem is formulated in terms of non-negative random variables), we deduce from (5.49) and (5.50):

$$P_i(w) - P_i(0) \sim -\frac{(1-\pi_i) b_i}{1-b_i} \frac{C_2}{\beta_2} \left(\frac{-w}{1-\rho_2}\right)^{\nu_2 - 1} l_2\left(\frac{-1}{w}\right), \qquad w \uparrow 0. \tag{5.51}$$

Using Equations (5.10), (5.11),

$$\frac{(1-\pi_2) b_2}{1-b_2} - \frac{(1-\pi_1) b_1}{1-b_1} = \frac{(1-\rho_1)\rho_2}{(1-b_1)(1-b_2)} \left(1 - \frac{1}{r_1} - \frac{1}{r_2}\right).$$

Combining the above two relations, we obtain, as $w \uparrow 0$,

$$\begin{aligned} &P_2(w) - P_2(0) - (P_1(w) - P_1(0)) \\ &\sim -\frac{(1-\rho_1)\rho_2}{(1-b_1)(1-b_2)} \left(1 - \frac{1}{r_1} - \frac{1}{r_2}\right) \frac{C_2}{\beta_2} \left(\frac{-w}{1-\rho_2}\right)^{\nu_2-1} l_2\left(\frac{-1}{w}\right). \end{aligned}$$

Substituting this into (5.43), we have, for $s \downarrow s_1$,

$$\begin{aligned} &\psi_2(s) - \psi_2(s_1) \\ &\sim -\frac{1}{(r_1-1)r_2} \frac{(1-\rho_1)\rho_2}{1-b_1} \frac{C_2}{\beta_2} \left(\frac{-\delta_1^{-1}(s)}{1-\rho_2}\right)^{\nu_2-1} l_2\left(\frac{-1}{\delta_1^{-1}(s)}\right). \end{aligned} \tag{5.52}$$

It remains to calculate the term $\psi_2(s_1) - \psi_2(0)$ in (5.38). From (5.6), noting that $\delta_1^{-1}(0) = 0$,

$$\psi_2(0) - \psi_0 = \frac{1}{r_1} \frac{\psi_0}{1 - \frac{1}{r_1} - \frac{1}{r_2}} \left(1 - e^{-R_1(0)+R_2(0)}\right). \tag{5.53}$$

Combining (5.40) with (5.53) and using (5.8) and (5.42),

$$\begin{aligned}
\psi_2(s_1) - \psi_2(0) &= \psi_2(s_1) - \psi_0 - (\psi_2(0) - \psi_0) \\
&= \frac{1}{r_1} \frac{\psi_0}{1 - \frac{1}{r_1} - \frac{1}{r_2}} \left(e^{-R_1(0)+R_2(0)} - e^{P_1(0)-P_2(0)} \frac{r_1}{(r_1-1)r_2}\right) \\
&= \frac{1}{r_1} \frac{1}{1 - \frac{1}{r_1} - \frac{1}{r_2}} \left(e^{-P_1(0)-R_1(0)} - e^{-P_2(0)-R_2(0)} \frac{r_1}{(r_1-1)r_2}\right) \\
&= \frac{1}{r_1} \frac{1}{1 - \frac{1}{r_1} - \frac{1}{r_2}} \left(1 - b_1 - \frac{(1-b_2)r_1}{(r_1-1)r_2}\right) \\
&= \frac{1-\rho_1}{r_1 - 1}.
\end{aligned} \tag{5.54}$$

Substituting (5.52) and (5.54) into (5.38) with

$$\delta_1^{-1}(s) \sim \frac{\rho_1 - 1}{\rho_1}(s - s_1),$$

as in (5.39), and observing that

$$\frac{1 + \beta_1 s}{\beta_1 s} \to \frac{\rho_1}{\rho_1 - 1}, \qquad s \downarrow s_1,$$

we obtain

$$\begin{aligned}
&(s - s_1)\frac{1 - \mathbb{E}e^{-sV_1}}{s} \\
&= \frac{\rho_1 \rho_2}{(1-b_1)r_2} \frac{C_2}{\beta_2} \left(\frac{1-\rho_1}{\rho_1(1-\rho_2)}(s - s_1)\right)^{\nu_2 - 1} l_2\left(\frac{1}{s - s_1}\right) + \mathrm{O}(s - s_1).
\end{aligned}$$

In terms of Theorem 2.2.2, we now have a series expansion of $(s - s_1)\phi(s)$ with $a_{01} = 0$, $\gamma_1 = \nu_2 - 1$, and

$$b_{01} = \frac{\rho_1 \rho_2}{(1-b_1)r_2} \frac{C_2}{\beta_2} \left(\frac{1-\rho_1}{\rho_1(1-\rho_2)}\right)^{\nu_2 - 1}.$$

Applying Theorem 2.2.2, using that

$$\Gamma(\nu_2 - 1)\frac{\sin(\pi(\nu_2 - 1))}{\pi} = \frac{1}{\Gamma(2 - \nu_2)}$$

and $(1 - b_1)r_2 = K_2 - \rho_2$, then gives the statement of the theorem.

There is still one problem left in applying Theorem 2.2.2: verifying Condition (iii). It follows from (5.38) and (5.54) that

$$\phi(s) = \frac{1}{s} + (r_1 - 1)\frac{1 + \beta_1 s}{\beta_1 s}\frac{\psi_2(s) - \psi_2(s_1)}{s - s_1}. \tag{5.55}$$

The condition that $\int |\phi(s)| \mathrm{d}y$ converges at $y = \pm\infty$ is not fulfilled. However, as observed in [36], a detailed analysis of Sutton's proof of Theorem 2.2.2 reveals that this convergence condition is not strictly necessary. It is sufficient that $\phi(s) = \phi(x + \mathrm{i}y)$ and its derivative w.r.t. y are bounded in absolute value by L_1/y and L_2/y^2, respectively, for x near s_1 and for y sufficiently large, for some positive constants L_1 and L_2. It easily follows from (5.55) that these conditions are indeed satisfied. □

Remark 5.4.2 *In deriving the coefficient b_{01}, we have disregarded the slowly varying function $l_2(\cdot)$. Formally, this causes a problem with the application of Theorem 2.2.2. The treatment of that – technical – problem is outside the scope of this monograph. The problem does not arise in the slightly less general case of regular variation with $l_2(\cdot) \equiv 1$.*

Proof of Theorem 5.4.2 The proof is largely similar to that of Theorem 5.4.1. As in the proof of Theorem 5.3.2, the main difference between this proof and the proof of Theorem 5.4.1, is that $\mathbb{P}(\hat{P}_2 < \infty) = 1/\rho_2$ (the distribution of $\hat{P}_2$ is defective). Using (5.9), (5.12), (5.13), (5.14), we now have, instead of (5.51),

$$P_i(w) - P_i(0) \to -\sum_{n=1}^{\infty} \frac{b_i^n}{n} \mathbb{P}\left(\sigma_n^{(i)} = -\infty\right), \qquad w \uparrow 0,$$

with

$$\mathbb{P}\left(\sigma_n^{(i)} = -\infty\right) = \sum_{k=0}^{n} \binom{n}{k} \pi_i^k (1-\pi_i)^{n-k} \mathbb{P}\left(\sum_{j=1}^{n-k} \hat{P}_{2,j} - \sum_{j=1}^{k} \hat{P}_{1,j} = \infty\right). \tag{5.56}$$

Since $\mathbb{P}(\hat{P}_1 < \infty) = 1$, we may write

$$\mathbb{P}\left(\sum_{j=1}^{n-k} \hat{P}_{2,j} - \sum_{j=1}^{k} \hat{P}_{1,j} = \infty\right) = 1 - \left(\mathbb{P}\left(\hat{P}_2 < \infty\right)\right)^{n-k} = 1 - \left(\frac{1}{\rho_2}\right)^{n-k}.$$

Using the above relation in (5.56), we obtain

$$P_i(w) - P_i(0) \to -\ln\left(\frac{\rho_2 - \rho_2 b_i \pi_i - b_i + b_i \pi_i}{\rho_2(1 - b_i)}\right), \qquad w \uparrow 0,$$

yielding

$$e^{P_2(w) - P_2(0) - (P_1(w) - P_1(0))} \to \frac{-r_2 + \rho_1 r_2 - \rho_1 + \rho_2}{(-r_1 + \rho_2 r_1 - \rho_2 + \rho_1)(r_2 - 1)}, \qquad w \uparrow 0.$$

Substituting this into (5.41),

$$\psi_2(s) - \psi_2(s_1) \to -\frac{\rho_2 - 1}{(r_1 - 1)(r_2 - 1)}, \qquad s \downarrow s_1. \tag{5.57}$$

Substituting (5.57) and (5.54) into (5.38), and observing that

$$\frac{1 + \beta_1 s}{\beta_1 s} \to \frac{\rho_1}{\rho_1 - 1}, \qquad s \downarrow s_1,$$

we obtain

$$(s - s_1)\frac{1 - \mathbb{E}e^{-sV_1}}{s} = \rho_1 H_2 + \mathrm{O}\left((s - s_1)^{\nu_2 - 1}\right).$$

In terms of Theorem 2.2.2, we now have a series expansion of $(s - s_1)\phi(s)$ with $a_{01} = \rho_1 H_2$ and $\gamma_1 = \nu_2 - 1$. Finally, applying Theorem 2.2.2 completes the proof. □

5.5 Q1 heavy-tailed and Q2 heavy-tailed: overload

5.5.1 Assumptions and main result

In this section we consider the case where the service time distribution at both Q_1 and Q_2 is regularly varying of index $-\nu_1$ and $-\nu_2$, respectively. We assume that $\rho_1 > 1$ so that the lower service speed alone is not sufficient to ensure stability of Q_1. Thus, Q_1 relies on surplus capacity from Q_2, which forces $\rho_2 < 1$. The next theorem shows that the workload distribution at Q_1 is then determined by the heaviest of the service time distributions at Q_1 and Q_2 (Theorem 6.2 in [20]).

Theorem 5.5.1 *If $\mathbb{P}(B_i > t) \in \mathcal{R}_{-\nu_i}$ as given in (5.1), $1 < \nu_i < 2$, $i = 1, 2$, and if $\rho_1 > 1$, $\rho_2 < 1$, then $\mathbb{P}(V_1 > t) \in \mathcal{R}_{1-\min\{\nu_1,\nu_2\}}$:*
If $\nu_1 < \nu_2$, then

$$\begin{aligned}\mathbb{P}(V_1 > t) &\sim \frac{1}{K_1 - \rho_1}\frac{\lambda_1 C_1}{\Gamma(2 - \nu_1)} t^{1-\nu_1} l_1(t) \\ &\sim \frac{\rho_1}{K_1 - \rho_1}\mathbb{P}(B_1^{\mathrm{r}} > t), \qquad t \to \infty;\end{aligned} \tag{5.58}$$

If $\nu_1 > \nu_2$, then

$$\begin{aligned}\mathbb{P}(V_1 > t) &\sim \frac{r_1 - 1}{K_1 - \rho_1}\frac{\lambda_2 C_2}{\Gamma(2 - \nu_2)}\left(\frac{1 - \rho_2}{\rho_1 - 1}t\right)^{1-\nu_2} l_2(t) \\ &\sim \frac{(r_1 - 1)\rho_2}{K_1 - \rho_1}\mathbb{P}\left(B_2^{\mathrm{r}} > \frac{1 - \rho_2}{\rho_1 - 1}t\right), \qquad t \to \infty;\end{aligned} \tag{5.59}$$

If $\nu_1 = \nu_2 = \nu$, then

$$\mathbb{P}(V_1 > t) \sim \frac{1}{K_1 - \rho_1}\frac{\lambda_1 C_1}{\Gamma(2 - \nu)} t^{1-\nu} l_1(t)$$

$$+\frac{r_1-1}{K_1-\rho_1}\frac{\lambda_2 C_2}{\Gamma(2-\nu)}\left(\frac{1-\rho_2}{\rho_1-1}t\right)^{1-\nu} l_2(t), \qquad t\to\infty. \tag{5.60}$$

Here $K_1 := \rho_2 + (1-\rho_2)r_1$ represents the average service rate at Q_1 when continuously backlogged.

Remark 5.5.1 *Qualitatively, the results for $\nu_1 < \nu_2$ and $\nu_1 > \nu_2$ in Theorem 5.5.1 are similar to Theorems 4.1 and 5.1 in [21], respectively.*

5.5.2 Discussion and interpretation

Theorem 5.5.1 shows that the workload distribution at Q_1 is determined by the heaviest of the service time distributions at Q_1 and Q_2. If the tail of B_1 is heavier than that of B_2, then the workload at Q_1 behaves exactly as if Q_1 is served in isolation at a constant rate K_1. The explanation is largely similar to that of Theorem 5.3.1 provided in Section 5.3.2. Specifically, the most likely scenario for the workload at Q_1 to reach a large level is again that a class-1 customer arrives with a large service time, so that Q_1 becomes backlogged for a long period of time. As a result, Q_2 will not receive any surplus capacity, and hence be empty only a fraction $1-\rho_2$ of the time. Thus, the average service rate at Q_1 while backlogged is $K_1 = \rho_2 + (1-\rho_2)r_1$ as before. Since $\rho_1 > 1$, a large workload at Q_1 may also occur when Q_2 is backlogged for a long period of time. The latter scenario is however far less likely when the tail of B_2 is lighter than that of B_1.

In contrast, if the tail of B_2 is heavier than that of B_1, then the tail of B_2 determines that of the workload distribution at Q_1. Using a result of [47] (see also (5.20)), we can rewrite (5.59) for $\nu_1 > \nu_2$ as

$$\mathbb{P}(V_1 > t) \sim (1-\rho_2)\frac{(r_1-1)\rho_2}{K_1-\rho_1}\mathbb{P}\left(P_2^{\mathrm{r}} > \frac{t}{\rho_1-1}\right), \tag{5.61}$$

with P_2 a random variable representing the busy period at Q_2 when served in isolation at unit rate. This may be heuristically explained as follows. Large-deviations arguments suggest that the most likely scenario for the workload at Q_1 to reach a large level is that class 2 generates a large amount of traffic, while class 1 *itself* shows average behavior. Specifically, suppose that a class-2 customer arrives with a large service time, so that Q_2 becomes backlogged for a long period of time. Then Q_1 is only served at unit rate, while class 1 generates traffic at an average rate $\rho_1 > 1$. Thus, the workload at Q_1 has a positive drift, and Q_1 will soon become backlogged too (if it not already is), and remain so for as long as Q_2 is backlogged. Consequently, Q_2 will experience a busy period as if it were served at unit rate. During that period, the workload at Q_1 will roughly grow at rate $\rho_1 - 1 > 0$. Only after Q_2 empties, Q_1 will start to drain again at approximately rate $r_1 - \rho_1$. Of course, the workload at Q_1 may also build up when class 1 *itself* generates a large amount of traffic. However, this scenario is highly implausible compared to the build-up that occurs during a busy period at Q_2, because of the relatively smooth nature of class-1 traffic.

Thus, the workload at Q_1 basically behaves as that in a queue of capacity $r_1 - \rho_1$ which is fed by an on-off source with as on and off periods the busy and idle periods at Q_2, respectively, and with inflow rate $r_1 - 1 > 0$ when on, and fraction of off time $1 - \rho_2$. In particular, the traffic intensity of the on-off source equals $(r_1 - 1)\rho_2$. Using a result of [67], it may be verified that the workload behavior in this queue is exactly as indicated by (5.61).

Remark 5.5.2 *Interestingly, the high service speed r_2 does not show up in either Theorem 5.3.1, Theorem 5.3.2, or Theorem 5.5.1 at all. In particular, the results coincide with those in [19] for the case $r_2 = 1$. This may be explained by observing that if the workload at Q_1 is large, then Q_2 operates at the lower service speed all the time, so that the high service speed is irrelevant. Similarly, the high service speed r_1 does not show up in either Theorem 5.4.1 or Theorem 5.4.2. In order for the workload at Q_1 to be large, Q_1 must have operated at the lower service speed all the time.*

5.5.3 Proof

In this section we present the proof of Theorem 5.5.1. We give the entire proof because the proof of Theorem 5.6.1 will be similar. Some parts of the proof are also used in the proof of Theorem 5.4.1.

Proof of Theorem 5.5.1 Starting point for studying the tail distribution of the workload V_1 is again Relation (5.5) for its LST, but we can no longer use (5.6) for the term $\psi_2(s)$. The reason for this is the following. We want to let $s \to 0$, but $\delta_1(w) \to \delta_1(0) \neq 0$ for $w \to 0$ if $\rho_1 > 1$. Let us therefore take a closer look at the zeroes of $f_1(s, w)$, cf. (5.15). In [42] it is observed that $\frac{\mathrm{d}}{\mathrm{d}s} f_1(s, w)$ has, for real $s \geq 0$, no zero if $\rho_1 < 1$, one zero $s_0 = 0$ if $\rho_1 = 1$, and one zero $s_0 > 0$ if $\rho_1 > 1$. If $\rho_1 \geq 1$, then the point $w_0 := s_0 - \lambda_1(1 - \beta_1\{s_0\})$ is a second-order branch-point of the analytic continuation of $\delta_1(w)$, Re $w \geq 0$, into Re $w < 0$. For $\rho_2 < 1$, $\rho_1 \geq 1$, and $w \in [w_0, 0]$, the *two* zeroes of $f_1(s, w)$ in $[0, \delta_1(0)]$ will be denoted by $\epsilon_1(w)$ and $\delta_1(w)$, such that:

- $\epsilon_1(w)$ maps $[w_0, 0]$ one-to-one onto $[0, s_0]$;
- $\delta_1(w)$ maps $[w_0, 0]$ one-to-one onto $[s_0, \delta_1(0)]$.

If $\rho_1 \geq 1$, then (6.24) of [42] yields

$$\left((1 - \frac{1}{r_1})\frac{w}{\delta_2(w)} - \frac{1}{r_1}\frac{w}{\epsilon_1(w)}\right)\left(\frac{1}{r_2}(\psi_2(\epsilon_1(w)) - \psi_0) - \frac{\psi_0}{r_1 r_2}\frac{1}{1 - \frac{1}{r_1} - \frac{1}{r_2}}\right) =$$
$$-\left((1 - \frac{1}{r_2})\frac{w}{\epsilon_1(w)} - \frac{1}{r_2}\frac{w}{\delta_2(w)}\right)\left(\frac{1}{r_1}(\psi_1(\delta_2(w)) - \psi_0) - \frac{\psi_0}{r_1 r_2}\frac{1}{1 - \frac{1}{r_1} - \frac{1}{r_2}}\right). \tag{5.62}$$

To determine the behavior of $\psi_2(\epsilon_1(w))$ for $w \uparrow 0$ (which eventually will give us the behavior of $\mathbb{E}e^{-sV_1}$ for $s \downarrow 0$, and hence that of $\mathbb{P}(V_1 > t)$ for $t \to \infty$), we need to determine the behavior, for $w \uparrow 0$, of $\epsilon_1(w)$, $\delta_2(w)$ and $\psi_1(\delta_2(w))$ – the terms that appear in (5.62). Take $w < 0$, $w \uparrow 0$, then (cf. (5.21)),

$$\epsilon_1(w) + \frac{w}{\rho_1 - 1} \sim \frac{\lambda_1 C_1}{\rho_1 - 1}\left(\frac{-w}{\rho_1 - 1}\right)^{\nu_1} l_1\left(\frac{-1}{w}\right), \quad w \uparrow 0. \tag{5.63}$$

As derived in (5.47) in the proof of Theorem 5.4.1,

$$\delta_2(w) + \frac{w}{1 - \rho_2} \sim -\frac{\lambda_2 C_2}{1 - \rho_2}\left(\frac{-w}{1 - \rho_2}\right)^{\nu_2} l_2\left(\frac{-1}{w}\right), \quad w \uparrow 0. \tag{5.64}$$

For $\rho_2 < 1$, $\psi_1(\delta_2(w))$ is specified by (6.23) of [42] (notice the symmetry with (5.6)),

$$\frac{1}{r_1}\left(\psi_1(\delta_2(w)) - \psi_0\right) = \frac{1}{r_1 r_2}\frac{\psi_0}{1 - \frac{1}{r_1} - \frac{1}{r_2}}\left(1 - e^{P_1(w) - P_2(w)}\right), \ \mathrm{Re}\, w \le 0. \tag{5.65}$$

We now turn to the study of $P_i(w)$. Compared to the study of $R_i(w)$, there is a slight difference. Now that $\rho_1 > 1$, the distribution of the busy period $\hat{P}_1$ is defective: $\mathbb{P}(\hat{P}_1 < \infty) = 1/\rho_1$. As in (5.31), (5.32), we obtain, with

$$p_i(t) := \sum_{n=1}^{\infty} \frac{b_i^n}{n} \mathbb{P}\left(-t < X_{i1} + \cdots + X_{in} < 0\right),$$

$$P_2(0) - p_2(t) \sim \frac{b_2(1 - \pi_2)}{1 - b_2\left(\frac{\pi_2}{\rho_1} + 1 - \pi_2\right)} \mathbb{P}\left(\hat{P}_2 > t\right) = \frac{\rho_2}{1 - \rho_2}\mathbb{P}\left(\hat{P}_2 > t\right).$$

Similarly,

$$P_1(0) - p_1(t) \sim \frac{b_1(1 - \pi_1)}{1 - b_1\left(\frac{\pi_1}{\rho_1} + 1 - \pi_1\right)} \mathbb{P}\left(\hat{P}_2 > t\right) = \frac{\rho_2}{1 - \rho_2}\mathbb{P}\left(\hat{P}_2 > t\right).$$

Using the counterpart of (5.22) for $\hat{P}_2$, and Theorem 2.2.1, it finally follows that (cf. (5.33)),

$$P_1(w) - P_1(0) \sim P_2(w) - P_2(0) \sim -\frac{\lambda_2 C_2}{1 - \rho_2}\left(\frac{-w}{1 - \rho_2}\right)^{\nu_2 - 1} l_2\left(\frac{-1}{w}\right), \quad w \uparrow 0.$$

Substituting into (5.65), we have

$$\psi_1(\delta_2(w)) - \psi_1(0) = \mathrm{o}\left(w^{\nu_2 - 1} l_2\left(\frac{-1}{w}\right)\right), \quad w \uparrow 0. \tag{5.66}$$

As a brief intermezzo, we make the following observation. From (5.66) it follows that $\mathbb{P}(V_1 = 0, V_2 > t) = \mathrm{o}(t^{1-\nu_2} l_2(t))$, $t \to \infty$. This may be surprising in view of the fact that if Q_2 were an M/G/1 queue in isolation, then $\mathbb{P}(V_2 > t) \sim C t^{1-\nu_2} l_2(t)$, cf. [39]. The explanation is the following. The workload at Q_1

has a positive drift $\rho_1 - 1$ when $V_2 > 0$. Therefore $\mathbb{P}(V_1 = 0|V_2 > t) = o(1)$ for $t \to \infty$: when the workload at Q_2 is very large, it is highly unlikely that Q_1 is empty.

The above result for the behavior of $\psi_1(\delta_2(w))$ for $w \uparrow 0$ allows us to determine the behavior of $\psi_2(\epsilon_1(w))$ for $w \uparrow 0$. Using (5.62) for the relation between $\psi_2(\epsilon_1(w))$ and $\psi_1(\delta_2(w))$, along with the asymptotic results (5.63) and (5.64) for $\epsilon_1(w)$ and $\delta_2(w)$ respectively, it follows after some calculations that

$$\psi_2(\epsilon_1(w)) - \psi_2(0) \sim -\frac{\rho_1 - 1}{(1-b_2)r_1}\lambda_2 C_2 \left(\frac{-w}{1-\rho_2}\right)^{\nu_2 - 1} l_2\left(\frac{-1}{w}\right) \mathrm{I}_{\{\nu_1 > \nu_2\}}$$
$$- \frac{(1-\rho_2)}{(1-b_2)r_1}\lambda_1 C_1 \left(\frac{w}{1-\rho_1}\right)^{\nu_1 - 1} l_1\left(\frac{-1}{w}\right) \mathrm{I}_{\{\nu_1 < \nu_2\}}, \; w \uparrow 0.$$

Using (5.63) once more,

$$\psi_2(s) - \psi_2(0) \sim -\frac{\rho_1 - 1}{(1-b_2)r_1}\lambda_2 C_2 (s\frac{\rho_1 - 1}{1-\rho_2})^{\nu_2 - 1} l_2(\frac{1}{s}) \mathrm{I}_{\{\nu_1 > \nu_2\}}$$
$$- \frac{(1-\rho_2)}{(1-b_2)r_1}\lambda_1 C_1 s^{\nu_1 - 1} l_1(\frac{1}{s}) \mathrm{I}_{\{\nu_1 < \nu_2\}}, \quad s \downarrow 0. \quad (5.67)$$

Finally we are ready to determine the tail behavior of the workload V_1 at Q_1. The LST of V_1 is given by (5.5). The first factor in its right-hand side is the LST of the workload distribution at Q_1 when served in isolation at unit speed (the Pollaczek-Khintchine workload LST in the M/G/1 queue). This factor would give a $t^{1-\nu_1}$ tail behavior, cf. [39]. Using (5.67), the second factor in the right-hand side of (5.5) is seen to yield either a $t^{1-\nu_1}$ or a $t^{1-\nu_2}$ tail behavior. To see which of these two terms dominates, we have to distinguish between three cases: (i) $\nu_1 < \nu_2$, (ii) $\nu_1 > \nu_2$ and (iii) $\nu_1 = \nu_2 =: \nu$.

(i) $\nu_1 < \nu_2$. In this case the heavier tail of B_1 dominates, and Formula (5.30) still holds,

$$\mathbb{E}e^{-sV_1} - 1 \sim -\frac{\lambda_1 C_1}{K_1 - \rho_1} s^{\nu_1 - 1} l_1\left(\frac{1}{s}\right), \quad s \downarrow 0, \quad (5.68)$$

with $K_1 = \rho_2 + (1-\rho_2)r_1$.

(ii) $\nu_1 > \nu_2$. In this case the heavier tail of B_2 dominates, resulting in

$$\mathbb{E}e^{-sV_1} - 1 \sim -\frac{r_1 - 1}{K_1 - \rho_1}\lambda_2 C_2 \left(s\frac{\rho_1 - 1}{1-\rho_2}\right)^{\nu_2 - 1} l_2\left(\frac{1}{s}\right), \; s \downarrow 0. \quad (5.69)$$

(iii) $\nu_1 = \nu_2 = \nu$. In this case, addition of the right-hand sides of (5.68) and (5.69) gives the right asymptotic behavior of $\mathbb{E}e^{-sV_1} - 1$.

The proof of the theorem is completed by yet another application of Theorem 2.2.1. □

5.6 One queue heavy-tailed, other lighter tail: overload

5.6.1 Introduction and main result

In this section we consider an extension of the model analyzed in the previous section. We show that the results obtained there, in fact continue to hold when the service time distribution at one of the queues is regularly varying and the service time distribution at the other queue has a lighter, general tail. More precisely, we consider the following two cases. For case (i) we assume the service times at Q_1 to be regularly varying of index $-\nu_1$, $1 < \nu_1 < 2$, see (5.17) for the distribution function and (5.18) for the behavior of its LST. The service times at Q_2 have a lighter tail than those at Q_1, more precisely, we assume that $\mathbb{P}(B_2 > t) = o(t^{-\nu_1} l_1(t))$ for $t \to \infty$. For case (ii) we assume the service times at Q_2 to be regularly varying of index $-\nu_2$, $1 < \nu_2 < 2$, similar to B_1 in case (i), and the service times at Q_1 to have a lighter tail, i.e., $\mathbb{P}(B_1 > t) = o(t^{-\nu_2} l_2(t))$ for $t \to \infty$. As in the previous section, we assume $\rho_1 > 1$ and $\rho_2 < 1$, so the lower speed at Q_2 is sufficient for stability, whereas Q_1 needs the surplus capacity from Q_2 to remain stable. We show that (5.58) still holds for case (i) and that (5.59) remains valid for case (ii).

Theorem 5.6.1 *For $\rho_1 > 1$, $\rho_2 < 1$:*
(i) if $\mathbb{P}(B_1 > t) \in \mathcal{R}_{-\nu_1}$, $1 < \nu_1 < 2$, and $\mathbb{P}(B_2 > t) = o(t^{-\nu_1} l_1(t))$, then

$$\begin{aligned}\mathbb{P}(V_1 > t) &\sim \frac{1}{K_1 - \rho_1} \frac{\lambda_1 C_1}{\Gamma(2-\nu_1)} t^{1-\nu_1} l_1(t) \\ &\sim \frac{\rho_1}{K_1 - \rho_1} \mathbb{P}(B_1^{\mathrm{r}} > t), \qquad t \to \infty;\end{aligned}$$

(ii) if $\mathbb{P}(B_2 > t) \in \mathcal{R}_{-\nu_2}$, $1 < \nu_2 < 2$, and $\mathbb{P}(B_1 > t) = o(t^{-\nu_2} l_2(t))$, then

$$\begin{aligned}\mathbb{P}(V_1 > t) &\sim \frac{r_1 - 1}{K_1 - \rho_1} \frac{\lambda_2 C_2}{\Gamma(2-\nu_2)} \left(\frac{1-\rho_2}{\rho_1 - 1} t\right)^{1-\nu_2} l_2(t) \\ &\sim \frac{(r_1 - 1)\rho_2}{K_1 - \rho_1} \mathbb{P}\left(B_2^{\mathrm{r}} > \frac{1-\rho_2}{\rho_1 - 1} t\right), \qquad t \to \infty.\end{aligned}$$

Just as in the previous section, the tail of the service time distribution of one queue dominates that of the other. Therefore, the interpretation and explanation of the above results are the same as given in Section 5.5.2.

5.6.2 Proof

Proof of Theorem 5.6.1 The proof is similar to that of Theorem 5.5.1. Again we use (5.5) to determine the tail distribution of V_1. Also, because $\rho_1 > 1$ we have to use (6.24) of [42] (see (5.62)) to find the behavior of $\psi_2(\epsilon_1(w))$ for $w \uparrow 0$. This means that we have to determine for both cases the behavior of $\epsilon_1(w)$, $\delta_2(w)$ and $\psi_1(\delta_2(w))$ for $w \uparrow 0$.

Before going into detail, we derive some results which we need in the remainder of the proof. If we assume for the distribution of a generic service time B that

$$\mathbb{P}(B > t) = \mathrm{o}\left(t^{-\nu} l(t)\right), \quad t \to \infty, \quad 1 < \nu < 2,$$

then applying Theorem 2.2.1 with $D = 0$ gives the following behavior of its LST,

$$\beta\{s\} = 1 - \beta s + \mathrm{o}\left(s^{\nu} l\left(\frac{1}{s}\right)\right), \quad s \downarrow 0, \tag{5.70}$$

with β denoting the first moment of B. Now consider an isolated stable M/G/1 queue with generic service time B and arrival rate λ. Rewriting the well-known formula for the LST of the workload V^1, and using (5.70), we obtain

$$1 - \mathbb{E}e^{-sV^1} = \mathrm{o}\left(s^{\nu-1} l\left(\frac{1}{s}\right)\right), \quad s \downarrow 0. \tag{5.71}$$

Ad case (i). Similar to the proof of Theorem 5.5.1 (see (5.63)), we have

$$\epsilon_1(w) + \frac{w}{\rho_1 - 1} \sim \frac{\lambda_1 C_1}{\rho_1 - 1}\left(\frac{-w}{\rho_1 - 1}\right)^{\nu_1} l_1\left(\frac{-1}{w}\right), \quad w \uparrow 0.$$

The behavior of $\delta_2(w)$ is determined as follows. According to the definition of $f_2(s, w)$ in (5.16),

$$\delta_2(w) = \lambda_2(1 - \beta_2\{\delta_2(w)\}) - w.$$

Using (5.70) for B_2, it follows that

$$\delta_2(w) = \frac{-w}{1 - \rho_2} + \mathrm{o}\left((\delta_2(w))^{\nu_1} l_1\left(\frac{-1}{w}\right)\right), \quad w \uparrow 0,$$

and hence

$$\delta_2(w) = \frac{-w}{1 - \rho_2} + \mathrm{o}\left((-w)^{\nu_1} l_1\left(\frac{-1}{w}\right)\right), \quad w \uparrow 0. \tag{5.72}$$

It remains to determine the behavior of $\psi_1(\delta_2(w))$. Observe that $\psi_1(s) - \psi_1(0)$ is the LST of $\mathbb{P}(V_1 = 0, V_2 > t)$. Obviously,

$$\mathbb{P}\left(V_1 = 0, V_2 > t\right) \leq \mathbb{P}\left(V_2 > t\right).$$

Recall that Q_2 is guaranteed to receive a minimum service rate of 1 when backlogged. Hence, the following inequality holds,

$$\mathbb{P}\left(V_2 > t\right) \leq \mathbb{P}\left(V_2^1 > t\right).$$

Using (5.71) it is readily seen that

$$\psi_1(\delta_2(w)) - \psi_1(0) = \mathrm{o}\left(w^{\nu_1 - 1} l_1\left(\frac{-1}{w}\right)\right), \quad w \uparrow 0. \tag{5.73}$$

Now we have to substitute (5.63), (5.72) and (5.73) into (5.62) to find the behavior of $\psi_2(\epsilon_1(w))$ for $w \uparrow 0$. After some calculations we obtain

$$\psi_2(\epsilon_1(w)) - \psi_2(0) \sim -\frac{1-\rho_2}{(1-b_2)r_1}\lambda_1 C_1 \left(\frac{-w}{\rho_1 - 1}\right)^{\nu_1 - 1} l_1\left(\frac{-1}{w}\right), \qquad w \uparrow 0.$$

Using (5.63) once again results in

$$\psi_2(s) - \psi_2(0) \sim -\frac{1-\rho_2}{(1-b_2)r_1}\lambda_1 C_1 s^{\nu_1 - 1} l_1\left(\frac{1}{s}\right), \qquad s \downarrow 0.$$

Now we have gathered all the ingredients necessary in (5.5) to find the tail behavior of V_1, except the term $\mathbb{E}e^{-sV_1^1}$. Observing that this is the same factor as in the proof of Theorem 5.5.1, we obtain (5.68).

Ad Case (ii). Now we can use (5.64) as given in the proof of Theorem 5.5.1,

$$\delta_2(w) + \frac{w}{1-\rho_2} \sim -\frac{\lambda_2 C_2}{1-\rho_2}\left(\frac{-w}{1-\rho_2}\right)^{\nu_2} l_2\left(\frac{-1}{w}\right), \qquad w \uparrow 0.$$

To find the behavior of $\epsilon_1(w)$, we can apply the same method as we used to determine the behavior of $\delta_2(w)$ in case (i). According to the definition of $f_1(s, w)$ (see (5.15)),

$$\epsilon_1(w) = \lambda_1(1 - \beta_1\{\epsilon_1(w)\}) + w. \tag{5.74}$$

Using a second-order Taylor approximation for $\epsilon_1(w)$ around $w = 0$ (note that $\epsilon_1(0) = 0$) together with (5.70) in (5.74), we obtain

$$\epsilon_1(w) = \frac{-w}{\rho_1 - 1} + \mathrm{o}\left(w^{\nu_2} l_2\left(-\frac{1}{w}\right)\right), \qquad w \uparrow 0. \tag{5.75}$$

The behavior of $\psi_1(\delta_2(w))$ is the same as that in the proof of Theorem 5.5.1 (see (5.66)),

$$\psi_1(\delta_2(w)) - \psi_1(0) = \mathrm{o}\left(w^{\nu_2 - 1} l_2\left(-\frac{1}{w}\right)\right), \qquad w \uparrow 0.$$

Substituting (5.64), (5.66), and (5.75) into (5.62) to find the asymptotic behavior of $\psi_2(\epsilon_1(w))$ for $w \uparrow 0$, we obtain after some calculations:

$$\psi_2(\epsilon_1(w)) - \psi_2(0) \sim -\frac{\rho_1 - 1}{(1-b_2)r_1}\lambda_2 C_2 \left(\frac{-w}{1-\rho_2}\right)^{\nu_2 - 1} l_2\left(\frac{-1}{w}\right), \qquad w \uparrow 0.$$

Using (5.75) once again gives

$$\psi_2(s) - \psi_2(0) \sim -\frac{\rho_1 - 1}{(1-b_2)r_1}\lambda_2 C_2 \left(s\frac{\rho_1 - 1}{1-\rho_2}\right)^{\nu_2 - 1} l_2\left(\frac{1}{s}\right), \qquad s \downarrow 0.$$

Observing that the term $\mathbb{E}e^{-sV_1^1}$ satisfies (5.71), we obtain (5.69). □

Appendix

5.A Proof of Equation (5.40)

For convenience we restate Equation (5.40) as a lemma.
Lemma *For* $s \downarrow s_1$,

$$\psi_2(s) - \psi_0 \sim \frac{1}{r_1} \frac{\psi_0}{1 - \frac{1}{r_1} - \frac{1}{r_2}} \left(1 - e^{P_1(\delta_1^{-1}(s)) - P_2(\delta_1^{-1}(s))} \frac{r_1}{(r_1 - 1)r_2}\right).$$

Proof Using the fact that $\mathbb{P}(V_1 > t) \le \mathbb{P}(V_1^1 > t)$, it may be verified that $\psi_2(s)$ is analytic for Re $s > s_1$. Analytic continuation of Formula (6.2) in [42], taking $w = \delta_1^{-1}(s)$, then shows that for $s \downarrow s_1$,

$$\begin{aligned} &\psi_2(s) - \psi_0 \\ = \;& \frac{1}{r_1} \frac{\psi_0}{1 - \frac{1}{r_1} - \frac{1}{r_2}} \\ &+ \left(\psi_1(\delta_2(\delta_1^{-1}(s))) - \psi_0 - \frac{\psi_0}{r_2} \frac{1}{1 - \frac{1}{r_1} - \frac{1}{r_2}}\right) \frac{s + (r_2 - 1)\delta_2(\delta_1^{-1}(s))}{s(r_1 - 1) + \delta_2(\delta_1^{-1}(s))}. \end{aligned} \tag{5.76}$$

Using Relations (5.39), (5.47), for $s \downarrow s_1$,

$$\delta_2(\delta_1^{-1}(s)) \sim \delta_2\left(\frac{\rho_1 - 1}{\rho_1}(s - s_1)\right) \sim -\frac{\rho_1 - 1}{\rho_1(1 - \rho_2)}(s - s_1). \tag{5.77}$$

Taking $w = \delta_1^{-1}(s)$ in (5.7), we obtain, for $s \downarrow s_1$,

$$\begin{aligned} &\psi_1(\delta_2(\delta_1^{-1}(s))) - \psi_0 - \frac{\psi_0}{r_2} \frac{1}{1 - \frac{1}{r_1} - \frac{1}{r_2}} \\ = \;& -\frac{\psi_0}{r_2} \frac{1}{1 - \frac{1}{r_1} - \frac{1}{r_2}} e^{P_1(\delta_1^{-1}(s)) - P_2(\delta_1^{-1}(s))}. \end{aligned} \tag{5.78}$$

Substituting (5.77) and (5.78) into (5.76) completes the proof. □

Chapter 6

Tandem and priority queue with Gaussian inputs

The models that we consider in this chapter differ in two ways from the ones that we studied in the previous chapters. First of all, we do not study the GPS mechanism, and secondly, we make different assumptions on the input processes.

We assume the input flows to behave as Gaussian processes, i.e., the cumulative amount of traffic follows a Gaussian distribution. Such a Gaussian process can either have light-tailed or heavy-tailed properties, depending on the correlation structure associated with the variance function. In order to obtain the Gaussian equivalent of an input process, we use a Gaussian process whose variance function corresponds to the variance function of the given process. A well-known input process with light-tailed features is the so-called Ornstein-Uhlenbeck process, which is the Gaussian equivalent of an exponential on-off source. A frequently used Gaussian model with heavy-tailed characteristics is fractional Brownian motion (fBm), which has a variance function that increases at a rate faster than linear. See [4] for more examples. The Gaussian model is justified by several theoretical results. Among these we mention [131], in which the superposition of many heavy-tailed on-off sources is shown to converge to fractional Brownian motion (fBm) (after appropriate rescaling). It was recently shown in [45] that this convergence carries over to the queueing process, justifying the choice of fBm as a good approximation of traffic inputs.

This chapter focuses on a two-node feed-forward tandem network, and a two-queue priority system. A two-queue GPS system is considered in the next chapter. The reason for considering the tandem and priority model first is that they provide some ideas on how to analyze the GPS model. The techniques that we apply in this chapter and in the next chapter have so far only been used to find logarithmic asymptotics for the workload in a single queue in [4]. In [95, 96] nice heuristics are given for the workload in the priority model and the GPS model respectively, but they do not provide explicit formulas for the logarithmic asymptotics.

This chapter is completely based on [90]. Part of it has appeared in [93]. We extend the work in [4] and derive the logarithmic many-sources asymptotics for the workload distribution in a tandem network and a priority system. More specifically, we examine in detail the second queue in a two-node tandem network, and the low-priority queue in a two-queue priority system. It turns out that these models are closely related. The results for two queues can be trivially extended to a network with m queues in tandem and to a priority system with m queues, $m > 2$. In the tandem model we assume that n i.i.d. Gaussian sources feed into the first queue. Both queues have a service rate, which is scaled by n. In the priority model we assume all queues to be fed by n i.i.d. Gaussian sources, and the service rate to be scaled by n.

A vast body of results exists for single queues under the many-sources scaling. Most notably, under very mild conditions on the source behavior, it is possible to calculate the *exponential* decay of the probability that the queue (fed by n sources, and emptied at a rate nc) exceeds level nb. Early references in this large-deviations framework are the logarithmic asymptotics found in [32, 43]. We remark that exact asymptotics for the workload in a single queue with Gaussian inputs were recently found in [46]. The goal of the present chapter is to find similar asymptotics for the workload in tandem and priority queues.

The work in this chapter fits in the framework of a series of papers that have appeared recently, i.e., [4, 95, 96, 97]. These papers examine queues with Gaussian sources, using a generalized version of Schilder's theorem. We use this theorem to show that the heuristics in [95] are typically close, but that there is a gap with the exact solution. For both the tandem and priority queue we derive a lower bound on the decay rate of the overflow probability. In addition, we present an explicit condition under which this lower bound matches the exact value of the decay rate. Notice that lower bounds for the decay rate are of practical interest, as typically the network has to be designed such that overflow is sufficiently rare. We mention that for priority systems in discrete time, different bounds were found in [139].

This chapter is organized as follows. Section 6.1 introduces the tandem model, and presents preliminaries on (sample-path) large deviations. Section 6.2 analyzes the two-queue tandem system. The analysis is numerically illustrated in Section 6.3. Section 6.4 studies the priority system, addressing the decay rate of overflow in the low-priority queue.

6.1 Model and preliminaries

This section introduces the tandem model that is analyzed in Section 6.2. In addition, we present preliminaries on the many-sources scaling and large-deviations results.

6.1.1 Tandem model

Consider a two-queue tandem model, with service rate nc_1 for the first queue and nc_2 for the second queue. We assume that $c_1 > c_2$, in order to exclude the trivial case where the workload at the second queue cannot build up.

We consider n sources (whose characteristics are specified in Section 6.1.2) that feed into the first queue. Traffic of these sources that has been served at the first queue immediately flows into the second queue – we assume no additional sources to feed the second queue. We are interested in the steady-state probability of the buffer content of the second queue $Q_{2,n}$ exceeding a certain threshold nb, $b > 0$, when the number of sources gets large, or, more specifically, its logarithmic asymptotics:

$$J(b) := -\lim_{n\to\infty} \frac{1}{n} \log \mathbb{P}(Q_{2,n} \geq nb). \tag{6.1}$$

With abuse of terminology we often follow the convention of referring to the probability $\mathbb{P}(Q_{2,n} \geq nb)$ as the overflow probability, whereas it is formally an exceedance probability.

6.1.2 Gaussian processes

We assume the n sources to behave as i.i.d. Gaussian processes with stationary increments. Define $A_i(s,t)$ as the amount of traffic generated by the ith source in $(s,t]$, with $s < t$ and $s,t \in \mathbb{R}$. Denote by $A(s,t)$ the generic random variable corresponding to a single source, i.e., $A_i(s,t) \stackrel{\rm d}{=} A(s,t)$ (recall that $\stackrel{\rm d}{=}$ denotes equality in distribution). The Gaussian sources are characterized by their mean rate $\mu \geq 0$ and their variance function $v(\cdot)$. More precisely, for all s,t with $s < t$, $\mathbb{E}A(s,t) = \mu(t-s)$ (with μ smaller than c_2) and $\mathbb{V}{\rm ar}A(s,t) = v(t-s)$. We also define the centered process $\bar{A}(\cdot)$ by $\bar{A}(t) := A(0,t) - \mu t$. We make the following assumption for the variance function. This assumption is also needed in the next chapter.

Assumption 6.1.1 *We assume that* (A1) $v(\cdot)$ *is continuous, differentiable on* $(0,\infty)$*;* (A2) $\sqrt{v(\cdot)}$ *is strictly increasing and strictly concave;* (A3) *for some* $\alpha < 2$ *it holds that* $v(t)t^{-\alpha} \to 0$ *as* $t \to \infty$.

Note that we need (A1) and (A3) to prove Theorem 6.2.1, while to derive Theorem 6.2.3 also (A2) is required. In Section 6.3 we present several examples illustrating that a broad class of processes satisfies all three assumptions.

We will frequently use the bivariate random variable $(A(-t,0), A(-s,0))$, with $s \in (0,t]$. It obviously obeys a two-dimensional Normal distribution with mean vector $(\mu t, \mu s)$ and covariance matrix $\Sigma(s,t)$. Defining the covariance function $\Gamma(s,t) := \mathbb{C}{\rm ov}[A(0,s), A(0,t)]$, we have

$$\Sigma(s,t) := \begin{pmatrix} v(t) & \Gamma(s,t) \\ \Gamma(s,t) & v(s) \end{pmatrix} \quad \text{and } \Gamma(s,t) = \frac{v(t) - v(t-s) + v(s)}{2}.$$

6.1.3 Many-sources scaling

In this section we show that the probability of interest can be written in terms of the *empirical mean process* $n^{-1}\sum_{i=1}^{n} A_i(\cdot,\cdot)$. The following lemma exploits the fact that we know both a representation of the first queue $Q_{1,n}$ and a representation of the *total* queue $Q_{1,n} + Q_{2,n}$.

Lemma 6.1.1 $\mathbb{P}(Q_{2,n} \geq nb)$ *equals*

$$\mathbb{P}\left(\bigcup_{b'\geq 0}\left\{\exists t>0:\forall s\in(0,t]:\sum_{i=1}^{n}\frac{A_i(-t,0)}{n}-c_2t\geq b+b',\right.\right.$$
$$\left.\left.\sum_{i=1}^{n}\frac{A_i(-s,0)}{n}-c_1s\leq b'\right\}\right).$$

Proof The event $\{Q_{2,n} \geq nb\}$ is trivially equivalent to

$$\bigcup_{b'\geq 0}\{Q_{1,n}+Q_{2,n}\geq n(b+b'), Q_{1,n}\leq nb'\}.$$

The stationary workload in the first queue $Q_{1,n}$ is represented by $Q_{1,n}(0)$, and the stationary workload in the total queue $Q_{1,n}+Q_{2,n}$ by $Q_{1,n}(0)+Q_{2,n}(0)$, with

$$Q_{1,n}(0)=\sup_{s>0}\left\{\sum_{i=1}^{n}A_i(-s,0)-nc_1s\right\},$$

and

$$Q_{1,n}(0)+Q_{2,n}(0)=\sup_{t>0}\left\{\sum_{i=1}^{n}A_i(-t,0)-nc_2t\right\}.$$

Hence, the event of interest reduces to

$$\bigcup_{b'\geq 0}\left\{\sup_{t>0}\left\{\sum_{i=1}^{n}\frac{A_i(-t,0)}{n}-c_2t\right\}\geq b+b',\ \sup_{s>0}\left\{\sum_{i=1}^{n}\frac{A_i(-s,0)}{n}-c_1s\right\}\leq b'\right\}.$$

Recall that minus the optimizing t (s) has the interpretation of the start of the last busy period of the total queue (the first queue) in which time 0 is contained. A positive first queue induces a positive total queue, implying that we can restrict ourself to $s \in (0,t]$. □

The above lemma implies that for analyzing $\mathbb{P}(Q_{2,n} \geq nb)$, we only have to focus on the behavior of the empirical mean process $n^{-1}\sum_{i=1}^{n}A_i(\cdot,\cdot)$. We denote by $f(\cdot)$ a path of this process, i.e., $f(r)$ is a *realization* of $\sum_{i=1}^{n}A_i(0,r)/n$ for $r \geq 0$, and similarly $f(r)$ is a realization of $\sum_{i=1}^{n}A_i(r,0)/n$ for $r < 0$. For $s,t \in \mathbb{R}, s \leq t$, we define $A[f](s,t)$ to be the value of $\sum_{i=1}^{n}A_i(s,t)/n$ for the (given) path $f(\cdot)$, i.e., $A[f](s,t) := f(t)-f(s)$. Now we can introduce the *scaled versions* of the buffer contents. Define

$$Q_1[f](0) := \sup_{s>0}\{A[f](-s,0)-c_1s\}, \tag{6.2}$$

which can be interpreted as the buffer content of a queue with service rate c_1, that is fed by the (deterministic) arrival process $A[f](\cdot,\cdot)$. The same can be done for the total queue length $Q_1[f](0)+Q_2[f](0)$. The scaling formalism can be used to analyze buffer overflow probabilities by applying sample-path large deviations, as discussed in the next subsection.

6.1.4 Sample-path large deviations

The analysis in the next sections relies on a sample-path large-deviations principle (LDP) for (centered) Gaussian processes. This subsection is devoted to a brief description of the main theorem in this field, the generalized version of *Schilder's theorem* [9]. However, we start by recalling the multivariate version of the well-known *Cramér's theorem*, see [50, Thm. 2.2.30].

Theorem 6.1.1 [Multivariate Cramér] *Let $X_i \in \mathbb{R}^d$ be i.i.d. d-dimensional random vectors, distributed as a random vector X. Then $n^{-1}\sum_{i=1}^n X_i$ satisfies the following LDP:*
(a) for any closed set $F \subset \mathbb{R}^d$,

$$\limsup_{n\to\infty} \frac{1}{n} \log \mathbb{P}\left(\frac{1}{n}\sum_{i=1}^n X_i \in F\right) \leq -\inf_{x\in F} \Lambda(x);$$

(b) for any open set $G \subset \mathbb{R}^d$,

$$\liminf_{n\to\infty} \frac{1}{n} \log \mathbb{P}\left(\frac{1}{n}\sum_{i=1}^n X_i \in G\right) \geq -\inf_{x\in G} \Lambda(x),$$

where the large-deviations rate function $\Lambda(\cdot)$ is given by

$$\Lambda(x) := \sup_{\theta\in\mathbb{R}^d} \left(\langle\theta, x\rangle - \log \mathbb{E}e^{\langle\theta,X\rangle}\right), \tag{6.3}$$

where $\langle\cdot,\cdot\rangle$ denotes the usual inner product: $\langle a, b\rangle := a^{\mathrm{T}}b = \sum_{i=1}^d a_i b_i$.

Remark 6.1.1 *Consider the specific case that $X \in \mathbb{R}^d$ has a multivariate Normal distribution with mean vector $\mu = (\mu_1, \ldots, \mu_d)^{\mathrm{T}} \in \mathbb{R}^d$ and $(d \times d)$ non-singular covariance matrix Σ. Using*

$$\log \mathbb{E}e^{\langle\theta,X\rangle} = \langle\theta,\mu\rangle + \frac{1}{2}\theta^{\mathrm{T}}\Sigma\theta,$$

it is not hard to derive that, with

$$(x-\mu)^{\mathrm{T}} \equiv (x_1-\mu_1, \ldots, x_d-\mu_d),$$

$$\theta^\star = \Sigma^{-1}(x-\mu) \quad \text{and} \quad \Lambda(x) = \frac{1}{2}(x-\mu)^{\mathrm{T}}\Sigma^{-1}(x-\mu), \tag{6.4}$$

where $\theta^\star$ optimizes (6.3). It is well-known that $\Lambda(\cdot)$ is convex (see for instance [50, Lemma 2.2.31]).

We now sketch the framework of Schilder's sample-path LDP, as established in [9], see also [53]. We restrict ourself to the aspects that are relevant in the present study; for more details we refer to [5, 4, 95]. Consider the n i.i.d. centered Gaussian processes $\bar{A}_i(\cdot) := \{\bar{A}_i(t), t \in \mathbb{R}\}$ with stationary increments, and covariance $\mathbb{C}\text{ov}[\bar{A}_i(s), \bar{A}_i(t)]$, which obviously equals $\Gamma(s,t)$ as defined in Section 6.1.2. Define the path space Ω as

$$\Omega := \left\{\omega : \mathbb{R} \to \mathbb{R}, \text{ continuous}, \ \omega(0) = 0, \ \lim_{t\to\infty} \frac{\omega(t)}{1+t} = \lim_{t\to-\infty} \frac{\omega(t)}{1+t} = 0\right\},$$

which is a separable Banach space by imposing the following norm, as explained in [95],

$$\|\omega\|_\Omega := \sup_{t\in\mathbb{R}} \left\{\frac{|\,\omega(t)\,|}{1+|\,t\,|}\right\}.$$

Next we introduce and define the *reproducing kernel Hilbert space* $R \subseteq \Omega$ – see [5] for a more detailed account – with the property that its elements are roughly as smooth as the covariance function $\Gamma(s,\cdot)$. We start from a 'smaller' space S, defined by

$$S := \left\{\omega : \mathbb{R} \to \mathbb{R}, \ \omega(\cdot) = \sum_{i=1}^{n} a_i\Gamma(s_i,\cdot), \ \ a_i, s_i \in \mathbb{R}, n \in \mathbb{N}\right\}.$$

The inner product on this space S is, for $\omega_a, \omega_b \in S$, defined as

$$\langle \omega_a, \omega_b\rangle_R := \left\langle \sum_{i=1}^{n} a_i\Gamma(s_i,\cdot), \sum_{j=1}^{n} b_j\Gamma(s_j,\cdot)\right\rangle_R = \sum_{i=1}^{n}\sum_{j=1}^{n} a_i b_j \Gamma(s_i, s_j); \tag{6.5}$$

notice that this implies $\langle \Gamma(s,\cdot), \Gamma(\cdot,t)\rangle_R = \Gamma(s,t)$. This inner product has the following useful property, which we refer to as the *reproducing kernel* property,

$$\omega(t) = \sum_{i=1}^{n} a_i\Gamma(s_i,t) = \left\langle \sum_{i=1}^{n} a_i\Gamma(s_i,\cdot), \Gamma(t,\cdot)\right\rangle_R = \langle \omega(\cdot), \Gamma(t,\cdot)\rangle_R. \tag{6.6}$$

From this we then define the norm $||\omega||_R := \sqrt{\langle \omega, \omega\rangle_R}$. The closure of S under this norm is defined as the space R. Now we can define the rate function of the sample-path LDP:

$$I(\omega) := \begin{cases} \frac{1}{2}||\omega||_R^2 & \text{if } \omega \in R; \\ \infty & \text{otherwise.} \end{cases} \tag{6.7}$$

Under the above assumptions, e.g. (A1) and (A3), the following sample-path LDP holds.

Theorem 6.1.2 [Generalized Schilder] $n^{-1}\sum_{i=1}^{n} \bar{A}_i(\cdot)$ *satisfies the following LDP:*
(a) for any closed set $F \subset \Omega$,

$$\limsup_{n\to\infty} \frac{1}{n} \log \mathbb{P}\left(\frac{1}{n}\sum_{i=1}^{n} \bar{A}_i(\cdot) \in F\right) \leq -\inf_{\omega\in F} I(\omega);$$

(b) for any open set $G \subset \Omega$,

$$\liminf_{n\to\infty} \frac{1}{n} \log \mathbb{P}\left(\frac{1}{n}\sum_{i=1}^{n} \bar{A}_i(\cdot) \in G\right) \geq -\inf_{\omega\in G} I(\omega).$$

The major difficulty of Schilder's theorem is its 'implicitness', as the explicit minimization of the rate function $I(\cdot)$ over the set of interest is often intractable. In [4] the reproducing kernel property is exploited to give a sample-path analysis of overflow in a single queue (with service rate nc) fed by Gaussian inputs. With Q_n denoting the stationary buffer content, it is derived that (centering the Gaussian arrival processes):

$$\lim_{n\to\infty}\frac{1}{n}\log\mathbb{P}(Q_n \geq nb) = -\inf_{t>0}\frac{(b+(c-\mu)t)^2}{2v(t)}. \tag{6.8}$$

If $t^\star$ denotes the minimizing t, the optimizing path for $r \in [-t^\star, 0)$ is

$$\begin{aligned} f^\star(r) &= \mu r - \frac{\Gamma(-r,t^\star)}{v(t^\star)}(b+(c-\mu)t^\star) \\ &= \mu r - \frac{v(t^\star)-v(t^\star+r)+v(-r)}{2v(t^\star)}(b+(c-\mu)t^\star). \end{aligned} \tag{6.9}$$

If on average an amount of input traffic equal to this path is sent into the buffer, then the buffer starts to fill at time $-t^\star$, and reaches overflow at time 0 (it is not hard to check that $f^\star(-t^\star) = -b - ct^\star$, and $f^\star(0) = 0$, giving $A[f](-t^\star, 0) = b + ct^\star$). Notice that the path of the centered processes (i.e., attaining value $\bar{f}^\star(s) := f^\star(s) - \mu s$ at time s) is indeed in R (in fact even in S). This optimizing path in the target set (i.e., the set of all paths leading to overflow), intersected with R, is usually referred to as the *most likely path to overflow*. It has the interpretation that given that the rare event of overflow occurs, with high probability it happens according to this trajectory. Recall that we rewrote our problem in terms of the empirical mean process. Given that $Q_n \geq nb$, the average amount of traffic behaves most likely as $f^\star(r)$, for $r \in \mathbb{R}$. The optimal $t^\star$ has the interpretation of the most likely duration of the busy period preceding overflow.

6.2 Analysis

In this section we analyze the logarithmic asymptotics of $\mathbb{P}(Q_{2,n} \geq nb)$. In Section 6.2.1 we first derive a lower bound on the decay rate (6.1), applying Cramér's and Schilder's results. It turns out that this lower bound has an insightful interpretation, which is given in Section 6.2.2. Section 6.2.3 presents conditions under which the lower bound is *tight* (meaning that the decay rate and lower bound match). Finally, in Section 6.2.4 we prove and explain the properties of the most likely path that we found.

6.2.1 Lower bound on the decay rate

The main result of this section is a tractable lower bound for the decay rate (6.1) of $\mathbb{P}(Q_{2,n} \geq nb)$, which is given in Corollary 6.2.2. Importantly, all steps towards this corollary yield equalities, except the supremum bound in Theorem 6.2.1.

The analysis is based on the insight that we derived in Section 6.1.3: to determine the asymptotics of the overflow probability in the second queue, we only need to consider the empirical mean of the input process. Hence we can apply the LDP for Gaussian processes, as given in Theorem 6.1.2, which yields

$$J(b) = \inf_{f\in\mathcal{F}_b} I(f), \tag{6.10}$$

where $\mathcal{F}_b$ is the set of paths that is defined as

$$\mathcal{F}_b := \{f \mid Q_2[f](0) \geq b\}.$$

In fact, the infimum over the set of paths $\mathcal{F}_b$ is strongly related to the infimum over the set of paths $\mathcal{F}_{b,0}$, with

$$\mathcal{F}_{b,0} := \{f \mid Q_2[f](0) \geq b,\ Q_1[f](0) = 0\},$$

i.e., $\mathcal{F}_{b,0}$ is the set of paths that lead to an empty first queue when the second queue reaches overflow. The intuition behind this is as follows. Suppose we have found a path such that at time 0 the second queue has a buffer content level of nb and the first queue is non-empty. Then, without any effort, the buffer content of the second queue at time $\delta > 0$ would be even higher than nb. Hence, this is not the 'cheapest' way to reach a buffer content nb. This principle is formalized in the next lemma.

Lemma 6.2.1 *For all $b > 0$,*

$$J(b) = \inf_{f\in\mathcal{F}_b} I(f) = \inf_{b'\geq b}\inf_{f\in\mathcal{F}_{b',0}} I(f).$$

Proof Observe that $J(b') \geq J(b)$ for all $b' > b$ and $J(b) \to \infty$ as $b \to \infty$. We prove the result in two steps. (i) We first show that for all $b > 0$ there exists $b_1 \geq b$ such that

$$\inf_{f\in\mathcal{F}_{b_1}} I(f) = \inf_{f\in\mathcal{F}_{b_1,0}} I(f). \tag{6.11}$$

(ii) Then we show that b_1 is the minimizing value of b' in the right-hand side of the result.

Ad (i). For all $b > 0 : \exists b_1 \geq b$ such that

$$b_1 := \arg\sup\left\{b'' : J(r) = J(b), \forall r \in [b, b'')\right\}.$$

It is easily verified that $J(b_1+\epsilon) > J(b_1)$ for all $\epsilon > 0$. Then (6.11) can be shown to hold using the following argument. Suppose it does not hold, then $f^\star := \arg\inf_{f\in\mathcal{F}_{b_1}} I(f)$ and $f^\star \notin \mathcal{F}_{b_1,0}$, meaning that $Q_1[f^\star](0) = a$ for some $a > 0$ (recall that $\mathcal{F}_{b_1,0} \subseteq \mathcal{F}_{b_1}$). Using that $c_1 > c_2$, there are positive δ, ϵ such that *without any additional effort* the path $f^\star$ yields $Q_2[f^\star](\delta) = b_1 + \epsilon$. This implies that $J(b_1+\epsilon) = J(b_1)$, which is in contradiction with the definition of b_1.

Ad (ii). Again using that for all $b' > 0$ we have $\mathcal{F}_{b',0} \subseteq \mathcal{F}_{b'}$, we find

$$\inf_{f \in \mathcal{F}_{b'}} I(f) \leq \inf_{f \in \mathcal{F}_{b',0}} I(f),$$

with equality for $b' = b_1$ because of (6.11). Hence, using this for the first equality, then using (6.11) for the second equality, and the definition of b_1 for the last equality, we obtain

$$\inf_{b' \geq b} \inf_{f \in \mathcal{F}_{b',0}} I(f) = \inf_{f \in \mathcal{F}_{b_1,0}} I(f) = \inf_{f \in \mathcal{F}_{b_1}} I(f) = J(b_1) = J(b).$$

This immediately gives the stated. □

Applying the same arguments as in Lemma 6.1.1, we get the following property.

Corollary 6.2.1 *Defining the set of paths $\mathcal{A}_b$ as*

$$\mathcal{A}_b := \{f \mid \exists t > 0 : \forall s \in (0, t] : A[f](-t, 0) \geq b + c_2 t,\ A[f](-s, 0) \leq c_1 s\},$$

we have

$$\inf_{f \in \mathcal{F}_{b,0}} I(f) = \inf_{f \in \mathcal{A}_b} I(f).$$

Remark 6.2.1 *In order to apply Schilder's theorem, the Gaussian processes involved need to be centered. Hence, (6.10) is formally not justified. This problem can be easily solved however, by considering the centered process $\bar{A}(\cdot)$, and subtracting μ from both service rates c_1 and c_2. We will return to this issue in Lemma 6.2.5.*

Observe that

$$\mathcal{A}_b = \bigcup_{t>0} \bigcap_{s \in (0,t]} \mathcal{A}_b^{s,t} \text{ with}$$
$$\mathcal{A}_b^{s,t} := \{f \mid A[f](-t, 0) \geq b + c_2 t,\ A[f](-s, 0) \leq c_1 s\}.$$

The probability of the union of events is reflected in the infimum of the decay rate; a similar rule, however, does not apply to the intersection of events. This gives the final exact expression for the decay rate, i.e.,

$$J(b) = \inf_{b' \geq b} \inf_{t>0} \inf_{f \in \bigcap_{s \in (0,t]} \mathcal{A}_{b'}^{s,t}} I(f). \tag{6.12}$$

In the next theorem we (conservatively) approximate the intersection over s by the supremum over s, i.e., the intersection is contained in the *least* likely of the individual events, which yields the following lower bound on the decay rate.

Theorem 6.2.1 *The following lower bound applies:*

$$J(b) \geq \inf_{b' \geq b} \inf_{t>0} \sup_{s \in (0,t]} \inf_{f \in \mathcal{A}_{b'}^{s,t}} I(f).$$

Proof Observe that for all $s \in (0,t]$ we have

$$\bigcap_{r\in(0,t]} \mathcal{A}_b^{r,t} \subseteq \mathcal{A}_b^{s,t}.$$

Hence, for all $s \in (0,t]$,

$$\inf_{f\in \bigcap_{r\in(0,t]} \mathcal{A}_b^{r,t}} I(f) \geq \inf_{f\in\mathcal{A}_b^{s,t}} I(f).$$

Now take the supremum over s in the right-hand side. Together with (6.12) this completes the proof. □

In the next lemma we show how the infimum over the set $\mathcal{A}_b^{s,t}$ can be computed. Recalling (6.4), the bivariate large-deviations rate function $\Lambda(\cdot,\cdot)$ is for $y,z \in \mathbb{R}$ and $s \in (0,t)$ defined as

$$\Lambda(y,z) := \frac{1}{2}\,(y-\mu t, z-\mu s)\,\Sigma(s,t)^{-1}\begin{pmatrix} y-\mu t \\ z-\mu s \end{pmatrix}.$$

We also introduce the notation

$$k(s,t) := \mu s + \frac{\Gamma(s,t)}{v(t)}(b+(c_2-\mu)t).$$

Lemma 6.2.2 *For* $s \in (0,t)$,

$$\inf_{f\in\mathcal{A}_b^{s,t}} I(f) = \Upsilon_b(s,t) := \begin{cases} \Lambda(b+c_2t, c_1s), & \text{if } k(s,t) > c_1 s; \\ (b+(c_2-\mu)t)^2/(2v(t)), & \text{if } k(s,t) \leq c_1 s. \end{cases}$$

Proof Observe that paths in $\mathcal{A}_b^{s,t}$ correspond to a bivariate Normal random variable. Hence we can use Theorem 6.1.1, which yields

$$\inf_{f\in\mathcal{A}_b^{s,t}} I(f) = \inf_{y\geq b+c_2t,\; z\leq c_1 s} \Lambda(y,z).$$

Using that $\Lambda(\cdot,\cdot)$ is convex, the infimum over y and z can be found using the Lagrangian

$$\mathcal{L}(y,z,\xi_1,\xi_2) := \Lambda(y,z) - \xi_1(y-b-c_2t) + \xi_2(z-c_1s),$$

with non-negative ξ_i. Two cases may occur, as illustrated in Figure 6.1. The crucial difference between the graphs is that in the left figure, the contour that touches the line $y = b+c_2t$ has a z value lower than c_1s, whereas in the right figure the opposite is the case. This is formalized as follows. Let z_0 solve

$$\left.\frac{\partial\Lambda(y,z)}{\partial z}\right|_{y=b+c_2t} = 0, \quad \text{i.e., } z_0 = \mu s + \frac{\Gamma(s,t)}{v(t)}(b+(c_2-\mu)t) = k(s,t). \quad (6.13)$$

The left panel shows that in the regime $z_0 \le c_1 s$, the optimum is attained at some point (y, z) in $\{b + c_2 t\} \times [0, c_1 s)$, whereas, according to the right panel, in the regime $z_0 > c_1 s$ an optimum is attained in $(y, z) = (b + c_2 t, c_1 s)$. In the regime $z_0 \le c_1 s$ we have

$$\Lambda(b + c_2 t, z_0) = (b + (c_2 - \mu)t)^2/(2v(t)),$$

which is indeed independent of s. □

Interestingly, the Lagrange multipliers ξ_i as used above are related to $\theta^\star$ of Theorem 6.1.1, with optimal values (6.4). Note that we use the notation $\theta_i^\star(s,t)$, $i = 1, 2$, to stress the dependence on s and t. This fact is used in the following lemma.

Lemma 6.2.3 *For $s \in (0, t)$, $\Upsilon_b(s,t)$ is increasing in b.*

Proof Observe that if $k(s,t) > c_1 s$, both conditions are binding in the proof of Lemma 6.2.2, implying that both ξ_i are strictly positive. Applying (6.4), and doing the Lagrangian calculations of the proof of Lemma 6.2.2, trivial calculus yields that

$$\begin{pmatrix} \xi_1 \\ -\xi_2 \end{pmatrix} = \begin{pmatrix} \theta_1^\star(s,t) \\ \theta_2^\star(s,t) \end{pmatrix} = \Sigma(s,t)^{-1} \begin{pmatrix} b + (c_2 - \mu)t \\ (c_1 - \mu)s \end{pmatrix}. \tag{6.14}$$

Direct calculations give

$$\frac{\partial \Upsilon_b(s,t)}{\partial b} = \frac{\partial \Lambda(b + c_2 t, c_1 s)}{\partial b} = \theta_1^\star(s,t),$$

which is positive due to the first equality in (6.14). If on the other hand $k(s,t) \le c_1 s$, then evidently $\theta_1^\star(s,t) = \xi_1 > 0$ and $\xi_2 = 0$. It turns out that

$$\frac{\partial \Upsilon_b(s,t)}{\partial b} = \theta_1^\star(s,t) = \frac{b + (c_2 - \mu)t}{v(t)} > 0.$$

This implies the stated. □

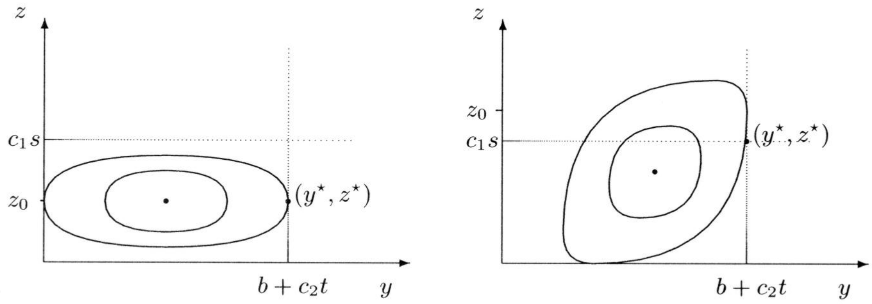

Figure 6.1: Contour lines of the (two-dimensional) rate function.

It is easily seen that $\Upsilon_b(s,t)$ is continuous at $s \downarrow 0$ and $s \uparrow t$; in particular

$$\Upsilon_b(0,t) = \frac{(b + (c_2 - \mu)t)^2}{2v(t)},$$

and

$$\Upsilon_b(t,t) = \frac{((c_1 - \mu)t)^2}{2v(t)} \quad \text{if } t = \frac{b}{c_1 - c_2} \quad \text{(and } \infty \text{ otherwise).}$$

From the continuity, we find that

$$\inf_{t>0} \sup_{s\in(0,t]} \Upsilon_b(s,t) = \inf_{t>0} \sup_{s\in(0,t)} \Upsilon_b(s,t).$$

The right-hand side of the previous display increases in b, as $\Upsilon_b(s,t)$ does so for any $t > 0$ and $s \in (0,t)$. Combining this with Theorem 6.2.1 and Lemma 6.2.2, we immediately obtain the following.

Corollary 6.2.2 *The following lower bound applies:*

$$J(b) \geq \inf_{t>0} \sup_{s\in(0,t]} \Upsilon_b(s,t).$$

6.2.2 Interpretation of the lower bound

The results of the previous section have a useful interpretation, leading to two regimes for values of c_1. In Corollary 6.2.4 we show that the lower bound in Corollary 6.2.2 can be simplified considerably for c_1 in one of these regimes.

We start with examining the single-node result, as presented in (6.8). There, t has to be found such that

$$L_c(t) := \frac{(b + (c - \mu)t)^2}{2v(t)} = \frac{(b + ct - \mathbb{E}A(-t,0))^2}{2\mathbb{V}\text{ar}A(-t,0)}$$

is minimized. Let us denote this optimizing t by t_c^F. $L_c(t)$ can be interpreted as the cost of generating $b + ct$ in an interval of length t, and t_c^F as the time epoch yielding the 'lowest cost'.

Now we turn to the result of Lemma 6.2.2. Computing the optimum of $\Lambda(y,z)$ over all $y \geq b + c_2 t$ and all $z \leq c_1 s$, the optimum over y is always attained at $y = b + c_2 t$. However, we saw that there are two possible regimes for the optimal z. Now, according to standard identities for conditional Normal distributions, recognize (6.13):

$$\begin{aligned} \mathbb{E}[A(-s,0) \mid A(-t,0) = b + c_2 t] &= \mu s + \frac{\Gamma(s,t)}{v(t)}(b + (c_2 - \mu)t) \\ &= z_0 \\ &= k(s,t). \end{aligned} \tag{6.15}$$

Popularly speaking, this implies that in the regime $k(s,t) \leq c_1 s$ the most likely realization of $\sum_{i=1}^n A_i(-t,0) \geq nb + nc_2 t$ yields $\sum_{i=1}^n A_i(-s,0) \leq nc_1 s$

with high probability (n large). In the other regime, $k(s,t) > c_1 s$, the most likely realization of $\sum_{i=1}^n A_i(-t,0) \geq nb + nc_2 t$ does not automatically yield $\sum_{i=1}^n A(-s,0) \leq nc_1 s$ with high probability; significant extra effort is needed here.

The next decomposition result follows immediately from Lemma 6.2.2 and the above.

Corollary 6.2.3 *For $s \in (0,t)$, we have $\Upsilon_b(s,t) = L_{c_2}(t) + L(s \mid t)$, with*

$$L(s \mid t) := \frac{\max^2\{\mathbb{E}[A(-s,0) \mid A(-t,0) = b + c_2 t] - c_1 s, 0\}}{2\mathbb{V}\mathrm{ar}[A(-s,0) \mid A(-t,0) = b + c_2 t]}. \tag{6.16}$$

Similar to the interpretation of the single-node result, we can interpret $\Upsilon_b(s,t)$ as the cost of generating the required amount of traffic. Denoting by $s^\star$ and $t^\star$ the optimizing arguments in Corollary 6.2.2, the intuition is as follows.

- $L_{c_2}(t)$ is just the cost of generating $b + c_2 t$ in the interval of length t. By taking the *infimum* over t (to get $t^\star$) we find the *most likely* epoch to meet the constraint.
- $L(s \mid t)$ is the cost of generating no more than $c_1 s$ in the interval $(-s, 0]$, *conditional* on the event $A(-t,0) = b + c_2 t$. We can interpret the interval $(-s^\star, 0]$ as the time period during which *most* effort has to be done to fulfill this requirement. This is of course reflected by the fact that in Corollary 6.2.2 we have to take the *supremum* over all s in $(0,t)$. Evidently, if $k(s,t) \leq c_1 s$ for all $s \in (0,t)$, this term is 0.

It is intuitively clear that for large values of c_1, $k(s,t)$ will be smaller than $c_1 s$ for all $s \in (0,t)$, since it does not depend on c_1. In this case the second term in Corollary 6.2.3 vanishes. If this holds for the t that maximizes the first term, i.e., $t^F_{c_2}$, then

$$\inf_{t>0} \sup_{s \in (0,t)} \Upsilon_b(s,t) = L_{c_2}(t^F_{c_2}). \tag{6.17}$$

It implies that, for these large values of c_1, the lower bound on the decay rate of $\mathbb{P}(Q_{2,n} \geq nb)$ in Corollary 6.2.2 coincides with the result of a single queue with service rate c_2. The intuition behind this is as follows. If c_1 is large compared to c_2, then all traffic entering the first queue is served immediately and goes directly into the second queue. Note that as a consequence of this large c_1, traffic is not 'reshaped' by the first queue, meaning that the second queue can be viewed as an ordinary single queue with service rate c_2.

If c_1 becomes smaller and approaches c_2, then c_1 will reach a value for which the first queue *does* play a role in delaying and reshaping the traffic before it can enter the second queue. The following corollary determines a *critical service rate* c_1^F and gives a simplification of the lower bound in Corollary 6.2.2 for $c_1 \geq c_1^F$. It is an immediate consequence of the fact that

$$K(t) := \sup_{s \in (0,t]} \{k(s,t) - c_1 s\}$$

is continuous and decreasing in c_1.

Corollary 6.2.4 *For all* $c_1 \geq c_1^F := \inf\left\{c_1 : K\left(t_{c_2}^F\right) \leq 0\right\}$, *Equation* (6.17) *applies.*

6.2.3 Tightness of the decay rate

In Corollary 6.2.2 a lower bound on the decay rate is given. Of course, this bound is only useful if it is relatively close to the actual decay rate $J(b)$, or, even better, coincides with $J(b)$. In the latter case we say that the lower bound is *tight*.

Recall that $s^\star$ and $t^\star$ are the optimizers in Corollary 6.2.2. Observe that we can prove tightness of the lower bound, by showing that the most probable path in $\mathcal{A}_b^{s^\star,t^\star}$ is in the set $\mathcal{A}_b$. We will distinguish between (A) $c_1 \geq c_1^F$, and (B) $c_1 < c_1^F$.

Case A: c_1 larger than the critical service rate

In this situation, we know from Corollary 6.2.4 that the lower bound in Corollary 6.2.2 reduces to the decay rate in a single queue as determined in [4]. The following result then easily follows.

Theorem 6.2.2 *If* $c_1 \geq c_1^F$ *then*

$$J(b) = \inf_{t>0} \sup_{s\in(0,t]} \Upsilon_b(s,t) = L_{c_2}(t_{c_2}^F),$$

and the most probable path is, for $r \in [-t_{c_2}^F, 0)$,

$$f^\star(r) = -\mathbb{E}[A(r,0) \mid A(-t_{c_2}^F, 0) = b + c_2 t_{c_2}^F] = -k\left(b, -r, t_{c_2}^F\right). \tag{6.18}$$

Proof As shown in Section 6.2.2, in this regime $t^\star = t_{c_2}^F$, while the choice of $s^\star$ is irrelevant. Recall that it is shown in [4] that the most probable path in the single queue with service rate c_2 is given by (6.9), with $c \equiv c_2$. With elementary formulae for conditional bivariate Normal distributions, we derive that (6.9) coincides with $f^\star(r) = -\mathbb{E}[A(r,0) \mid A(-t^\star,0) = b + c_2 t^\star]$. We now claim that $f^\star(\cdot) \in \mathcal{A}_b$, or more precisely, (i) there exists $t > 0$ such that $A[f^\star](-t,0) = -f^\star(-t^\star) \geq b + c_2 t$ and (ii) for all $r \in [-t^\star, 0)$ it holds that $A[f^\star](r,0) = -f^\star(r) \leq -c_1 r$. Assertion (i) follows trivially from $f(-t^\star) = -b - c_2 t^\star$. Claim (ii) follows because for $c_1 \geq c_1^F$ it holds that $K(t^\star) \leq 0$, by definition of c_1^F (see Corollary 6.2.4). This completes the proof. □

Remark 6.2.2 *When defining most probable paths, we only give explicit formulae for the interval* $[-t^\star, 0)$. *The paths can be trivially extended to all* $r \in \mathbb{R}$ *by taking the conditional expectation over the relevant interval.*

We want to stress that the above theorem holds for all Gaussian processes, regardless of the specific shape of the variance function. Consequently, the result is also valid for long-range dependent processes, such as fractional Brownian motion.

Case B: c_1 smaller than the critical service rate

We follow the same approach as in Case A: first we derive (in Lemma 6.2.5) the most probable path in $\mathcal{A}_b^{s,t}$, and then we verify (in Theorem 6.2.3) whether the most probable path in $\mathcal{A}_b^{s^\star,t^\star}$ is in $\mathcal{A}_b$. It turns out that we have to impose certain additional conditions to make the lower bound of Corollary 6.2.2 tight. We start by proving a technical lemma.

Lemma 6.2.4 *If $c_1 < c_1^F$, then $k(s^\star, t^\star) > c_1 s^\star$.*

Proof The lemma is proven in three steps.

- In [46, Lemma 3.1] it is shown that under Assumptions (A2) and (A3), $L_{c_2}(t)$ decreases for all $t < t_{c_2}^F$, whereas it increases for all $t > t_{c_2}^F$.
- By contradiction we prove that $K(t^\star) \geq 0$ for $c_1 < c_1^F$. Suppose $K(t^\star) < 0$, which corresponds to $k(s, t^\star) < c_1 s$ for all $s \in (0, t^\star]$. Assume first that $t^\star > t_{c_2}^F$. Now consider the decomposition of Corollary 6.2.3. Due to the fact that $K(t^\star)$ is *strictly* negative, it is possible to decrease $t^\star$ such that the term $L_{c_2}(\cdot)$ decreases (as we approach $t_{c_2}^F$ from above, see Step 1), and the second remains 0. Hence the sum of both terms decreases, implying that $t^\star$ cannot be optimal. So it cannot be that both
$$K(t^\star) = \sup_{s \in (0,t^\star]} k(s, t^\star) - c_1 s < 0 \quad \text{and} \quad t^\star > t_{c_2}^F.$$
Similarly $K(t^\star) < 0$ rules out $t^\star < t_{c_2}^F$. So $K(t^\star) < 0$ implies $t^\star = t_{c_2}^F$. We however assumed that $K(t_{c_2}^F) > 0$, as we chose c_1 smaller than c_1^F. Contradiction.
- Observe that there is an $s \in (0, t^\star]$ such that $k(s, t^\star) \geq c_1 s$ and that $s^\star$ is determined by maximizing $L(s \mid t^\star)$ over $s \in (0, t^\star]$, see (6.16).

This concludes the proof. □

Lemma 6.2.5 *If $k(s,t) > c_1 s$, then, for $r \in [-t, 0)$, the most probable path in $\mathcal{A}_b^{s,t}$ is*
$$-\mathbb{E}[A(r, 0) \mid A(-t, 0) = b + c_2 t, A(-s, 0) = c_1 s]. \tag{6.19}$$

Proof The proof is along the same lines as [4, Prop. 2.2]. We first notice, with standard properties of conditional multivariate Normal random variables, that (6.19) equals
$$\mu r - \begin{pmatrix} \mathbb{C}\mathrm{ov}[A(r,0), A(-t,0)] \\ \mathbb{C}\mathrm{ov}[A(r,0), A(-s,0)] \end{pmatrix}^{\mathrm{T}} \Sigma(s,t)^{-1} \begin{pmatrix} b + (c_2 - \mu)t \\ (c_1 - \mu)s \end{pmatrix} =$$
$$\mu r - \theta_1^\star(s,t)\Gamma(-r,t) - \theta_2^\star(s,t)\Gamma(-r,s), \tag{6.20}$$
using (6.14) and the notation of Section 6.1.4. Obviously, in self-evident notation,
$$\mathcal{A}_b^{s,t} = \bigcup_{y \geq b + c_2 t, z \leq c_1 s} \mathcal{A}_{y,z},$$

with

$$\begin{aligned}\mathcal{A}_{y,z} &:= \{f \mid A[f](-t,0) = y, A[f](-s,0) = z\}\\ &= \{\bar{f} \mid \bar{A}[\bar{f}](-t,0) = y - \mu t, \bar{A}[\bar{f}](-s,0) = z - \mu s\}.\end{aligned}$$

Notice that we state the minimization in terms of the *centered* process, as is required by Schilder's theorem. Obviously $f^\star(\cdot)$ and $\bar{f}^\star(\cdot)$ are related through $f^\star(r) = \mu r + \bar{f}^\star(r)$. We first find the most probable path $\bar{f}^\star_{y,z}(\cdot)$ in $\mathcal{A}_{y,z}$, and then take the infimum over $y \geq b + c_2 t$ and $z \leq c_1 s$. Because this path has to be found in the set $\mathcal{A}_{y,z} \cap R$, the reproducing kernel property (6.6) gives

$$\bar{f}^\star(-t) = \langle \bar{f}^\star(\cdot), \Gamma(-\cdot, t)\rangle_R \quad \text{and} \quad \bar{f}^\star(-s) = \langle \bar{f}^\star(\cdot), \Gamma(-\cdot, s)\rangle_R.$$

Combining this with the definition of $\mathcal{A}_{y,z}$, and recalling $\bar{A}[\bar{f}](-t,0) = -\bar{f}(-t)$, it follows that $\bar{f}^\star_{y,z}(\cdot)$ has to satisfy

$$\langle \bar{f}^\star(\cdot), \Gamma(-\cdot, t)\rangle_R = -y + \mu t, \quad \text{and} \quad \langle \bar{f}^\star(\cdot), \Gamma(-\cdot, s)\rangle_R = -z + \mu s. \tag{6.21}$$

Because both $\Gamma(-\cdot, t)$ and $\Gamma(-\cdot, s)$ are orthogonal to the hyperplanes defined by (6.21), the solution is a linear combination of these vectors, i.e.,

$$f^\star_{y,z}(r) = \mu r + \zeta_1 \Gamma(-r, t) + \zeta_2 \Gamma(-r, s),$$

for $r \in [-t, 0)$, where ζ_1 and ζ_2 clearly depend on y and z. Equation (6.20) now suggests

$$(\zeta_1, \zeta_2)^{\mathrm{T}} := -\Sigma(s,t)^{-1}(y,z)^{\mathrm{T}}.$$

It is easily verified that under this choice $f^\star_{y,z}(\cdot)$ indeed satisfies (6.21). With the inner product defined by (6.5), we derive after tedious calculations that $\frac{1}{2}\|f^\star_{y,z}\|^2_R = \Lambda(y,z)$. Theorem 6.1.2 now implies that $f^\star_{y,z}(\cdot)$ is indeed minimizer in $\mathcal{A}_{y,z}$. Due to Lemma 6.2.2, we have that if $k(s,t) > c_1 s$, then the infimum is attained at $y = b + c_2 t$ and $z = c_1 s$. □

Before presenting our tightness condition, we introduce some new notation. For $r_1 < r_2$,

$$\bar{\mathbb{E}}A(r_1, r_2) := \mathbb{E}[A(r_1, r_2) \mid A(-t^\star, 0) = b + c_2 t^\star],$$

with $\bar{\mathbb{V}\mathrm{ar}}[\cdot]$ and $\bar{\mathbb{C}\mathrm{ov}}[\cdot,\cdot]$ defined similarly. For $r \in (-t^\star, 0)$ we define the functions

$$\bar{m}(r) := \frac{\bar{\mathbb{E}}A(r,0) + c_1 r}{\sqrt{\bar{\mathbb{V}\mathrm{ar}}A(r,0)}}, \qquad m(r) := \frac{\bar{m}(r)}{\bar{m}(-s^\star)},$$

and

$$\rho(r) := \frac{\bar{\mathbb{C}\mathrm{ov}}[A(r,0), A(-s^\star,0)]}{\sqrt{\bar{\mathbb{V}\mathrm{ar}}A(r,0)\, \bar{\mathbb{V}\mathrm{ar}}A(-s^\star,0)}}.$$

Note that both $\rho(\cdot)$ and $m(\cdot)$ attain a maximum 1 at $r = -s^\star$. For $\rho(\cdot)$ this follows from the observation that $\rho(r)$ is a correlation coefficient; for $m(\cdot)$ from Corollary 6.2.3 and Lemma 6.2.4.

Theorem 6.2.3 *If $c_1 < c_1^F$, and iff for all $r \in (-t^\star, 0)$ it holds that $m(r) \leq \rho(r)$, then*

$$J(b) = \inf_{t>0} \sup_{s\in(0,t]} \Upsilon_b(s,t) = \Lambda(b + c_2 t^\star, c_1 s^\star),$$

and the most probable path is

$$f^\star(r) = -\mathbb{E}[A(r,0) \mid A(-t^\star,0) = b + c_2 t^\star, A(-s^\star,0) = c_1 s^\star].$$

Proof Using Lemma 6.2.4 we know that for $s^\star$ and $t^\star$ we are in the regime where $k(s^\star,t^\star) > c_1 s^\star$. We (again) have to show that $f^\star(\cdot)$ is in $\mathcal{A}_b$. It is easy to check that (i) $A[f^\star](-t^\star,0) = b + c_2 t^\star$ (use (6.14)). It is left to show that (ii) under $m(\cdot) \leq \rho(\cdot)$ it holds that $A[f^\star](r,0) \leq -c_1 r$ for all $r \in (-t^\star, 0)$. Now rewrite

$$\begin{aligned} A[f^\star](r,0) &= \bar{\mathbb{E}}\left[A(r,0) \mid A(-s^\star,0) = c_1 s\right] \\ &= \bar{\mathbb{E}}A(r,0) + \frac{\bar{\mathbb{Cov}}[A(r,0), A(-s^\star,0)]}{\bar{\mathbb{Var}}A(-s^\star,0)} \left(c_1 s^\star - \bar{\mathbb{E}}A(-s^\star,0)\right). \end{aligned}$$

For $r \in (-t^\star, 0)$ it is easily seen that $A[f^\star](r,0) \leq -c_1 r$ is equivalent to $m(r) \leq \rho(r)$. □

Although the (necessary and sufficient) condition $m(r) \leq \rho(r)$ for all $r \in (0, t^\star)$, required in Theorem 6.2.3, is stated in terms of the model parameters, as well as known statistics of the arrival process, it could be a tedious task to verify this condition for a specific situation. The next lemma presents a *necessary* condition for $m(r) \leq \rho(r)$ to hold.

Lemma 6.2.6 *A necessary condition for $m(r) \leq \rho(r)$, for all $r \in (-t^\star, 0)$, is*

$$m''(-s^\star) \leq \rho''(-s^\star), \tag{6.22}$$

or equivalently,

$$\theta_1^\star(s^\star,t^\star) \cdot (v''(t^\star - s^\star) - v''(s^\star)) \; + \; \theta_2^\star(s^\star,t^\star) \cdot (v''(0) - v''(s^\star)) \; \geq \; 0. \tag{6.23}$$

Proof First we show that (6.22) holds. As noted earlier, both $m(\cdot)$ and $\rho(\cdot)$ have a maximum 1 at $-s^\star$. This means that (6.22) is necessary to enforce $m(r) \leq \rho(r)$ for r in a neighborhood of $-s^\star$.

Next we show that (6.22) is equivalent to (6.23). First multiply both $m(\cdot)$ and $\rho(\cdot)$ by $h(\cdot)$, where

$$h(r) := \sqrt{\frac{\bar{\mathbb{Var}}A(r,0)}{\bar{\mathbb{Var}}A(-s^\star,0)}} \left(\bar{\mathbb{E}}A(-s^\star,0) - c_1 s^\star\right).$$

Since $h(r) : (-t^\star, 0) \to \mathbb{R}_+$, this yields the requirement $\pi(r) \leq n(r)$ for all $r \in (-t^\star, 0)$, with

$$\begin{aligned} \pi(r) &:= \bar{\mathbb{E}}A(r,0) + c_1 r, \\ n(r) &:= \frac{\bar{\mathbb{Cov}}[A(r,0), A(-s^\star,0)]}{\bar{\mathbb{Var}}A(-s^\star,0)} \left(\bar{\mathbb{E}}A(-s^\star,0) - c_1 s^\star\right). \end{aligned}$$

Recall that $m(\cdot)$ and $\rho(\cdot)$ have the same function value and derivative at $-s^\star$. It is easy to derive that this implies that $(m\cdot h)(-s^\star) = (\rho\cdot h)(-s^\star)$ (where we denote by $(m\cdot h)(\cdot)$ the function obtained after multiplying $m(\cdot)$ with $h(\cdot)$) and $(m\cdot h)'(-s^\star) = (\rho\cdot h)'(-s^\star)$. Therefore, the necessary condition becomes $\pi''(-s^\star) \leq n''(-s^\star)$.

Using standard formulae for conditional means of multivariate Normal random variables,

$$\bar{\mathbb{E}}A(r,0) = -\mu r + \frac{\Gamma(-r,t^\star)}{v(t^\star)}\,(b + (c_2-\mu)t^\star),$$

leading to

$$\frac{\mathrm{d}}{\mathrm{d}r^2}\left(\bar{\mathbb{E}}A(r,0) + c_1 r\right)\Big|_{r=-s^\star} = \frac{b+(c_2-\mu)t^\star}{2v(t^\star)}\,(v''(s^\star) - v''(t^\star - s^\star)).$$

Assuming $r \leq -s^\star$,

$$\bar{\mathbb{C}\mathrm{ov}}[A(r,0), A(-s^\star,0)] = \frac{v(-r)+v(s^\star)-v(-r-s^\star)}{2} - \frac{\Gamma(-r,t^\star)\Gamma(s^\star,t^\star)}{v(t^\star)},$$

such that

$$\frac{\mathrm{d}^2}{\mathrm{d}r^2}\bar{\mathbb{C}\mathrm{ov}}[A(r,0), A(s^\star,0)]\Big|_{r=-s^\star} = \frac{v''(s^\star)-v''(0)}{2} - \frac{v''(s^\star)-v''(t^\star-s^\star)}{2}\,\frac{v(s^\star,t^\star)}{v(t^\star)}.$$

Note that the same result holds when the derivative is calculated for $r > -s^\star$. Now it is a straightforward but tedious computation to prove that this implies that $\pi''(-s^\star) \leq n''(-s^\star)$ is equivalent to (6.23). □

6.2.4 Properties of the most probable path

So far, we have analyzed trajectories $f(\cdot)$ and $A[f](\cdot,\cdot)$, i.e., the path of the *cumulative* amount of traffic generated. In this section we turn our attention to the first derivative of $f(\cdot)$, which can be interpreted as the path of the *input rate*. We define this input rate path by $g(\cdot)$, i.e.,

$$g(r) := -\frac{\mathrm{d}}{\mathrm{d}r}A[f](r,0) = f'(r).$$

As before, we have to consider two regimes: (A) $c_1 \geq c_1^F$, and (B) $c_1 < c_1^F$. In both regimes we compute $g^\star(\cdot) := -(f^\star)'(\cdot)$. In case (A), with $t^\star = t^F_{c_2}$, and $r \in [-t^\star, 0)$,

$$g^\star(r) = \mu + \frac{b+(c_2-\mu)t^\star}{2v(t^\star)}\,(v'(r+t^\star) + v'(-r)),$$

whereas in case (B) it turns out that, with $r \in [-t^\star, -s^\star)$,

$$g^\star(r) = \mu + \frac{v'(r+t^\star)+v'(-r)}{2}\,\theta_1^\star(s^\star,t^\star) + \frac{-v'(-r-s^\star)+v'(-r)}{2}\,\theta_2^\star(s^\star,t^\star),$$

and with $r \in [-s^\star, 0)$,

$$g^\star(r) = \mu + \frac{v'(r+t^\star) + v'(-r)}{2}\,\theta_1^\star(s^\star,t^\star) + \frac{v'(r+s^\star)+v'(-r)}{2}\,\theta_2^\star(s^\star,t^\star).$$

Assuming $v'(0) = 0$, the path $g^\star(\cdot)$ has some nice properties. This assumption holds for many Gaussian processes. Note however that standard Brownian motion (Bm) does not fulfill this requirement, as $v(t) = t$. The special structure of Bm allows an explicit analysis, see Section 6.3.1. Fractional Brownian motion (fBm), with $v(t) = t^{2H}$, has $v'(0) = 0$ only for $H \in (\frac{1}{2}, 1]$, see Section 6.3.2.

Proposition 6.2.1 *If $c_1 \geq c_1^F$ and $v'(0) = 0$ then $g^\star(0) = g^\star(-t^\star) = c_2$.*

Proof Notice that, due to (6.8), $t^\star$ satisfies

$$2(c_2 - \mu)\,\frac{v(t^\star)}{v'(t^\star)} = b + (c_2-\mu)t^\star.$$

Using $v'(0) = 0$ yields the stated. (As an aside, we mention that $g^\star(\cdot)$ is clearly symmetric around $-t^\star/2$.) □

Just as we exploited properties of $t^\star$ in the proof of Proposition 6.2.1, we need conditions for $s^\star$ and $t^\star$ for the regime $c_1 < c_1^F$. These are derived in the next lemma.

Lemma 6.2.7 *If $c_1 < c_1^F$, then $s^\star$ and $t^\star$ satisfy the following equations:*

$$\begin{aligned} 2(c_2-\mu) &= \theta_1^\star(s^\star,t^\star)\cdot v'(t^\star) + \theta_2^\star(s^\star,t^\star)\cdot(v'(t^\star) - v'(t^\star - s^\star));\\ 2(c_1-\mu) &= \theta_2^\star(s^\star,t^\star)\cdot v'(s^\star) + \theta_1(s^\star,t^\star)\cdot(v'(s^\star) + v'(t^\star - s^\star)). \end{aligned}$$

Proof By Lemma 6.2.4, $k(s^\star,t^\star) > c_1 s^\star$. Observe that $\Upsilon_b(s,t) = \Lambda(b + c_2 t, c_1 s)$ can be rewritten as

$$\theta^{\mathrm T} x(s,t) - \frac{1}{2}\theta^{\mathrm T}\Sigma(s,t)\theta, \quad \text{where } x(s,t) := \begin{pmatrix} b + (c_2-\mu)t \\ (c_1-\mu)s \end{pmatrix}; \tag{6.24}$$

here we abbreviate $\theta \equiv [\theta_1^\star(s,t)\ \theta_2^\star(s,t)]^{\mathrm T}$. We write ∂_t and ∂_s for the partial derivatives with respect to t and s, respectively. The optimal $s^\star$ and $t^\star$ necessarily satisfy the first-order conditions, obtained by differentiating (6.24) with respect to t and s, and equating them to 0. Direct calculations yield

$$\begin{aligned} \begin{pmatrix} \theta_1(c_2-\mu) \\ \theta_2(c_1-\mu) \end{pmatrix} &= \begin{pmatrix} \partial_t\theta_1 & \partial_t\theta_2 \\ \partial_s\theta_1 & \partial_s\theta_2 \end{pmatrix} (\Sigma(s,t)\theta - x(s,t)) \\ &\quad + \begin{pmatrix} \frac{1}{2}\theta_1^2 v'(t) + \partial_t\Gamma(s,t)\theta_1\theta_2 \\ \frac{1}{2}\theta_2^2 v'(s) + \partial_s\Gamma(s,t)\theta_1\theta_2 \end{pmatrix}. \end{aligned}$$

The second equality in (6.14) gives $x(s,t) = \Sigma(s,t)\theta$. Now the stated follows directly. □

Proposition 6.2.2 *If* $c_1 < c_1^F$ *and* $v'(0) = 0$ *then* (i) $g^\star(-t^\star) = c_2$, *and* (ii) $g^\star(-s^\star) = c_1$. *Also, the necessary condition (6.23) is equivalent to* $(g^\star)'(-s^\star) \geq 0$.

Proof Claims (i) and (ii) follow directly from $v'(0) = 0$ and Lemma 6.2.7. The last statement follows directly after some calculations. □

Proposition 6.2.2 can be interpreted as follows. To cause buffer overflow in the second queue at time 0, this buffer starts to fill at time $-t^\star$. During this trajectory, the first queue starts a busy period at time $-s^\star$ and is empty again at time 0, if the conditions of Theorem 6.2.3 apply.

Remark 6.2.3 *The approach we have followed in this section to analyze the two-node tandem network, can be easily applied to an* m*-node tandem network, with strictly decreasing service rates, i.e.,* $c_1 > \ldots > c_m$ *– nodes* i *for which* $c_i \leq c_{i+1}$ *can be ignored, cf. [71]. Note that* $\sum_{i=1}^{k} Q_{i,n}$ *is equivalent to the queue length in a single node fed by the same sources and emptied at rate* c_k*. This means that we have the characteristics of both* $\sum_{i=1}^{m-1} Q_{i,n}$ *and* $\sum_{i=1}^{m} Q_{i,n}$*, which enables the analysis of* $Q_{m,n}$*, just as in the two-node tandem case.*

6.3 Examples

One of the reasons for considering Gaussian input processes, is that they cover a broad range of correlation structures. Choosing the variance function appropriately, we can make the input process exhibit for instance long-range dependent behavior. In this section we computate the decay rate explicitly for various variance functions. We also discuss in detail the condition in Theorem 6.2.3.

6.3.1 Standard Brownian motion

The variance function for Brownian motion (Bm) is given by $v(t) = t$. Using (6.8), it is easily found that $t_{c_2}^F = b/(c_2 - \mu)$. According to Corollary 6.2.4, c_1^F is the largest value of c_1 such that for all $s \in (0, t_{c_2}^F]$,

$$\mu s + \frac{s}{t_{c_2}^F}\left(b + (c_2 - \mu)t_{c_2}^F\right) - c_1 s \leq 0,$$

i.e., $c_1^F = 2c_2 - \mu$. Hence, using Theorem 6.2.2, we have for $c_1 \geq 2c_2 - \mu$, that $J(b) = 2b(c_2 - \mu)$, with a constant input rate $g^\star(r) = 2c_2 - \mu$ for $r \in (-t_{c_2}^F, 0)$ and $g^\star(r) = \mu$ for $r \in \mathbb{R} \setminus \{[-t_{c_2}^F, 0]\}$.

Now we turn to the case where $c_1 < 2c_2 - \mu$. The optimizing $s^\star$ and $t^\star$ are determined by solving the first-order equations for s and t of $\Lambda(b + c_2 t, c_1 s)$ (see Lemma 6.2.2). We immediately obtain that $t^\star = b/(c_1 - c_2)$ and $s^\star = 0$. Obviously, for this regime the service rate of the first queue *does* play a role. The most probable path reads $g^\star(r) = c_1$, for $r \in (-t^\star, 0)$ and $g^\star(r) = \mu$ for $r \in \mathbb{R} \setminus \{[-t^\star, 0]\}$. Because $A[f^\star](r, 0) = -c_1 r$, it is easily verified that the most probable path $f^\star(\cdot)$ is an element of the set $\mathcal{A}_b$, making the decay rate as found in Theorem 6.2.3 tight. In other words,

$$J(b) = \Lambda(b + c_2 t^\star, c_1 s^\star) = \frac{b(c_1-\mu)^2}{2(c_1-c_2)}.$$

Observe that, interestingly, Bm apparently changes its rate instantaneously, as reflected by the most likely path. This is a consequence of the independence of the increments.

6.3.2 Fractional Brownian motion

The variance function for fBm is given by $v(t) = t^{2H}$, with H the Hurst parameter as defined in Section 2.1. An interesting regime is that with $H > \frac{1}{2}$, for which the process exhibits long-range dependence. For general H we obtain from (6.8):

$$t^F_{c_2} = \frac{b}{c_2-\mu}\frac{H}{1-H}.$$

By Theorem 6.2.2,

$$J(b) = \frac{1}{2}\left(\frac{b}{1-H}\right)^{2-2H}\left(\frac{c_2-\mu}{H}\right)^{2H}$$

for all $c_1 \geq c_1^F$. Unfortunately, for general H there does not exist a closed-form expression for c_1^F. Let us consider the case $c_1 < c_1^F$. Lemma 6.2.6 states that (6.23) is a necessary condition for tightness to hold. Observe that $v''(t) = (2H-1)2Ht^{2H-2}$ and hence $v''(0) = \infty$. Combining this with $\theta_2^\star(s^\star, t^\star) \leq 0$ means that in this case (6.23) is not satisfied. Therefore the lower bound on $J(b)$ is *not* tight.

6.3.3 M/G/∞ input

A versatile traffic model is the so-called M/G/∞ input process. In this model sessions arrive according to a Poisson process with rate λ, and stay in the system for some random duration D. During this period they generate traffic at a unit rate. By choosing specific session-length distributions D, both short-range and long-range dependent inputs can be modeled. For more results on queues with M/G/∞ input traffic processes, see e.g. [54, 110]. Below we approximate the M/G/∞ inputs by their 'Gaussian counterpart', i.e., Gaussian sources with the same mean and variance as the M/G/∞ input. This procedure is motivated in [4].

Let the mean session length be finite, say δ, such that the mean input rate equals $\lambda\delta$. We denote by $F_D(\cdot)$ the distribution function of D and by $F_{D^r}(\cdot)$ the distribution function of the *residual* session length D^r (recall that $F_{D^r}(x) = \delta^{-1}\int_0^x(1-F_D(y))dy$). We denote the corresponding densities by $f_D(\cdot)$ and $f_{D^r}(\cdot)$.

We now show how to compute the variance $v(t)$ of $a(t)$, which denotes the amount of traffic generated by a single M/G/∞ input source in an interval of length t. We will do this by first deriving the moment generating function of $a(t)$. In fact two types of sources contribute.

- Sources that were already present at the start of the interval. The number of these sources has a Poisson distribution with mean $\lambda\delta$. Their residual duration has density $f_{D^r}(\cdot)$; with probability $(1 - F_{D^r}(t))$ they transmit traffic during the entire interval.
- Sources that arrive during the interval. Their number follows a Poisson distribution with mean λt. Given an arrival, this happens at an epoch that is uniformly distributed over the interval (with density t^{-1}). Their duration has density $f_D(\cdot)$.

Straightforward computations now yield, cf. [89]

$$\log \mathbb{E}e^{\theta a(t)} = \lambda\delta(M_t(\theta) - 1) + \lambda t(N_t(\theta) - 1),$$

with

$$M_t(\theta) := \int_0^t e^{\theta x} f_{D^r}(x)\mathrm{d}x + e^{\theta t}(1 - F_{D^r}(t))$$

and

$$N_t(\theta) := \int_0^t \int_u^t \frac{1}{t} e^{\theta(x-u)} f_D(x-u)\mathrm{d}x\mathrm{d}u + \int_0^t \frac{1}{t} e^{\theta(t-u)}(1 - F_D(t-u))\mathrm{d}u.$$

Taking the second derivative of the log moment generating function with respect to θ and then substituting 0 for θ, gives the variance $v(t)$ of $a(t)$:

$$\lambda\delta\left(\int_0^t x^2 f_{D^r}(x)\mathrm{d}x + t^2(1 - F_{D^r}(t))\right)$$

$$+\lambda\left(\int_0^t \int_u^t (x-u)^2 f_D(x-u)\mathrm{d}x\mathrm{d}u + \int_0^t (t-u)^2(1 - F_D(t-u))\mathrm{d}u\right).$$

For fBm we could *a priori* rule out tightness of the lower bound due to $v''(0) = \infty$, see Lemma 6.2.6. For M/G/∞ inputs we show in the following lemma that $v''(0)$ is finite, even for heavy-tailed D. It implies that the condition $m(r) \leq \rho(r)$ for all $r \in (-t^\star, 0)$ needs to be checked to verify tightness.

Lemma 6.3.1 *For $\delta < \infty$ and finite $f_D(\cdot)$, both* (i) $v'(0) = 0$ *and* (ii) $v''(0) < \infty$.

Proof Using standard rules for differentiation of integrals,

$$v'(t) = \lambda\delta 2t\,(1 - F_{D^r}(t)) + \lambda \int_0^t 2(t-u)(1 - F_D(t-u))\mathrm{d}u,$$

and hence $v'(0) = 0$. Similarly,

$$\begin{aligned} v''(t) &= 2\lambda\delta\left(1 - F_{D^r}(t) - t f_{D^r}(t)\right) \\ &\quad + 2\lambda \int_0^t \left(1 - F_D(t-u) - (t-u) f_D(t-u)\right) \mathrm{d}u \\ &= 2\lambda \int_t^\infty (1 - F_D(s)) \mathrm{d}s. \end{aligned}$$

Hence, $v''(0) = 2\lambda\delta < \infty$. □

Now we consider some examples of session-length distributions. In all the examples we take $b = 0.5$, $\lambda = 0.125$, $\delta = 2$ and $c_2 = 1$.

Exponential. Using the above formula for $v(\cdot)$, we get

$$v(t) = 2\lambda\delta^3 \left(\frac{t}{\delta} + \exp\left(-\frac{t}{\delta}\right) - 1 \right).$$

Numerical computations then give $c_1^F = 1.195$. Now taking $c_1 = 1.1$ results in $s^\star = 4.756$, $t^\star = 5.169$ and $m(r)$, $\rho(r)$ as given in Figure 6.2. The left panel shows $m(r)$ and $\rho(r)$ for $r \in (-t^\star, 0)$ and in the right panel the figure is magnified around $-s^\star$. We see that indeed $m(\cdot) \leq \rho(\cdot)$ on the desired interval, so the decay rate is tight. The input rate path is given in Figure 6.3 and satisfies the properties as indicated in Proposition 6.2.2.

Hyperexponential. In case D has a hyperexponential distribution, it is easily verified that

$$v(t) = 2\lambda \frac{p}{\nu_1^3} \left(\nu_1 t - 1 + e^{-\nu_1 t} \right) + 2\lambda \frac{1-p}{\nu_2^3} \left(\nu_2 t - 1 + e^{-\nu_2 t} \right),$$

with $\nu_2 = (1-p)/(\delta - p/\nu_1)$. For $p = 0.25$ and $\nu_1 = 5$, we find $c_1^F = 1.173$, and $s^\star = 4.700$, $t^\star = 5.210$, when using $c_1 = 1.1$. Also for this example $m(\cdot) \leq \rho(\cdot)$ as can be verified in Figure 6.4. The input rate path is given in Figure 6.5.

Pareto. If D has a Pareto distribution, the variance function is given by

$$v(t) = \frac{2\lambda}{(3-\alpha)(2-\alpha)(1-\alpha)} \left(1 - (t+1)^{3-\alpha} + (3-\alpha)t \right),$$

with $\alpha = (1+\delta)/\delta$, excluding $\delta = \frac{1}{2}$, $\frac{1}{3}$ or $\frac{1}{4}$. Notice that we have $\alpha = 1\frac{1}{2}$, yielding $v(t) \sim t\sqrt{t}$, which corresponds to long-range dependent traffic. Numerical calculations show that $c_1^F = 1.115$, and for $c_1 = 1.1$ we obtain $s^\star = 4.373$, $t^\star = 5.432$. Again $m(\cdot)$ is majorized by $\rho(\cdot)$, as can be seen in Figure 6.6. The input rate path is given in Figure 6.7. We empirically found that there is not always tightness in the M/Par/∞ case. If b is larger, for instance $b = 1$, then (6.23) is not met.

6.4 Priority queues

In Section 6.2 we analyzed overflow in the second queue of a tandem system. This analysis was enabled by the fact that we have explicit knowledge of both the *first queue* of the series and the *total queue*, i.e., we have expressions of both in terms of the arrival processes and the service rates. An explicit expression of the second queue in terms of arrival process and service rates, requires a characterization of the output process of the first queue. However, we used the observation that the second queue is simply the difference of the total queue and the first queue, which turned out to be a good starting point for the analysis. Now observe that the same principle holds for the two-queue priority system. There we have expressions in terms of arrival processes and service rates for both the high-priority queue and the total queue. Although an explicit expression for the low-priority queue requires the output process of the high-priority queue, we can analyze the low-priority queue simply by writing it as the difference of the total queue and the high-priority queue.

6.4.1 Analysis

We consider a priority queue, with service rate nc, which is fed by traffic of two classes, each with its own queue. Traffic of class 1 does not 'see' class 2 at all, and consequently we know how the *high-priority queue* $Q_n^{\rm h}$ behaves. Also, due to the work-conserving property of the system, the *total queue* length $Q_n^{\rm h} + Q_n^{\ell}$ can be characterized. Now we are able, applying the same arguments as for the tandem queue, to analyze the decay rate of the probability of exceeding some buffer threshold in the low-priority queue.

We let the system be fed by n i.i.d. high-priority (hp) sources, and an equal number of i.i.d. low-priority (lp) sources; both classes are independent. We assume that both hp and lp sources are Gaussian, and satisfy the requirements imposed in Section 6.1. Define the means by $\mu^{\rm h}$ and μ^{ℓ}, and the variance functions by $v^{\rm h}(\cdot)$ and $v^{\ell}(\cdot)$, respectively; also $\mu := \mu^{\rm h} + \mu^{\ell}$ and $v(\cdot) := v^{\rm h}(\cdot) + v^{\ell}(\cdot)$. We denote the amount of traffic from the ith hp source in $(s,t]$, with $s < t$, by $A_i^{\rm h}(s,t)$; we define $A_i^{\ell}(s,t)$ analogously. Also $\Gamma^{\rm h}(s,t)$ and $\Gamma^{\ell}(s,t)$ are defined as before. (Notice that this setting also covers the case that the numbers of sources of both classes are *not* equal. Assume for instance that there are $n\alpha$ lp sources. Multiplying μ^{ℓ} and $v^{\ell}(\cdot)$ by α and applying the fact that the Normal distribution is infinitely divisible, we arrive at n i.i.d. sources.) The analysis of the low-priority queue is justified by the 'two-dimensional Schilder' framework. The formal description of this framework can be found in the next chapter (Section 7.1 (see also [95])), where we need it to analyze the two-queue GPS system. In this section we give the results for the priority queue, without the formal proofs, as they follow in a straightforward manner from the tandem results. In Section 7.1 a large-deviations rate function $I(\cdot)$ is introduced, with two-dimensional argument $f(\cdot) \equiv (f^{\rm h}(\cdot), f^{\ell}(\cdot))$. Let $A^{\rm h}[f]$ and $A^{\ell}[f]$ be defined similarly as before.

Analogously to Lemma 6.1.1, we obtain that $\mathbb{P}(Q_n^{\ell} \geq nb)$ equals

$$\mathbb{P}\left(\bigcup_{x\geq 0, t>0} \bigcap_{s\in(0,t]} \left\{ \sum_{i=1}^{n} \frac{A_i^{\rm h}(-t,0) + A_i^{\ell}(-t,0)}{n} - ct \geq b + x; \right.\right.$$

$$\left.\left.\sum_{i=1}^{n} \frac{A_i^{\rm h}(-s,0)}{n} - cs \leq x\right\}\right).$$

Define the exponential decay rate of this probability by $J_p(b)$. Although we expect that the most likely path is such that the hp queue is empty at the epoch the lp queue reaches overflow, we have not succeeded in proving the counterpart of Lemma 6.2.1. Define

$$\mathcal{A}_{b,x}^{s,t} := \{f \mid A^{\rm h}[f](-t,0) + A^{\ell}[f](-t,0) \geq b + x + ct, A^{\rm h}[f](-s,0) \leq cs + x\},$$

and

$$\mathcal{A}_{b,x} := \bigcup_{t\geq 0} \bigcap_{s\in(0,t]} \mathcal{A}_{b,x}^{s,t}.$$

The following results are proven similarly to those in Section 6.2.

Theorem 6.4.1 *The following lower bound applies:*

$$J_p(b) = \inf_{x\geq 0} J_p(b,x),$$

where

$$J_p(b,x) := \inf_{f\in\mathcal{A}_{b,x}} I(f) \geq \inf_{t>0} \sup_{s\in(0,t]} \inf_{f\in\mathcal{A}_{b,x}^{s,t}} I(f). \tag{6.25}$$

The infimum over all $f(\cdot) \in \mathcal{A}_{b,x}^{s,t}$ can be computed explicitly. To present the counterpart of Lemma 6.2.2, define

$$k(x,s,t) := \mu^{\rm h} s + \frac{\Gamma^{\ell}(s,t)}{v(t)} \left(b + x + (c-\mu)t\right),$$

and

$$\Lambda(y,z) := \frac{1}{2}\left(y - \mu t,\ z - \mu^{\rm h} s\right) \begin{pmatrix} v(t) & \Gamma^{\ell}(s,t) \\ \Gamma^{\ell}(s,t) & v^{\ell}(s) \end{pmatrix} \begin{pmatrix} y - \mu t \\ z - \mu^{\rm h} s \end{pmatrix}.$$

Lemma 6.4.1 *For $s \in (0,t)$,*

$$\begin{aligned} &\inf_{f\in\mathcal{A}_{b,x}^{s,t}} I(f) \\ = \quad &\Upsilon_{b,x}(s,t) := \begin{cases} \Lambda(b+x+ct, x+cs), & \text{if } k(x,s,t) \geq cs + x; \\ (b+x+(c-\mu)t)^2/(2v(t)), & \text{if } k(x,s,t) < cs + x. \end{cases} \end{aligned}$$

Denote by $s_x^\star$ and $t_x^\star$ the optimizing s and t in (6.25). Similar to the tandem system, we distinguish between (A) the case in which, conditional on a large value of the total queue length, the hp queue will be relatively empty, and (B) the case in which, under the same condition, the hp queue is not automatically small. To this end, we first define

$$L_c(t,x) := \frac{(b+x+(c-\mu)t)^2}{v(t)},$$

which is minimized by $t^F_{c,x}$. We also introduce the set of 'below-critical' service rates $C^F(x)$, i.e., those c such that

$$\sup_{s\in[0,t^F_{c,x}]} \left\{\mathbb{E}\left[A^{\rm h}(-s,0) \mid A^{\rm h}(-t^F_{c,x},0) + A^{\ell}(-t^F_{c,x},0) = b + x + ct^F_{c,x}\right] - cs\right\} \leq x.$$

In order to state our tightness result, we introduce some new notation. For $r_1 < r_2$,

$$\bar{\mathbb{E}}_x A(r_1,r_2) := \mathbb{E}[A(r_1,r_2) \mid A^{\rm h}(-t^\star_x,0) + A^{\ell}(-t^\star_x,0) = b + x + ct^\star_x],$$

with $\bar{\mathbb{V}}\mathrm{ar}_x[\cdot]$ and $\bar{\mathbb{C}}\mathrm{ov}_x[\cdot,\cdot]$ defined similarly. For $r \in (-t^\star_x, 0)$ we define the functions

$$\bar{m}(r,x) := \frac{\bar{\mathbb{E}}_x A^{\rm h}(r,0) + cr}{\sqrt{\bar{\mathbb{V}}\mathrm{ar}_x A(r,0)}}, \qquad m(r',x) := \frac{\bar{m}(r,x)}{\bar{m}(-s^\star_x,x)},$$

and

$$\rho(r,x) := \frac{\bar{\mathbb{C}}\mathrm{ov}_x[A(r,0), A(-s^\star_x,0)]}{\sqrt{\bar{\mathbb{V}}\mathrm{ar}_x A(r,0)\,\bar{\mathbb{V}}\mathrm{ar}_x A(-s^\star_x,0)}}.$$

Theorem 6.4.2 (A) *If $c \in C^F(x)$, then $t^\star_x = t^F_{c,x}$ and $J_p(b,x) = L_c(t^F_c, x)$.*
(B) *If $c \notin C^F(x)$, and if for all $r \in (-t^\star_x, 0)$ it holds that $m(r,x) \leq \rho(r,x)$, then*

$$J_p(b,x) = \inf_{t>0} \sup_{s\in(0,t]} \Upsilon_{b,x}(s,t).$$

The most probable paths can be found in the same fashion as in Section 6.2.3.

6.4.2 Discussion

Large deviations for priority queues have been studied in several papers. We mention here the work in [95] and [139]. We briefly review the results, and compare them with the analysis of this chapter. As expected, it is likely that in general the minimum over x is achieved for $x = 0$. In other words, under this conjecture, the minimization can be done over all paths in

$$\mathcal{A}_{b,0} := \{f \mid \exists t > 0 : \forall s \in (0,t] : A^{\rm h}[f](-t,0) \leq ct,$$
$$A^{\rm h}[f](-t,0) + A^{\ell}[f](-t,0) \geq b + ct\}.$$

Our lower bound then reads

$$J_p^{(\mathrm{I})}(b) := \inf_{t\geq 0} \sup_{s\in(0,t]} \Upsilon_{b,0}(s,t).$$

Just as we did, in [95] two cases are identified. The same solution is found for our case (A) above, i.e., the situation in which, given a long total queue length, the hp queue is relatively short, see also the empty-buffer approximation in [11]. In case (B) the hp queue tends to be large, given that the total queue is long.

In order to prevent this, [95] proposes a heuristic that minimizes the decay rate over

$$\{f \mid \exists t > 0 : A^{\rm h}[f](-t,0) \le ct,\ A^{\rm h}[f](-t,0) + A^{\ell}[f](-t,0) \ge b + ct\}. \quad (6.26)$$

It is not hard to see that the resulting decay rate, denoted by $K_p^{(\rm II)}(b)$, yields a lower bound: set (6.26) is contained in $\mathcal{A}_{b,0}$. As our approach maximizes over $s \in (0,t]$, whereas [95] picks $s = t$, our lower bound will be closer to the real decay rate:

$$J_p^{(\rm II)}(b) := \inf_{t \ge 0} \Upsilon_{b,0}(t,t) \le \inf_{t>0} \sup_{s\in(0,t]} \Upsilon_{b,0}(s,t) = J_p^{(\rm I)}(b).$$

In the simulation experiments performed in [95], the lower bound $J_p^{(\rm II)}(b)$ is usually close to the exact value. Our numerical experiments (analogous to the examples on the tandem queue in Section 6.3) show that the hp buffer usually starts to fill shortly after the total queue starts its busy period. This means that in many cases the error made by taking $s = t$ is relatively small. It explains why the heuristic based on set (6.26) performs well.

The results in [139] are for discrete time, and allow more general traffic than just Gaussian sources. Translated into continuous time, in regime (B), the lower bound on the decay rate in [139, Theorem 14] minimizes the decay rate over

$$\{f \mid \exists t > 0 : \exists s \in (0,t] : A^{\rm h}[f](-s,0) \le cs,\\ A^{\rm h}[f](-t,0) + A^{\ell}[f](-t,0) \ge b + ct\}, \quad (6.27)$$

resulting in the lower bound $J_p^{(\rm III)}(b)$. A straightforward comparison gives that $\mathcal{A}_{b,0}$ is contained in (6.27). Again, our lower bound $J_p^{(\rm I)}(b)$ is closer to the actual decay rate:

$$J_p^{(\rm III)}(b) := \inf_{t\ge 0} \inf_{s\in(0,t]} \Upsilon_{b,0}(s,t) \le \inf_{t\ge 0} \sup_{s\in(0,t]} \Upsilon_{b,0}(s,t) = J_p^{(\rm I)}(b).$$

Remark 6.4.1 *Just like in the tandem case, there are model instances for which the lower bound $J_p^{(\rm I)}(b)$ is not tight. This discrepancy is due to the (conservative) estimation of the decay rate of the intersection of a number of events by the supremum over the individual decay rates. A better lower bound is obtained by imposing a tighter approximation on the intersection:*

$$\inf_{t>0} \sup_{0<s_1\le\cdots\le s_n\le t} \inf_{f\in\mathcal{A}_{b,0}^{\bar{s},t}} I(f),$$

with

$$\bar{s} := (s_1, \ldots, s_n),$$

and

$$\mathcal{A}_{b,0}^{\bar{s},t} := \{f \mid A^{\rm h}[f](-t,0) + A^{\ell}[f](-t,0) \ge b + ct,\\ A^{\rm h}[f](-s_i,0) \le cs_i \text{ for } i = 1,\ldots,n\}.$$

Clearly, the higher n, the better the intersection is approximated. In fact, Theorem 6.4.2 gives sufficient conditions under which $n = 1$ suffices. If $n = 1$ does not work (i.e., the condition in Theorem 6.4.2 is not satisfied, e.g., for fBm input), it is not clear a priori what value of n gives tightness.

6.5 Concluding remarks

A direction for further research is the analysis for the Gaussian models that do not satisfy the condition under which tightness holds. An example of such a model is fractional Brownian motion. As the analysis for the two-node tandem network already causes a lot of problems, we started the analysis of a similar, but slightly easier model: the busy-period analysis in the single queue. In [101] an upper and a lower bound are derived for the logarithmic asymptotics of the probability that the busy period of a single queue with fBm input gets large. As a first step towards the exact asymptotic decay rate in the tandem model we consider the exact asymptotic decay rate for the busy period (see [87]).

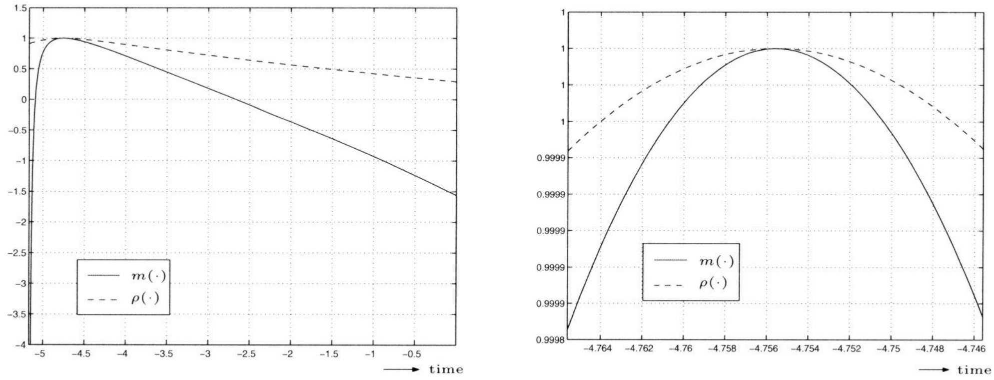

Figure 6.2: $m(\cdot)$ and $\rho(\cdot)$ for M/exp/∞ input process.

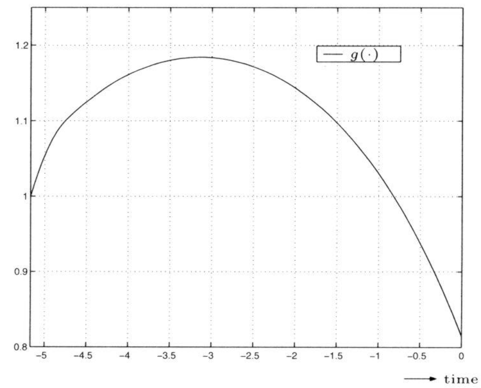

Figure 6.3: Input rate path for M/exp/∞ input process.

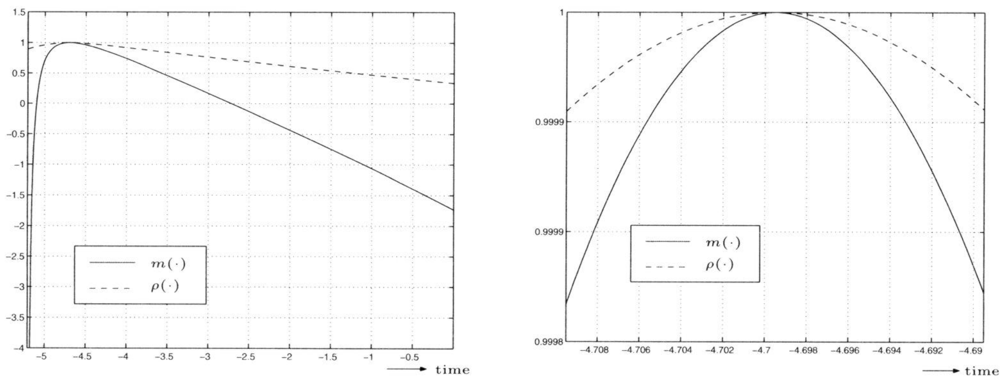

Figure 6.4: $m(\cdot)$ and $\rho(\cdot)$ for M/H2/∞ input process.

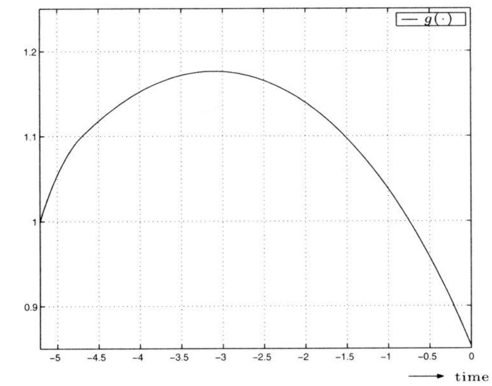

Figure 6.5: Input rate path for M/H2/∞ input process.

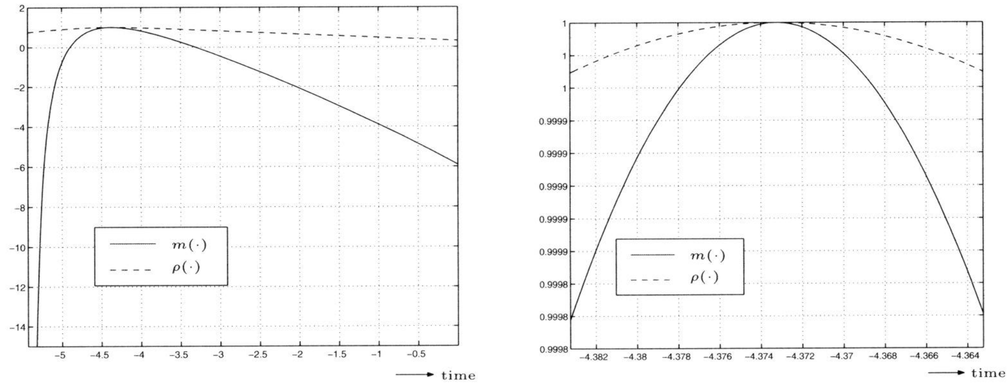

Figure 6.6: $m(\cdot)$ and $\rho(\cdot)$ for M/Par/∞ input process.

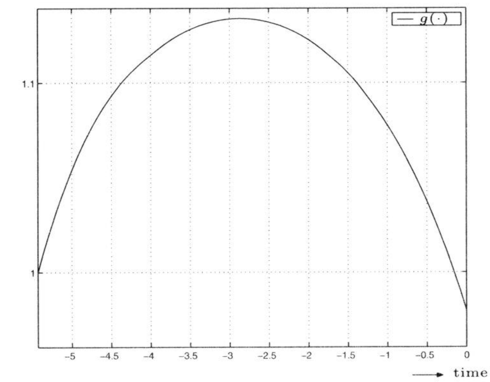

Figure 6.7: Input rate path for M/Par/∞ input process.

Chapter 7

GPS with Gaussian inputs

In this chapter we extend the results of the previous chapter and consider a two-class GPS system. Similar to the priority model, we assume both classes to consist of n i.i.d. Gaussian sources. The variance function and the mean rate of the sources in one class can differ from those of the sources in the other class.

Only few large-deviations results are known for queues operating under a non-FIFO scheduling discipline. In [95] the study of the priority mechanism with many Gaussian inputs was initiated, and in [96] that of the GPS discipline with many Gaussian inputs. In both papers useful intuition and heuristics were developed. The results in [95] on the priority mechanism were sharpened in the previous chapter. The main goal of this chapter is to obtain similar rigorous many-sources large-deviations results for the two-class GPS system. We recall that more results have been derived for GPS in the large-buffer regime (rather than the many-sources regime), see Chapters 3, 4 and 5 and the references in Section 1.6.

The results of this chapter can be summarized as follows. In the first place, we derive upper and lower bounds for the overflow probabilities. These are generic in that they do not only apply to Gaussian inputs, but in fact to any traffic source model. Then we evaluate these bounds in the many-sources framework, i.e., we derive their exponential decay rates (in the number of sources n), after rescaling the link speed $C \equiv nc$ as well as the buffer threshold $B \equiv nb$. We do this by using large-deviations machinery, in particular the multi-dimensional version of the classical Cramér result for sample means, and the pathwise large-deviations principle of Schilder. We then prove tightness of the derived bounds under certain conditions, and present intuitive arguments why tightness can be expected more generally. Finally we address the problem of finding appropriate weights of the GPS queues. This problem should be solved, in order to fully benefit from GPS. It is however not a straightforward task, and not much literature on this subject has appeared so far (the only attempts were [60, 80]). We focus on the operational issue of finding weights such that the QoS requirement is met for all combinations of sources within some predefined region.

This chapter is organized as follows. Section 7.1 deals with preliminaries on GPS, Gaussian sources, and large deviations. Section 7.2 presents the generic upper and lower bounds on the overflow probability of, say (without loss of generality), queue 1. We first focus on the regime in which the mean rate of the type-2 sources, $n\mu_2$, is below their guaranteed rate $n\phi_2 c$; lower and upper bounds on the decay rate are derived in Sections 7.3 and 7.4. Section 7.5 deals with the somewhat easier case $n\mu_2 \geq n\phi_2 c$. A discussion of the results is given in Section 7.6. It turns out that three generic regimes can be distinguished. Section 7.7 addresses weight setting algorithms.

7.1 Model and preliminaries

In Section 7.1.1 we introduce the two-class GPS model with some additional notation. Then we discuss in Section 7.1.2 Gaussian sources. The multivariate version of the large-deviations theorem of Schilder will be presented in Section 7.1.3.

7.1.1 Generalized Processor Sharing

We consider a system where traffic is served according to the GPS mechanism, consisting of two queues sharing a link of capacity nc. We assume the system to be fed by traffic from two classes, where class i uses queue i, for $i = 1, 2$. Without loss of generality it is assumed that both classes consist of n flows (see Remark 7.1.1). We assign a weight $\phi_i \geq 0$ to class i and, again without loss of generality, assume that these add up to 1, i.e., $\phi_1 + \phi_2 = 1$.

Without loss of generality, we focus on the workload of the first queue. The goal of this chapter is to derive the logarithmic asymptotics for the probability that the stationary workload exceeds a threshold nb. Denoting by $Q_{i,n} \equiv Q_{i,n}(0)$ the stationary workload in the i-th GPS queue at time 0, the probability of interest reads

$$\mathbb{P}(Q_{1,n} \geq nb).$$

We denote by $A_{j,i}(s,t)$ the amount of traffic generated by the j-th flow of class i in the interval $(s,t]$, $j = 1, \ldots, n$, $i = 1, 2$. Defining $B_{i,n}(s,t)$ as the total service that was available for class i in the interval $(s,t]$, we introduce the following identity (see also (1.2)):

$$Q_{i,n}(t) = Q_{i,n}(s) + \sum_{j=1}^{n} A_{j,i}(s,t) - B_{i,n}(s,t), \ \forall s < t, \text{ with } s,t \in \mathbb{R}. \tag{7.1}$$

The supremum representation is then given by:

$$Q_{i,n}(0) = \sup_{t>0} \left\{ \sum_{j=1}^{n} A_{j,i}(-t,0) - B_{i,n}(-t,0) \right\}, \tag{7.2}$$

where minus the optimizing t corresponds to the beginning of the busy period that includes time 0. In Section 7.2 we rewrite our problem in terms of the *empirical mean processes* $n^{-1}\sum_{j=1}^{n} A_{j,i}(\cdot,\cdot)$, with $i = 1, 2$. The realization of $n^{-1}\sum_{j=1}^{n} A_{j,i}(0,r)$ is then defined by $f_i(r)$, i.e., $f_i(\cdot)$ is the *path* of the empirical mean process of class i. By $A_i[f_i](s,t)$ we then denote the value of $n^{-1}\sum_{j=1}^{n} A_{j,i}(s,t)$ for the (given) path $f_i(\cdot)$, i.e., $A_i[f_i](s,t) := f_i(t) - f_i(s)$. For notational convenience we use $f(\cdot)$ to denote the two-dimensional path $(f_1(\cdot), f_2(\cdot))$.

7.1.2 Gaussian processes

We assume the n flows of class i to be i.i.d. Gaussian processes with stationary increments. Let $A_{j,i}(s,t)$ be distributed as $A_i(s,t)$, where $A_i(s,t)$ can be considered as the 'generic' random variable corresponding to the amount of traffic of a single class-i flow arriving in the interval $(s,t]$. We denote the corresponding mean traffic rate and variance *function* by μ_i and $v_i(\cdot)$ respectively: for all $s < t$, $\mathbb{E}A_i(s,t) = \mu_i(t-s)$ and $\mathbb{V}\text{ar}A_i(s,t) = v_i(t-s)$. We also define the aggregate mean rate $\mu := \mu_1 + \mu_2$ and the aggregate variance function $v(\cdot) := v_1(\cdot) + v_2(\cdot)$. To guarantee stability, we assume that $\mu < c$. To apply Schilder's theorem (see Theorem 6.1.2 for the one-dimensional version and Theorem 7.1.1 for the multivariate version), we also need to introduce the *centered* process $\bar{A}_i(t) := A_i(0,t) - \mu_i t$. The covariance function $\Gamma_i(s,t)$ is for all $s < t$ defined by

$$\begin{aligned}\Gamma_i(s,t) &:= \mathbb{C}\text{ov}[A_i(0,s), A_i(0,t)] = \mathbb{C}\text{ov}[\bar{A}_i(s), \bar{A}_i(t)] \\ &= \frac{1}{2}(v_i(s) + v_i(t) - v_i(t-s)).\end{aligned}$$

We make the same assumptions on the variance function as in the previous chapter, and restate them for convenience.

Assumption 7.1.1 *We assume that, for $i = 1, 2$, (A1) $v_i(\cdot)$ is continuous, differentiable on $(0,\infty)$; (A2) $\sqrt{v_i(\cdot)}$ is strictly increasing and strictly concave; (A3) for some $\alpha < 2$ it holds that $v_i(t)t^{-\alpha} \to 0$ as $t \to \infty$.*

Assumptions (A1) and (A3) are required to apply Schilder's theorem. Assumption (A2) is needed in the proof of Lemma 7.4.4.

Remark 7.1.1 *Above we assumed that both classes consist of n sources, but the analysis can be easily extended to the case of an unequal number of sources. The scenario with class i having $n\beta_i$ flows, mean μ_i and variance $v_i(\cdot)$ is equivalent to a scenario where class i has n flows, mean $\beta_i\mu_i$ and variance $\beta_i v_i(\cdot)$, due to the infinite divisibility of the Gaussian distribution, cf. Section 6.4.*

7.1.3 Sample-path large deviations

The analysis in this chapter relies on a sample-path LDP for (centered) Gaussian processes. This section is devoted to a brief description of the multivariate version of the main theorem in this field, the generalized version of *Schilder's* *theorem* [9]. For the one-dimensional version of Schilder's theorem and the multivariate version of the well-known Cramér's theorem we refer to Section 6.1.4.

We restrict ourself to the aspects that are relevant in the present study; for more details we refer to [4, 95, 103]. Consider n i.i.d. centered Gaussian processes $\bar{A}_{j,i}(\cdot)$, with stationary increments and covariance $\mathbb{C}\mathrm{ov}[\bar{A}_{j,i}(s), \bar{A}_{j,i}(t)] = \Gamma_i(s,t)$. In the remainder of this section $i = 1, 2$. Define, the path space Ω_i as

$$\Omega_i := \left\{ \omega_i : \mathbb{R} \to \mathbb{R},\ \text{continuous},\ \omega_i(0) = 0,\ \lim_{t\to\infty} \frac{\omega_i(t)}{1+t} = \lim_{t\to-\infty} \frac{\omega_i(t)}{1+t} = 0 \right\},$$

which is a separable Banach space by imposing a specific norm, as explained in [95]. We adhere to the approach in [95] by choosing $\Omega = \Omega_1 \times \Omega_2$ as our path space, where $\left\{\bar{A}_{j,1}(\cdot)\right\}_{j=1}^n$ and $\left\{\bar{A}_{j,2}(\cdot)\right\}_{j=1}^n$ are independent.

Next we introduce and define the *reproducing kernel Hilbert space* $R_i \subseteq \Omega_i$, which has elements that are roughly as smooth as the covariance functions $\Gamma_i(s, \cdot)$. We start from a 'smaller' space S_i, defined by

$$S_i := \left\{ \omega_i : \mathbb{R} \to \mathbb{R},\ \omega_i(\cdot) = \sum_{j=1}^{n} a_{j,i}\Gamma_i(s_j, \cdot),\ \ a_{j,i}, s_j \in \mathbb{R}, j = 1, \dots, n; n \in \mathbb{N} \right\}.$$

The inner product on this space S_i is, for $\omega_{a,i}, \omega_{b,i} \in S_i$, defined as

$$\langle \omega_{a,i}, \omega_{b,i} \rangle_{R_i} := \left\langle \sum_{j=1}^{n} a_{j,i}\Gamma_i(s_j, \cdot), \sum_{k=1}^{n} b_{k,i}\Gamma_i(s_k, \cdot) \right\rangle_{R_i} = \sum_{j=1}^{n}\sum_{k=1}^{n} a_{j,i} b_{k,i} \Gamma_i(s_j, s_k); \tag{7.3}$$

notice that this implies $\langle \Gamma_i(s, \cdot), \Gamma_i(\cdot, t) \rangle_{R_i} = \Gamma_i(s,t)$. Similar to the reproducing kernel property in the previous chapter, we can define this property for the multivariate case as follows. For $\omega_i \in S_i$:

$$\omega_i(t) = \sum_{j=1}^{n} a_{j,i}\Gamma_i(s_j, t) = \left\langle \sum_{j=1}^{n} a_{j,i}\Gamma_i(s_j, \cdot), \Gamma_i(t, \cdot) \right\rangle_{R_i} = \langle \omega_i(\cdot), \Gamma_i(t, \cdot) \rangle_{R_i}. \tag{7.4}$$

We now define the norm $||\omega_i||_{R_i} := \sqrt{\langle \omega_i, \omega_i \rangle_{R_i}}$. The closure of S_i under this norm is defined as the space R_i. Because we have assumed the processes $\bar{A}_{j,1}(\cdot)$ and $\bar{A}_{j,2}(\cdot)$ to be independent, we can define the reproducing kernel Hilbert space of the bivariate process $(\bar{A}_{j,1}(\cdot), \bar{A}_{j,2}(\cdot))$ by $R := R_1 \times R_2$. The inner product in R, with $\omega_{a,i}, \omega_{b,i} \in R_i$, obviously reads

$$\langle (\omega_{a,1}, \omega_{a,2}), (\omega_{b,1}, \omega_{b,2}) \rangle_R = \langle \omega_{a,1}, \omega_{b,1} \rangle_{R_1} + \langle \omega_{a,2}, \omega_{b,2} \rangle_{R_2}.$$

Similarly, the reproducing kernel property can be extended to the space R,

$$\langle (\omega_{a,1}, \omega_{a,2}), (\Gamma_1(s, \cdot), \Gamma_2(t, \cdot) \rangle_R = \omega_{a,1}(s) + \omega_{a,2}(t).$$

Now we can define the rate function of the sample-path LDP by

$$I(\omega) := \begin{cases} \frac{1}{2}||\omega||_R^2 & \text{if } \omega \in R; \\ \infty & \text{otherwise.} \end{cases} \tag{7.5}$$

Under assumptions (A1) and (A3) the following sample-path LDP holds.

Theorem 7.1.1 [Generalized Schilder] $n^{-1}\sum_{j=1}^{n}\bar{A}_{j,i}(\cdot)$ *satisfies the following LDP:*
(a) for any closed set $F \subset \Omega$,

$$\limsup_{n\to\infty}\frac{1}{n}\log\mathbb{P}\left(\frac{1}{n}\sum_{j=1}^{n}\bar{A}_{j,i}(\cdot)\in F\right)\leq -\inf_{\omega\in F}I(\omega);$$

(b) for any open set $G \subset \Omega$,

$$\liminf_{n\to\infty}\frac{1}{n}\log\mathbb{P}\left(\frac{1}{n}\sum_{j=1}^{n}\bar{A}_{j,i}(\cdot)\in G\right)\geq -\inf_{\omega\in G}I(\omega).$$

7.2 Generic upper and lower bound on the probability

As we explained, in a GPS framework the workloads of the queues are intimately related: it is not possible to write down an explicit expression for $Q_{i,n}(0)$, for $i = 1, 2$, without using the workload in the other queue at time instances before 0. In this section we derive explicit upper and lower bounds for $Q_{1,n}(0)$ in terms of the processes $\sum_{j=1}^{n} A_{j,i}(\cdot,\cdot)$, $i = 1, 2$.

In the remainder of this chapter, we have to distinguish between two regimes. The most involved regime is $\mu_2 < \phi_2 c$, which we refer to as *underload for class 2*. In this regime, class 2 is stable regardless of the behavior of the other class. The other regime is $\mu_2 \geq \phi_2 c$, the regime where class 2 is said to be in *overload*. Although the bounds that are derived in this section hold for both regimes, they are only useful in the regime with underload for class 2 – they will be exploited in Sections 7.3 and 7.4. The analysis for the regime with class 2 in overload is presented in Section 7.5.

Note that the results in this section hold regardless of the distribution of the inputs. We mention that bounds similar to the ones that we apply in the next lemmas, have been used in Chapters 3 and 4, and in [146].

Trivially, we can rewrite the overflow probability to

$$\begin{aligned}&\mathbb{P}\left(Q_{1,n}(0)\geq nb\right)\\ =\;&\mathbb{P}\left(\bigcup_{x\geq 0}\{Q_{1,n}(0)+Q_{2,n}(0)\geq nx+nb,\,Q_{2,n}(0)\leq nx\}\right).\end{aligned}\tag{7.6}$$

Because of the work-conserving nature of GPS, it is easily seen that the following relation holds for the total queue:

$$Q_{1,n}(0)+Q_{2,n}(0)=\sup_{t>0}\left\{\sum_{j=1}^{n}(A_{j,1}(-t,0)+A_{j,2}(-t,0))-nct\right\}.\tag{7.7}$$

Substituting this relation for $Q_{1,n}(0)+Q_{2,n}(0)$ in the right-hand side of (7.6), we find

$$\begin{aligned}&\mathbb{P}\left(Q_{1,n}(0)\geq nb\right)\\ =\ &\mathbb{P}\Big(\bigcup_{x\geq 0}\Big\{\sup_{t>0}\Big\{\sum_{j=1}^{n}(A_{j,1}(-t,0)+A_{j,2}(-t,0))-nct\Big\}\geq nx+nb,\\ &\qquad Q_{2,n}(0)\leq nx\Big\}\Big).\end{aligned} \tag{7.8}$$

We denote the optimizing t in the above supremum by $t^\star$. Following [117], $-t^\star$ can be interpreted as the beginning of the busy period of the total queue containing time 0. Next we consider $Q_{2,n}(0)$. Let us denote by $-s^\star$ the beginning of the busy period of queue 2 containing time 0. Then clearly, $s^\star\in[0,t^\star]$, since the busy period of the total queue always starts before or at the start of the busy period of queue 2. Now using the supremum relation (7.2), we obtain

$$Q_{2,n}(0)=\sup_{s\in(0,t^\star]}\left\{\sum_{j=1}^{n}A_{j,2}(-s,0)-B_{2,n}(-s,0)\right\}. \tag{7.9}$$

In order to find bounds for $\mathbb{P}(Q_{1,n}(0)\geq nb)$, it follows from (7.8) that we need to bound the class-2 workload at time 0, $Q_{2,n}(0)$. Given its representation in (7.9), this means that we have to find bounds on the service that was available for class 2 during the busy period containing time 0.

We introduce the following additional notation. We define the events

$$\mathcal{E}_n:=\left\{\begin{array}{l}\exists x\geq 0,t>0:\forall s\in(0,t]:\\ \quad n^{-1}\sum_{j=1}^{n}(A_{j,1}(-t,0)+A_{j,2}(-t,0))\geq x+b+ct,\\ \quad n^{-1}\sum_{j=1}^{n}A_{j,2}(-s,0)\leq x+\phi_2cs\end{array}\right\};$$

$$\mathcal{F}_n:=\left\{\begin{array}{l}\exists x\geq 0,t>0:\forall s\in(0,t]:\exists u\in[0,s):\\ \quad n^{-1}\sum_{j=1}^{n}(A_{j,1}(-t,0)+A_{j,2}(-t,0))\geq x+b+ct,\\ \quad n^{-1}\sum_{j=1}^{n}(A_{j,2}(-s,0)+A_{j,1}(-s,-u))\leq x+\phi_1cu-cs\end{array}\right\}.$$

In the next lemmas we derive the lower and upper bound for the overflow probability of class 1.

Lemma 7.2.1

$$\mathbb{P}\left(Q_{1,n}(0)\geq nb\right)\geq\mathbb{P}\left(\mathcal{E}_n\right).$$

Proof Recall that $-s^\star$ denotes the beginning of the busy period of queue 2 that contains time 0. Hence, the workload of class 2 is positive in the interval $(-s^\star,0]$, indicating that class 2 claims at least its guaranteed rate in this interval: $B_{2,n}(-s^\star,0)\geq n\phi_2cs^\star$. Using this lower bound in (7.9), we derive

$$Q_{2,n}(0)\leq\sup_{s\in(0,t^\star]}\left\{\sum_{j=1}^{n}A_{j,2}(-s,0)-\phi_2ncs\right\}. \tag{7.10}$$

The lower bound for $\mathbb{P}(Q_{1,n}(0) \geq nb)$ is now found by substituting (7.10) for $Q_{2,n}(0)$ in (7.8). □

Lemma 7.2.2

$$\mathbb{P}\left(Q_{1,n}(0) \geq nb\right) \leq \mathbb{P}\left(\mathcal{F}_n\right).$$

Proof Recall that $s^\star$ optimizes (7.9). From (7.9) it follows that we need an upper bound for $B_{2,n}(-s^\star, 0)$. We distinguish between two scenarios: (i) queue 1 is strictly positive during $(-s^\star, 0]$ and (ii) queue 1 has been empty at some time in $(-s^\star, 0]$.

(i) Since both queues are strictly positive during $(-s^\star, 0]$, both classes claim their guaranteed rate, i.e., $B_{2,n}(-s^\star, 0) = n\phi_2 cs^\star$.

(ii) Trivially, we have $B_{2,n}(-s^\star, 0) \leq ncs^\star - B_{1,n}(-s^\star, 0)$. Bearing in mind that queue 1 has been empty in $(-s^\star, 0]$, we define $u^\star := \inf\{u \in [0, s^\star) : Q_{1,n}(-u) = 0\}$. Hence both queues were strictly positive during $(-u^\star, 0]$, and consequently both classes are assigned their guaranteed rates. Together with (7.1) this yields

$$\begin{aligned} B_{1,n}(-s^\star, 0) &= B_{1,n}(-s^\star, -u^\star) + B_{1,n}(-u^\star, 0) \\ &= Q_{1,n}(-s^\star) + \sum_{j=1}^{n} A_{j,1}(-s^\star, -u^\star) + n\phi_1 cu^\star \\ &\geq \inf_{u \in [0, s^\star)} \left\{ \sum_{j=1}^{n} A_{j,1}(-s^\star, -u) + n\phi_1 cu \right\}. \end{aligned}$$

This implies

$$B_{2,n}(-s^\star, 0) \leq ncs^\star - \inf_{u \in [0, s^\star)} \left\{ \sum_{j=1}^{n} A_{j,1}(-s^\star, -u) + n\phi_1 cu \right\}. \tag{7.11}$$

As the right-hand side of (7.11) is larger than $n\phi_2 cs^\star$, we derive

$$B_{2,n}(-s^\star, 0) \leq ncs^\star - \inf_{u \in [0, s^\star)} \left\{ \sum_{j=1}^{n} A_{j,1}(-s^\star, -u) + n\phi_1 cu \right\}.$$

We now use this upper bound in (7.9) to obtain

$$\begin{aligned} &Q_{2,n}(0) \\ \geq\ & \sup_{s \in (0,t]} \left\{ \sum_{j=1}^{n} A_{j,2}(-s, 0) - ncs + \inf_{u \in [0,s)} \left\{ \sum_{j=1}^{n} A_{j,1}(-s, -u) + n\phi_1 cu \right\} \right\}. \end{aligned}$$

Substituting this for $Q_{2,n}(0)$ in (7.8) then yields the desired upper bound. □

Remark 7.2.1 *Compare the sets $\mathcal{E}_n$ and $\mathcal{F}_n$, then evidently $\mathcal{E}_n \subseteq \mathcal{F}_n$. Any path f of the sample-mean process in $\mathcal{F}_n$ defines epochs $s^\star$ and $u^\star$ (as identified in the proof of Lemma 7.2.2). It is not hard to see that if these epochs match, f is also in $\mathcal{E}_n$. From the proof of Lemma 7.2.2, taking $u^\star = s^\star$ means that scenario (i) applies, where queue 1 is strictly positive during the busy period of queue 2 containing time 0. These simple observations turn out to play a crucial role in the discussion presented in Section 7.6.*

7.3 Lower bound on the decay rate: class 2 in underload

Sections 7.3 and 7.4 concern the regime in which class 2 is in underload, i.e., the regime $\mu_2 < \phi_2 c$. In Section 7.3 we first determine the decay rate of the upper bound on $\mathbb{P}(Q_{1,n}(0) \geq nb)$ as presented in Lemma 7.2.2. Then in Section 7.4 we calculate the decay rate of the lower bound on $\mathbb{P}(Q_{1,n}(0) \geq nb)$ as presented in Lemma 7.2.1.

Because of Lemma 7.2.2,

$$-\limsup_{n\to\infty} \frac{1}{n} \log \mathbb{P}\left(Q_{1,n}(0) \geq nb\right) \geq -\limsup_{n\to\infty} \frac{1}{n} \log \mathbb{P}\left(\mathcal{F}_n\right).$$

We now investigate the decay rate in the right-hand side of the previous display. Defining the set of paths

$$\mathcal{A}_{b,x}^{s,t,u} := \left\{ f \,\middle|\, \begin{array}{l} A_1[f](-t,0) + A_2[f](-t,0) \geq x + b + ct, \\ A_2[f](-s,0) + A_1[f](-s,-u) \leq x - \phi_1 cu + cs \end{array} \right\},$$

Schilder's sample-path LDP yields

$$-\lim_{n\to\infty} \frac{1}{n} \log \mathbb{P}\left(\mathcal{F}_{\mathbf{n}}\right) = \inf_{x\geq 0} J^L(b,x),$$

where

$$J^L(b,x) := \inf_{t>0} \inf_{f\in \bigcap_{s\in(0,t]} \bigcup_{u\in[0,s)} \mathcal{A}_{b,x}^{s,t,u}} I(f).$$

Notice that the decay rate of a union of events is just the infimum of the individual decay rates. Unfortunately, we do not have such a relation for an intersection of events. However, it is possible to find an explicit lower bound, as presented in the next theorem. We remark that the proof of the second part resembles that of Theorem 6.2.1 in the previous chapter.

Theorem 7.3.1

$$-\limsup_{n\to\infty} \frac{1}{n} \log \mathbb{P}\left(Q_{1,n}(0) \geq nb\right) \geq -\inf_{x\geq 0} J^L(b,x)$$

where

$$J^L(b,x) \geq \inf_{t>0} \sup_{s\in(0,t]} \inf_{u\in[0,s)} \inf_{f\in\mathcal{A}_{b,x}^{s,t,u}} I(f). \tag{7.12}$$

Proof The first claim follows directly from the above. Now consider the second claim. Because for all $s \in (0, t]$, for given t,

$$\bigcap_{r\in(0,t]}\bigcup_{u\in[0,r]} \mathcal{A}_{b,x}^{r,t,u} \subseteq \bigcup_{u\in[0,s)} \mathcal{A}_{b,x}^{s,t,u},$$

we have for all $s \in (0, t]$,

$$\inf_{f\in \bigcap_{s\in(0,t]}\bigcup_{u\in[0,s)} \mathcal{A}_{b,x}^{s,t,u}} I(f) \geq \inf_{f\in \bigcup_{u\in[0,s)} \mathcal{A}_{b,x}^{s,t,u}} I(f).$$

Hence, it also holds for the maximizing s,

$$\inf_{f\in \bigcap_{s\in(0,t]}\bigcup_{u\in[0,s)} \mathcal{A}_{b,x}^{s,t,u}} I(f) \geq \sup_{s\in(0,t]} \inf_{f\in \bigcup_{u\in[0,s)} \mathcal{A}_{b,x}^{s,t,u}} I(f).$$

This implies the second claim. □

7.4 Upper bound on the decay rate: class 2 in underload

This section concentrates on the decay rate of the lower bound on $\mathbb{P}(Q_{1,n}(0) \geq nb)$ as given in Lemma 7.2.1 for the regime $\mu_2 < \phi_2 c$. The procedure turns out to be more involved than that of Section 7.3.

Because of Lemma 7.2.1,

$$-\liminf_{n\to\infty} \frac{1}{n} \log \mathbb{P}\left(Q_{1,n}(0) \geq nb\right) \leq -\liminf_{n\to\infty} \frac{1}{n} \log \mathbb{P}\left(\mathcal{E}_n\right).$$

We now investigate the decay rate in the right-hand side of the previous display. Define the set of paths

$$\mathcal{A}_{b,x}^{s,t} := \{f \mid A_1[f](-t,0) + A_2[f](-t,0) \geq x + b + ct, A_2[f](-s,0) \leq x + \phi_2 cs\}.$$

Similarly to Theorem 7.3.1, Schilder's sample-path LDP yields the following upper bound.

Lemma 7.4.1

$$-\liminf_{n\to\infty} \frac{1}{n} \log \mathbb{P}\left(Q_{1,n}(0) \geq nb\right) \leq \inf_{x\geq 0} J^U(b,x),$$

where

$$J^U(b,x) := \inf_{t>0} \inf_{f\in \bigcap_{s\in(0,t]} \mathcal{A}_{b,x}^{s,t}} I(f).$$

The objective of this section is to prove that, under some assumptions,

$$J^U(b,x) = \inf_{t>0} \sup_{s\in(0,t]} \inf_{f\in\mathcal{A}_{b,x}^{s,t}} I(f). \tag{7.13}$$

Again, because of the fact that an intersection is involved, no explicit expression for $J^U(b,x)$ is available. We therefore take the following approach: we first derive in Section 7.4.1 a lower bound for $J^U(b,x)$, and then in Section 7.4.2 we give conditions under which this lower bound matches the exact value of $J^U(b,x)$.

Remark 7.4.1 *If the optimizing s and u in (7.12) coincide, see also Remark 7.2.1, then the right-hand sides of (7.12) and (7.13) are the same.*

7.4.1 Lower bound on $J^U(b,x)$

The following lemma gives a lower bound for $J^U(b,x)$. We omit its proof because it is similar to that of Theorem 6.2.1 in the previous chapter, and resembles that of the second claim in Theorem 7.3.1.

Lemma 7.4.2

$$J^U(b,x) \geq \inf_{t>0} \sup_{s\in(0,t]} \inf_{f\in\mathcal{A}_{b,x}^{s,t}} I(f).$$

The lower bound in Lemma 7.4.2 can be calculated more explicitly. To this end, we first concentrate on computing the minimum of $I(f)$ over $f \in \mathcal{A}_{b,x}^{s,t}$, for fixed s and t. The result, as stated in Lemma 7.4.3, requires the introduction of two functions. First recall the large-deviations rate function $\Lambda(\cdot,\cdot)$, of the bivariate Normal random variable $(A_1(-t,0)+A_2(-t,0), A_2(-s,0))$, as given in (6.4) in Section 6.1.4 of the previous chapter,

$$\Lambda(y_1,y_2) := \frac{1}{2}\begin{pmatrix} y_1-\mu t \\ y_2-\mu_2 s\end{pmatrix}^{\mathrm{T}} \Sigma(s,t)^{-1}\begin{pmatrix} y_1-\mu t \\ y_2-\mu_2 s\end{pmatrix}$$

with

$$\Sigma(s,t) := \begin{pmatrix} v(t) & \Gamma_2(s,t) \\ \Gamma_2(s,t) & v_2(s)\end{pmatrix}.$$

We also define for $i=1,2$ (notice the similarity with $k(s,t)$ as introduced in Section 6.2.1):

$$k_i(x,s,t) := \mu_i s + \frac{(x+b+(c-\mu)t)}{v(t)}\,\Gamma_i(s,t).$$

Lemma 7.4.3 *For $s\in(0,t]$,*

$$\begin{aligned} &\inf_{f\in\mathcal{A}_{b,x}^{s,t}} I(f) \\ =\ & \Upsilon_{b,x}(s,t) := \begin{cases} \Lambda(x+b+ct, x+\phi_2 cs), & \textit{if } k_2(x,s,t) > x+\phi_2 cs; \\ (x+b+(c-\mu)t)^2/2v(t), & \textit{if } k_2(x,s,t) \leq x+\phi_2 cs.\end{cases}\end{aligned}$$

Proof Using Theorem 6.1.1 (see Section 6.1.4),

$$\inf_{f\in\mathcal{A}_{b,x}^{s,t}} I(f) = \inf_{y_1\geq x+b+ct,\, y_2\leq x+\phi_2 cs} \Lambda(y_1,y_2).$$

Because $\Lambda(y_1, y_2)$ is convex in y_1 and y_2, we can use the Lagrangian to find the infimum over y_1 and y_2:

$$\mathcal{L}(y_1, y_2, \xi_1, \xi_2) = \Lambda(y_1, y_2) - \xi_1(y_1 - x - b - ct) + \xi_2(y_2 - x - \phi_2 cs),$$

with $\xi_1, \xi_2 \geq 0$. Two cases may occur, depending on the specific values of x, s and t. (i) If x, s and t are such that $k_2(x, s, t) > x + \phi_2 cs$, then both constraints are binding, i.e., $y_1 = x + b + ct$ and $y_2 = x + \phi_2 cs$. (ii) If x, s and t are such that $k_2(x, s, t) \leq x + \phi_2 cs$, then only the first constraint is binding, i.e., $y_1 = x + b + ct$, and $y_2 = k_2(x, s, t)$. □

Remark 7.4.2 *Note that $\theta^\star$ in Theorem 6.1.1 is related to the Lagrange multipliers ξ_1 and ξ_2 that are used in the proof of Lemma 7.4.3 (like in the previous chapter we use the notation $\theta_i^\star(s,t)$, $i = 1, 2$). In case (i) $\theta_1^\star(s,t) = \xi_1 > 0$ and $\theta_2^\star(s,t) = -\xi_2 > 0$, whereas in case (ii) $\theta_1^\star(s,t) = \xi_1 > 0$ and $\theta_2^\star(s,t) = -\xi_2 = 0$.*

Observe that $\Upsilon_{b,x}(s,t)$ is continuous at $s \downarrow 0$, i.e.,

$$\Upsilon_{b,x}(0,t) = \frac{(x + b + (c - \mu)t)^2}{2v(t)}.$$

Now Lemmas 7.4.2 and 7.4.3 yield the final lower bound for $J^U(b, x)$, as stated in the next corollary.

Corollary 7.4.1

$$J^U(b, x) \geq \inf_{t>0} \sup_{s \in (0,t]} \Upsilon_{b,x}(s,t).$$

Interpretation of $\Upsilon_{b,x}(s,t)$. The decay rate $I(f)$ can be interpreted as the cost of having a path f, and, likewise, $\Upsilon_{b,x}(s,t)$ as the cost of generating a traffic pattern in the set $\mathcal{A}_{b,x}^{s,t}$.

The proof of Lemma 7.4.3 shows that the *first* constraint, i.e., $y_1 \geq x + b + ct$ is always binding, whereas the *second* constraint, i.e., $y_2 \leq x + \phi_2 cs$, is sometimes binding, depending on the value of $k_2(x, s, t)$ compared to $x + \phi_2 cs$. Observe that $k_2(x, s, t)$ is in fact a conditional expectation (see also (6.13)):

$$k_2(x, s, t) \equiv \mathbb{E}[A_2(-s, 0) \mid A_1(-t, 0) + A_2(-t, 0) = x + b + ct].$$

The two cases of Lemma 7.4.3 can now be interpreted as follows. (i) The optimal value for y_2 is $x + \phi_2 cs$. In this case, $k_2(x, s, t)$, which is the expected value of the amount of traffic sent by class 2 in $(-s, 0]$ given that in total $x + b + ct$ is sent during $(-t, 0]$, is larger than $x + \phi_2 cs$: with high probability the second constraint is *not met* just by imposing the first constraint. In terms of cost, this means that in this regime additional cost is incurred by imposing the second constraint. (ii) The optimal value for y_2 is precisely $k_2(x, s, t)$, and is smaller than $x + \phi_2 cs$: $A_1(-t, 0) + A_2(-t, 0) = x + b + ct$ implies $A_2(-s, 0) > x + \phi_2 cs$ with high probability. Intuitively this means that, given that the first constraint is satisfied, the second constraint is already met, with high probability.

Using this reasoning, it follows after some calculations that $\Upsilon_{b,x}(s,t)$ can be rewritten in a helpful way as shown in the next corollary. The first term accounts for the cost of satisfying the first constraint in $\mathcal{A}_{b,x}^{s,t}$, the second term (which is possibly 0) for the second constraint.

Corollary 7.4.2

$$\begin{aligned}\Upsilon_{b,x}(s,t) &= \frac{(x+b+ct-\mathbb{E}[A_1(-t,0)+A_2(-t,0)])^2}{2\mathbb{V}\mathrm{ar}[A_1(-t,0)+A_2(-t,0)]} \\ &+\frac{\max^2\{\mathbb{E}[A_2(-s,0)\mid A_1(-t,0)+A_2(-t,0)=x+b+ct]-x-\phi_2 cs,0\}}{2\mathbb{V}\mathrm{ar}[A_2(-s,0)\mid A_1(-t,0)+A_2(-t,0)=x+b+ct]}.\end{aligned}$$

Two regimes for ϕ_2. Corollary 7.4.2 implies that

$$\inf_{t>0}\sup_{s\in(0,t]}\Upsilon_{b,x}(s,t)\geq\inf_{t>0}\frac{(x+b+(c-\mu)t)^2}{2v(t)}. \tag{7.14}$$

Let the optimum in the right-hand side be attained in $t^F(x)$. However, in the remainder of this section we suppress x, as x is held fixed. Suppose that for all $s\in(0,t^F]$ it holds that $k_2(x,s,t^F)<x+\phi_2 cs$, then obviously the inequality in (7.14) is tight. This corresponds to a critical weight $\phi_2^{F,U}(x)$ above which there is tightness. This critical value is given by

$$\begin{aligned}\phi_2^{F,U}(x) &:= \inf\left\{\phi_2:\sup_{s\in(0,t^F]}\left\{k_2(x,s,t^F)-x-\phi_2 cs\right\}\leq 0\right\}\\ &\equiv \sup_{s\in(0,t^F]}\frac{k_2(x,s,t^F)-x}{cs}.\end{aligned} \tag{7.15}$$

The resulting two regimes can be intuitively described as follows.

Large ϕ_2. If $\phi_2>\phi_2^{F,U}(x)$, using the interpretation in terms of conditional expectations, the buffer content of queue 2 at time 0 is likely to be below nx. Hence, if in total (at least) $n(x+b+ct^F)$ is sent during $(-t^F,0]$, it is likely that at time 0, the buffer of class 1 has at least value nb.

Small ϕ_2. If $\phi_2<\phi_2^{F,U}(x)$ then the guaranteed rate for class 2 is relatively small, meaning that its buffer content may easily grow. Again (at least) $n(x+b+ct^F)$ has been sent during the interval $(-t^F,0]$, but now it is *not* obvious that most of it goes to the buffer of class 1. Class 2 has to be 'forced' to take *at most* its guaranteed rate during this interval.

7.4.2 Conditions for exactness

As the overflow *behavior* in case of $\phi_2\geq\phi_2^{F,U}(x)$ is essentially different from that in case of $\phi_2<\phi_2^{F,U}(x)$, we will consider in this section the two regimes separately. The procedure followed though will be the same for both regimes. Let us denote the optimizing s and t in Corollary 7.4.1 by $s^\star(x)$ and $t^\star(x)$, respectively. Again we suppress x however, and use the notation $s^\star$ and $t^\star$ in the remainder of this section. First we use Schilder's theorem to determine the most probable path in $\mathcal{A}_{b,x}^{s^\star,t^\star}$ for the regime of ϕ_2 under consideration. Denoting this optimal path by $f^\star$ we then check whether

$$f^\star \in \left(\bigcup_{t\geq 0} \bigcap_{s\in(0,t]} \mathcal{A}_{b,x}^{s,t} \right). \tag{7.16}$$

If so, the optimal path giving rise to the lower bound of Corollary 7.4.1, is in fact the optimal path for $J^U(b,x)$. Consequently, under the condition (7.16), the lower bound and $J^U(b,x)$ coincide. We remark that the optimal path $f^\star$ in regime A will turn out to differ from that in regime B.

Case A: ϕ_2 larger than critical weight

Because of the definition of $\phi_2^{F,U}(x)$, it holds for all $\phi_2 \geq \phi_2^{F,U}(x)$ that

$$\inf_{t>0} \sup_{s\in(0,t]} \Upsilon_{b,x}(s,t) = \frac{(x+b+(c-\mu)t^F)^2}{2v(t^F)},$$

as identified before, i.e., $t^\star = t^F$ and s does not play a role anymore. The next theorem states that, for these ϕ_2, the lower bound on $J^U(b,x)$ as given in Corollary 7.4.1 actually *equals* $J^U(b,x)$. We omit its proof because it strongly resembles the proof of Theorem 6.2.2 in the previous chapter.

Theorem 7.4.1 *If $\phi_2 \geq \phi_2^{F,U}(x)$, then*

$$J^U(b,x) = \inf_{t>0} \sup_{s\in(0,t]} \Upsilon_{b,x}(s,t) = \frac{(x+b+(c-\mu)t^F)^2}{2v(t^F)},$$

and the most probable paths are, for $r \in [-t^F, 0)$,

$$f_1^\star(r) = -\mathbb{E}[A_1(r,0) \mid A_1(-t^F,0) + A_2(-t^F,0) = x+b+ct^F] = -k_1(x,-r,t^F);$$

$$f_2^\star(r) = -\mathbb{E}[A_2(r,0) \mid A_1(-t^F,0) + A_2(-t^F,0) = x+b+ct^F] = -k_2(x,-r,t^F).$$

Case B: ϕ_2 smaller than critical weight

The analysis of this regime is more involved than that of case A. First we will show in the next lemma that in this regime both constraints in Lemma 7.4.3 are met with equality. Its proof is omitted here, as it is along the lines of Lemma 6.2.4 in the previous chapter. Note that Assumptions (A2) and (A3) are used in the proof. Recall that we denote the optimizing s and t in Corollary 7.4.1 by $s^\star$ and $t^\star$.

Lemma 7.4.4 *If $\phi_2 < \phi_2^{F,U}(x)$, then $k_2(x,s^\star,t^\star) > x + \phi_2 c s^\star$.*

The next lemma gives the most probable paths in the set $\mathcal{A}_{b,x}^{s,t}$ for the regime where for given x, s and t we have that $k_2(x,s,t) > x + \phi_2 cs$. We give the most probable paths for $r \in [-t, 0)$, but they can be trivially extended to the entire real axis.

Lemma 7.4.5 *If $k_2(x,s,t) > x + \phi_2 cs$, then, for $r \in [-t, 0)$, the most probable paths in $\mathcal{A}_{b,x}^{s,t}$ are*

$$f_1(r) = -\mathbb{E}[A_1(r,0) \mid A_1(-t,0) + A_2(-t,0) = x + b + ct, A_2(-s,0) = x + \phi_2 cs];$$

$$f_2(r) = -\mathbb{E}[A_2(r,0) \mid A_1(-t,0) + A_2(-t,0) = x + b + ct, A_2(-s,0) = x + \phi_2 cs].$$

Proof This is shown by using the arguments of the proof of Lemma 6.2.5 in the previous chapter. □

Easy calculations show that we can rewrite the paths obtained in Lemma 7.4.5 as

$$\begin{aligned} f_1(r) &= \mu_1 r - \theta_1^\star(s,t)\Gamma_1(-r,t); \\ f_2(r) &= \mu_2 r - \theta_1^\star(s,t)\Gamma_2(-r,t) - \theta_2^\star(s,t)\Gamma_2(-r,s), \end{aligned}$$

where $\theta^\star$ follows from Theorem 6.1.1 (see also Remark 7.4.2). Interestingly, only one covariance function is involved in the most probable path of class 1, meaning that its path will be symmetric around $-\frac{1}{2}t$.

Now we present conditions under which the lower bound of Corollary 7.4.1 matches $J^U(b,x)$, with an approach that is similar to Section 6.2.3. First we introduce new notation. For $r_1 < r_2$,

$$\bar{\mathbb{E}}A_i(r_1,r_2) := \mathbb{E}[A_i(r_1,r_2) \mid A_1(-t^\star,0) + A_2(-t^\star,0) = x + b + ct^\star], \quad i = 1,2,$$

with $\bar{\mathbb{V}}\mathrm{ar}[\cdot]$ and $\bar{\mathbb{C}}\mathrm{ov}[\cdot,\cdot]$ defined similarly. For $r \in (-t^\star, 0)$ we define the functions

$$\bar{m}_x(r) := \frac{\bar{\mathbb{E}}A_2(r,0) - x + \phi_2 cr}{\sqrt{\bar{\mathbb{V}}\mathrm{ar}A_2(r,0)}}, \qquad m_x(r) := \frac{\bar{m}_x(r)}{\bar{m}_x(-s^\star)},$$

$$\rho_x(r) := \frac{\bar{\mathbb{C}}\mathrm{ov}[A_2(r,0), A_2(-s^\star,0)]}{\sqrt{\bar{\mathbb{V}}\mathrm{ar}A_2(r,0)\,\bar{\mathbb{V}}\mathrm{ar}A_2(-s^\star,0)}}.$$

Again, we suppress the x and use the notation $m(\cdot)$ and $\rho(\cdot)$ in the remainder of this section. Both $m(\cdot)$ and $\rho(\cdot)$ attain a maximum 1 at $r = -s^\star$; for $m(\cdot)$ this follows from Corollary 7.4.2 and Lemma 7.4.4; for $\rho(\cdot)$ from the fact that it is a correlation coefficient.

Theorem 7.4.2 *If $\phi_2 < \phi_2^{F,U}(x)$, then*

$$J^U(b,x) = \inf_{t>0} \sup_{s\in(0,t]} \Upsilon_{b,x}(s,t) = \Lambda(x + b + ct^\star, x + \phi_2 cs^\star),$$

under the condition that $m(r) \le \rho(r)$ for all $r \in (-t^\star, 0)$. The corresponding most probable paths are

$$f_1^\star(r) = -\bar{\mathbb{E}}[A_1(r,0) \mid A_2(-s^\star,0) = x + \phi_2 cs^\star];$$

$$f_2^\star(r) = -\bar{\mathbb{E}}[A_2(r,0) \mid A_2(-s^\star,0) = x + \phi_2 cs^\star].$$

Proof We have to show that (7.16) holds. Straightforward calculations show that indeed $A_1[f^\star](-t^\star,0)+A_2[f^\star](-t^\star,0)=x+b+ct^\star$, as desired. Now it remains to be shown that, if we have $m(r)\leq\rho(r)$ for all $r\in(-t^\star,0)$, then $A_2[f^\star](r,0)\leq x-\phi_2cr$ for all $r\in(-t^\star,0)$. This follows immediately from the following (standard) decomposition:

$$\begin{aligned} & A_2[f^\star](r,0)=-f_2^\star(r) \\ = \; & \bar{\mathbb{E}}A_2(r,0)+\frac{\bar{\mathbb{Cov}}[A_2(r,0),A_2(-s^\star,0)]}{\bar{\mathbb{Var}}A_2(-s^\star,0)}\left(x+\phi_2cs^\star-\bar{\mathbb{E}}A_2(-s^\star,0)\right). \end{aligned}$$

The fact that the decay rate now equals $\Lambda(x+b+ct^\star,x+\phi_2cs^\star)$ is due to Lemma 7.4.4. This proves the stated. □

Remark 7.4.3 *Note that the condition $m(r)\leq\rho(r)$ for all $r\in(0,t^\star)$ only involves properties of the class-2 input process. The above Theorem 7.4.2 therefore holds for* any *class-1 Gaussian process with stationary increments.*

Remark 7.4.4 *Following the approach in Section 6.2.4 the optimal input rate paths $g_1^\star(\cdot)$ and $g_2^\star(\cdot)$, which are the first derivatives of $f_1^\star(\cdot)$ and $f_2^\star(\cdot)$, can be calculated. Assuming $v'(0)=0$, these paths exhibit similar properties as those in Section 6.2.4: $g_1^\star(-t^\star)+g_2^\star(-t^\star)=c$ and $g_2^\star(-s^\star)=\phi_2c$. Hence, at time $-t^\star$ the total input rate is c, making the server operate at full capacity. Then at time $-s^\star$ the total input rate of queue 2 is ϕ_2c, meaning that queue 2 starts claiming its guaranteed rate.*

7.5 Analysis of the decay rate: class 2 in overload

In this section we derive the logarithmic asymptotics of $\mathbb{P}(Q_{1,n}(0)\geq nb)$ in the 'overload regime', i.e., $\mu_2\geq\phi_2c$. We define t^o to be the optimizing t in

$$\inf_{t>0}\frac{(b+(\phi_1c-\mu_1)t)^2}{2v_1(t)}.$$

Theorem 7.5.1 *If $\phi_2\leq\mu_2/c$, then*

$$-\lim_{n\to\infty}\frac{1}{n}\log\mathbb{P}(Q_{1,n}(0)\geq nb)=\inf_{t>0}\frac{(b+(\phi_1c-\mu_1)t)^2}{2v_1(t)}. \tag{7.17}$$

The corresponding most probable paths are, for $r\in[-t^o,0)$

$$f_1^\star(r)=-\mathbb{E}[A_1(r,0)\mid A_1(-t^o,0)=b+\phi_1ct^o];$$
$$f_2^\star(r)=-\mathbb{E}[A_2(r,0)\mid A_1(-t^o,0)=b+\phi_1ct^o].$$

Proof We first show that the right-hand side of (7.17) is a lower bound. Denote by $Q_{i,n}^{nc}(0)$ the stationary workload of queue i if it is served in isolation at a constant rate nc. Then we have

$$\mathbb{P}(Q_{1,n}(0)\geq nb)\leq\mathbb{P}\left(\exists t>0:\frac{1}{n}\sum_{j=1}^{n}A_{j,1}(-t,0)\geq b+\phi_1ct\right),$$

due to $Q_{1,n}(0) \leq Q_{1,n}^{n\phi_1 c}(0)$, which shows that the right-hand side of (7.17) is a lower bound.

The upper bound is a matter of computing the rate function of a feasible path. Observe that $f_2^\star(r) = \mu_2 r$. Hence, as $\mu_2 \geq \phi_2 c$, the type-2 sources constantly claim their weight, such that exactly service rate $n\phi_1 c$ is left for the type-1 sources. This obviously leads to overflow in queue 1. The norm of $f_2^\star(\cdot)$ is 0 (i.e., the norm of the centered version of $f_2^\star(\cdot)$ is 0), as these sources are transmitting at mean rate. The rate function corresponding to $f_1^\star(\cdot)$ equals the desired expression. □

7.6 Discussion of the results

In this section we will discuss the results of the previous sections. We identify three regimes for the value of ϕ_2, corresponding to three generic overflow scenarios. Case (i) directly relates to the overload regime of Section 7.5; Cases (ii) and (iii) to the underload regime of Sections 7.3 and 7.4.

For Case (i) our analysis immediately yields the exact decay rate, see Theorem 7.5.1. For Cases (ii) and (iii), however, the situation is more complicated. Theorems 7.3.1, 7.4.1, and 7.4.2 provide *bounds* on the decay rate. We strongly believe, however, that under fairly general conditions these bounds coincide. This claim is justified (1) by heuristic arguments in Section 7.6.1, (2) by extensive numerical experiments, as reported in Section 7.6.2, and (3) by explicit results for the special case of Brownian motion input in Section 7.6.3. In this section we denote

$$J(b) := -\lim_{n\to\infty} \frac{1}{n} \log \mathbb{P}(Q_{1,n}(0) \geq nb),$$

assuming the latter decay rate exists.

7.6.1 Structure of the solution

Case (i): class 2 in overload

This case corresponds to the regime $\phi_2 \leq \mu_2/c$, for which we have shown that the type-2 sources claim their guaranteed rate $n\phi_2 c$ with overwhelming probability, so that overflow in queue 1 resembles overflow in a FIFO queue with link rate $n\phi_1 c$. This principle plays a crucial role in the proof of Theorem 7.5.1. We repeat it here to compare it with Cases (ii) and (iii).

For $\phi_2 \in [0, \mu_2/c]$:

$$J(b) = \inf_{t>0} \frac{(b + (\phi_1 c - \mu_1)t)^2}{2v_1(t)}.$$

Case (ii): class 2 in underload, with ϕ_2 small

As argued in Sections 7.3 and 7.4, in this regime it is not sufficient to require that $n(x+b+ct)$ traffic is generated in t units of time, since, with high probability, a considerable amount of traffic will be left in queue 2. Hence, additional effort is required to ensure that queue 2 stays below nx. Based on heuristic arguments, we present two claims: (A) regarding the optimal values of u and s, which we denote by $u^\star$ and $s^\star$ (observe that $u^\star$ and $s^\star$ as defined in the proof of Lemma 7.2.2 are the same as the optimizing u and s in (7.12)), and (B) regarding the optimal value of x.

Ad Claim (A)*: the optimal values of u and s.* Recall the probabilistic upper bound in Lemma 7.2.2. In the proof of that lemma, $-s^\star$ denotes the beginning of the busy period of the second queue, which contains time 0. Hence, the second queue remains backlogged during the interval $(-s^\star, 0]$ and claims at least its guaranteed rate $n\phi_2 c$, leaving at most rate $n\phi_1 c$ to the first queue. Parallelling the proof of Lemma 7.2.2, two scenarios are possible: In scenario (i) queue 1 was continuously backlogged during $(-s^\star, 0]$, whereas in scenario (ii), queue 1 has been empty after time $-s^\star$, i.e., queue 1 was empty at some time $-u^\star$ during the busy period of queue 2.

Scenario (ii) is not likely to be optimal, for the following reason. As queue 1 was empty at $-u^\star$, it does not benefit from any effort before $-u^\star$; queue 1 has to build up its entire buffer in the interval $(-u^\star, 0]$. Now recall that queue 2 already started to show deviant behavior from time $-s^\star < -u^\star$, claiming its guaranteed rate. However, this additional effort of queue 2 before time $-u^\star$ is of no 'benefit' for queue 1. In order for queue 1 to fully exploit that queue 2 takes its guaranteed rate during $(-s^\star, 0]$, it should be continuously backlogged during this interval, as in scenario (i). We therefore expect that in the most likely scenario $u^\star = s^\star$.

Ad Claim (B)*: the optimal value of x.* We introduced x in the left-hand side of (7.6). From this representation it follows immediately that nx can be interpreted as the amount of traffic left in queue 2 (at the epoch when the total queue size reaches $n(x+b)$).

We argued before that queue 2 has to claim its guaranteed rate during the interval $(-s^\star, 0]$. If a positive amount of traffic is left in queue 2 at time 0, the type-2 sources apparently 'generated too much traffic'; the guaranteed rate could have been claimed with less effort. We therefore expect that in the most likely scenario $x^\star = 0$. Notice that an essential condition here is that $\phi_2 > \mu_2/c$, as otherwise a build-up of traffic in queue 2 would not be 'wasted effort'.

Because of claims (A) and (B), we expect that this regime applies to $\phi_2 \in [\mu_2/c, \phi_2^F]$, with

$$\phi_2^F := \phi_2^{F,U}(0) = \sup_{s \in (0, t^F(0)]} \frac{k_2(0, s, t^F(0))}{cs}.$$

We also define

$$\begin{pmatrix} z_1(t) \\ z_2(s) \end{pmatrix} := \begin{pmatrix} b + (c-\mu)t \\ (\phi_2 c - \mu_2)s \end{pmatrix}.$$

We expect that the following relation holds.

For $\phi_2 \in [\mu_2/c, \phi_2^F]$:

$$J(b) = \frac{1}{2} \inf_{t>0} \sup_{s\in(0,t]} \begin{pmatrix} z_1(t) \\ z_2(s) \end{pmatrix}^{\mathrm{T}} \begin{pmatrix} v_1(t)+v_2(t) & \Gamma_2(s,t) \\ \Gamma_2(s,t) & v_2(s) \end{pmatrix}^{-1} \begin{pmatrix} z_1(t) \\ z_2(s) \end{pmatrix},$$

provided that for all $r \in (-t^\star(0), 0)$ it holds that $m_0(r) \leq \rho_0(r)$.

Case (iii): class 2 in underload, with ϕ_2 large

Here overflow of the total queue implies overflow of queue 1. Consequently we expect the following relation.

For $\phi_2 \in [\phi_2^F, 1]$:

$$J(b) = \inf_{t>0} \frac{(b+(c-\mu)t)^2}{2v(t)}.$$

In the appendix in Section 7.A it is formally shown that this result applies for

$$\phi_2 \in \left[\sup_{x\geq 0} \phi_2^{F,u}(x), 1\right].$$

Arguments similar to claim (B) above (and extensive numerical experiments) however suggest that

$$\sup_{x\geq 0} \phi_2^{F,u}(x) = \phi_2^F.$$

7.6.2 Numerical results

Section 7.6.3 verifies the claims of Section 7.6.1 for the special case of Brownian inputs. Extensive numerical experiments, however, suggest that the claims are valid under considerably more general conditions – we have not found any counterexamples so far. In this section we present two numerical examples.

Example 1

In this example type-1 sources are fBm with $\mu_1 = 0.2$ and $v_1(t) = t^{2H}$, with Hurst parameter $H = 0.75$, whereas type-2 sources are Ornstein-Uhlenbeck (OU) sources with $\mu_2 = 0.3$ and $v_2(t) = t + e^{-t} - 1$. Take $c = 1$ and $b = 1$. Here $\mu_2/c = 0.3$, while numerical computations yield that $\phi_2^F = 0.4914$. Empirically, it turns out that in Case (ii) where $\phi_2 \in [\mu_2/c, \phi_2^F]$ it holds that $m_0(r) \leq \rho_0(r)$ for all $r \in (-t^\star(0), 0)$ (notice the subscript 0 indicating that $x^\star$ will be 0). Hence we can compare the upper and lower bounds. As they turn out to match, we conclude that we found the exact value of the decay rate. Regarding Case (iii) where $\phi_2 \in [\phi_2^F, 1]$, we empirically find that indeed $\sup_{x\geq 0} \phi_2^{F,u}(x) = \phi_2^F$, implying the correctness of the relation that we expected.

A specific example is considered in the left panel of Figure 7.1. There we focus on a situation in which ϕ_2 is in Case (ii): $\phi_2 = 0.4$. Numerical computations yield that $x^\star = 0$, $t^\star = 6.1819$, while $s^\star = u^\star = 5.6853$. The figure shows the traffic rates of both classes as a function of time. The total buffer starts to build up at time $-t^\star$, whereas queue 2 starts a busy period at $-s^\star$. More detailed inspection yields that with these traffic rates, at time 0 the first queue has indeed overflow, whereas the second queue is empty – in other words: the path is feasible.

Example 2

In this example we interchange the two classes of Example 1. Now $\mu_2/c = 0.2$ and $\phi_2^F = 0.7232$. We again find $\sup_{x\geq 0} \phi_2^{F,u}(x) = \phi_2^F$, so that the expected relation for Case (iii) indeed holds.

For Case (ii) however, we do *not* find the exact decay rate. Consider the example $\phi_2 = 0.4$. In the computation of the lower bound we find $x^\star = 0$, $t^\star = 5.0723$, and $s^\star = u^\star = 5.0597$. Again we verified the 'exactness condition', but now we found $r \in (-t^\star, 0)$ such that $m(r) > \rho(r)$ – hence, the upper bound does not hold. The right panel of Figure 7.1 explains what happens. The corresponding input rate path of the fBm sources has a 'dip' at time $-s^\star$. Consequently this path is *not* feasible: it is true that the sources build up $b + ct^\star$ traffic, as desired, but a positive amount of traffic is left in the second queue at time 0.

Despite the fact that our approach does not yield the exact value of the decay rate in Case (ii), it still provides us with useful information. (1) In the first place, we do not have an upper bound, but fortunately the *lower* bound on the decay rate still applies. Such a lower bound corresponds to an upper bound on the probability of interest, which is of practical relevance, as typically communication networks have to be designed such that overflow is sufficiently rare. (2) Numerical experiments showed that the amount of fluid left in the second queue at time 0 is usually extremely small. This makes us believe that the lower bound is relatively close to the exact value. (3) (Rough) full-link approximations, as introduced in [96], optimize over paths f such that there is a $t > 0$ such that $A_1[f](-t, 0) + A_2[f](-t, 0)$ exceeds $b + ct$, while at the same time $A_2[f](-t, 0) \leq \phi_2 ct$. It is easily seen that this procedure provides a more conservative lower bound (as it *a priori* chooses $s = t$). The observations (1), (2), and (3) justify using the lower bound as an approximation in Case (ii), as is done in Section 7.7.1.

7.6.3 Brownian motion input

In this section we consider the special case that both types of sources correspond to Brownian motions: $v_1(t) = \lambda_1 t$, $v_2(t) = \lambda_2 t$. The formulae from the previous section can be evaluated explicitly, as shown in [91]. The result is given below. In particular, in the proof of this result it turns out that both Claim (A) and Claim (B) hold.

Theorem 7.6.1 *Suppose $v_i(t) = \lambda_i t$, $i = 1, 2$. Then, with*

$$\phi_2^F = 1 - \frac{\lambda_1 - \lambda_2}{\lambda_1 + \lambda_2}\left(1 - \frac{\mu_1 + \mu_2}{c}\right) - \frac{\mu_1}{c},$$

it holds that (i) for $\phi_2 \in [0, \mu_2/c]$,

$$J(b) = 2\,\frac{\phi_1 c - \mu_1}{\lambda_1}\,b;$$

(ii) for $\phi_2 \in [\mu_2/c, \phi_2^F]$,

$$t^\star = s^\star = u^\star = \frac{b}{\sqrt{(\phi_1 c - \mu_1)^2 + (\phi_2 c - \mu_2)^2 \frac{\lambda_1}{\lambda_2}}};$$

$$J(b) = \frac{1}{2}\left(\frac{(b + (\phi_1 c - \mu_1)t^\star)^2}{\lambda_1 t^\star} + \frac{(\phi_2 c - \mu_2))^2}{\lambda_2} t^\star\right);$$

(iii) for $\phi_2 \in [\phi_2^F, 1]$,

$$J(b) = 2\,\frac{c - \mu}{\lambda_1 + \lambda_2}\,b.$$

Notice that in all three Cases (i)-(iii) it holds that $J(b)$ is linear in b. For Case (ii) it takes some simple algebra to see this.

7.7 Weight setting

This section focuses on the operational issue of selecting appropriate weights in a two-class GPS system. With any set of weights $\phi \equiv (\phi_1, \phi_2)$ an *admissible region* $\mathcal{S}(\phi)$ can be associated. Such a region contains all combinations (n_1, n_2), with n_i the number of sources of class i, for which the required QoS for both classes is realized. Obviously the size and shape of $\mathcal{S}(\phi)$ depends critically on the weights ϕ chosen.

When selecting appropriate weights, various objectives could be chosen. In this section we investigate two approaches. Following [60], we assume in Section 7.7.2 that, for practical reasons, it should be avoided to switch between a large number of different weights – in fact, we require that just one set of weights ϕ is used. Therefore, we consider the situation that the user population fluctuates just mildly around some 'operating point' $(\bar{n}_1, \bar{n}_2)$, where $\bar{n}_i$ could be for instance the average number of active sources in class i. We develop an algorithm to find a ϕ such that some 'ball' around $\bar{n}_1, brn_2)$ is contained in $\mathcal{S}(\phi)$.

In Section 7.7.3 we take the opposite approach and allow *infinitely many* weight adaptations. We compute the resulting admissible region $\mathcal{S} = \cup_\phi \mathcal{S}(\phi)$. Both Sections 7.7.2 and 7.7.3 require fast and straightforward approximations of the overflow probabilities in the GPS system. We start with these in Section 7.7.1.

7.7.1 Approximation of the overflow probabilities

In this section we develop an approximation for the overflow probabilities in both queues of the GPS system. Recall that for $i = 1, 2$, n_i is the (typically large) number of sources of type i. We denote the stationary buffer content in this GPS model with unequal number of sources by Q_i, the service rate by C, and the buffer threshold of queue i by B_i. Invoking Remark 7.1.1, the GPS model with $n_1 \neq n_2$ is equivalent to a GPS model with n sources in both classes, mean rates $(n_i/n)\mu_i$ and variance functions $(n_i/n)v_i(\cdot)$. We scale the buffer threshold and service rate with n such that $nb_i \equiv B_i$ and $nc \equiv C$. Now we can apply our earlier results, where we assumed both classes to consist of n sources.

To approximate the overflow probabilities, three regimes are distinguished, as in Section 7.6. Again, we concentrate on the first queue; the second queue can be treated analogously. Define $\Delta_i(n_1, n_2) := -\log \mathbb{P}(Q_i \geq B_i)$. Then it holds that $\Delta_i(n_1, n_2) \equiv -\log \mathbb{P}(Q_{i,n} \geq nb_i)$, with its approximation given by $\bar{\Delta}_i(n_1, n_2) := nJ(b_i)$.

We denote the minimizer of

$$\frac{1}{2} \inf_{t>0} \frac{(B_1 + (C - n_1\mu_1 - n_2\mu_2)t)^2}{n_1 v_1(t) + n_2 v_2(t)}$$

by t^F, and we define

$$\phi_2^F := \sup_{s\in[0,t^F]} \left(\frac{\Gamma_2(t^F, s)}{C\, s\, v(t^F)} \right) \left(B_1 + (C - n_1\mu_1 - n_2\mu_2)t^F \right). \tag{7.18}$$

(i) If $\phi_2 \in [0, n_2\mu_2/C]$, then

$$\bar{\Delta}_1(n_1, n_2) = \frac{1}{2} \inf_{t>0} \frac{(B_1 + (\phi_1 C - n_1\mu_1)t)^2}{n_1 v_1(t)}.$$

(ii) If $\phi_2 \in (n_2\mu_2/C, \phi_2^F)$, then

$$\bar{\Delta}_1(n_1, n_2) =$$

$$\frac{1}{2} \inf_{t>0} \sup_{s\in(0,t]} \begin{pmatrix} z_1(t, n_1, n_2) \\ z_2(s, n_1, n_2) \end{pmatrix}^{\mathrm{T}} \begin{pmatrix} n_1 v_1(t) + n_2 v_2(t) & n_2\Gamma_2(s,t) \\ n_2\Gamma_2(s,t) & n_2 v_2(s) \end{pmatrix}^{-1} \begin{pmatrix} z_1(t, n_1, n_2) \\ z_2(s, n_1, n_2) \end{pmatrix},$$

where

$$\begin{pmatrix} z_1(t, n_1, n_2) \\ z_2(s, n_1, n_2) \end{pmatrix} := \begin{pmatrix} B_1 + (C - n_1\mu_1 - n_2\mu_2)t \\ (\phi_2 C - n_2\mu_2)s \end{pmatrix}. \tag{7.19}$$

(iii) If $\phi_2 \in [\phi_2^F, 1]$, then

$$\bar{\Delta}_1(n_1, n_2) = \frac{1}{2} \inf_{t>0} \frac{(B_1 + (C - n_1\mu_1 - n_2\mu_2)t)^2}{n_1 v_1(t) + n_2 v_2(t)}.$$

7.7.2 Weight setting algorithm

This subsection focuses on a procedure for finding weights (ϕ_1, ϕ_2) such that both classes receive the desired QoS, despite (mild) fluctuations in the number of sources present. More precisely, for specified (positive) numbers δ_i, we require that $\Delta_i(n_1, n_2) \geq \delta_i$ for all (n_1, n_2) in a 'ball' $\mathcal{B}(\bar{n}_1, \bar{n}_2)$ around $(\bar{n}_1, \bar{n}_2)$:

$$\mathcal{B}(\bar{n}_1, \bar{n}_2) := \left\{(n_1, n_2) \in \mathbb{N}^2 \mid \gamma_1(n_1 - \bar{n}_1)^2 + \gamma_2(n_2 - \bar{n}_2)^2 \leq 1\right\},$$

for positive γ_1, γ_2. It can be easily verified that the procedure described below works, in fact, for any 'target area' $\mathcal{B}$ that is finite and *convex*, rather than just these ellipsoidal sets.

To simplify our algorithm, we use the following expansion of $\bar{\Delta}_i(n_1, n_2)$ around $(n_1, n_2) = (\bar{n}_1, \bar{n}_2)$:

$$\begin{aligned} \bar{\Delta}_i(n_1, n_2) \;\; &\approx \;\; \bar{\Delta}_i(\bar{n}_1, \bar{n}_2) + (n_1 - \bar{n}_1) \left.\frac{\partial \bar{\Delta}_i(n_1, n_2)}{\partial n_1}\right|_{(n_1,n_2)=(\bar{n}_1,\bar{n}_2)} \\ &+ (n_2 - \bar{n}_2) \left.\frac{\partial \bar{\Delta}_i(n_1, n_2)}{\partial n_2}\right|_{(n_1,n_2)=(\bar{n}_1,\bar{n}_2)} . \end{aligned} \tag{7.20}$$

This approximation requires the evaluation of two partial derivatives, which can be done relatively explicitly, as described in Section 7.B of the appendix. Relying on (7.20), we have to verify whether for all $(n_1, n_2) \in \mathcal{B}(\bar{n}_1, \bar{n}_2)$ and $i = 1, 2$,

$$\bar{\Delta}_i(\bar{n}_1, \bar{n}_2) + (n_1 - \bar{n}_1, n_2 - \bar{n}_2)^{\mathrm{T}} e_i \geq \delta_i,$$

where

$$e_i \equiv (e_{i1}, e_{i2}) := \left(\left.\frac{\partial \bar{\Delta}_i(n_1, n_2)}{\partial n_1}\right|_{(n_1,n_2)=(\bar{n}_1,\bar{n}_2)}, \left.\frac{\partial \bar{\Delta}_i(n_1, n_2)}{\partial n_2}\right|_{(n_1,n_2)=(\bar{n}_1,\bar{n}_2)} \right).$$

Because of the convex shape of $\mathcal{B}(\bar{n}_1, \bar{n}_2)$, we only have to verify this condition for the two points on the boundary $\partial\mathcal{B}(\bar{n}_1, \bar{n}_2)$ having a tangent with slopes equal to $-e_{11}/e_{12}$ and $-e_{21}/e_{22}$ respectively. Denoting these points by $(n_{11}^\star, n_{12}^\star)$ and $(n_{21}^\star, n_{22}^\star)$ we have

$$(n_{i1}^\star, n_{i2}^\star) := \left(\bar{n}_1 + \sqrt{\left(\gamma_1 + \frac{e_{i2}^2}{e_{i1}^2}\frac{\gamma_1^2}{\gamma_2}\right)^{-1}}, \bar{n}_2 + \sqrt{\left(\gamma_2 + \frac{e_{i1}^2}{e_{i2}^2}\frac{\gamma_2^2}{\gamma_1}\right)^{-1}} \right), \quad i = 1, 2.$$

We say that ϕ is feasible if $K_i := \bar{\Delta}_i(\bar{n}_1, \bar{n}_2) + (n_{i1}^\star - \bar{n}_1, n_{i2}^\star - \bar{n}_2)^{\mathrm{T}} e_i \geq \delta_i$ for both class $i = 1$ and 2. Notice that K_i is a function of the weights; as $\phi_1 + \phi_2 = 1$, we can write $K_i(\phi_1)$. $K_1(\phi_1)$ will increase in ϕ_1, whereas $K_2(\phi_1)$ will decrease.

This suggests the following solution approach to the weight setting problem: (i) First find the smallest ϕ_1 such that $K_1(\phi_1) \geq \delta_1$. If this does not exist, then there is no solution. (ii) If it does exist, then verify whether for this ϕ_1 it holds that $K_2(\phi_1) \geq \delta_2$. If this is true, then the weight setting problem can be solved; otherwise there is no solution (i.e., there is no ϕ such that $\mathcal{B}(\bar{n}_1, \bar{n}_2) \subseteq \mathcal{S}(\phi)$).

Example 3

We first explain how requirements on the admissible numbers of sources naturally lead to a set of the type $\mathcal{B}(\bar{n}_1, \bar{n}_2)$.

- Our analysis assumes fixed numbers of sources of both types, but in practice this number fluctuates in time: sources arrive, and stay in the system for a random amount of time. Now suppose that sources of both types arrive according to Poisson processes (with rates λ_i, for $i = 1, 2$), and that, if admitted, these would require service for some random duration (with finite means $\mathbb{E}D_i$). If there were no admission control, the distributions of the number of jobs of both types are Poisson with means (and variances) $\bar{n}_i = \lambda_i \, \mathbb{E}D_i$.
- Suppose the system must be designed such that this mean $(\bar{n}_1, \bar{n}_2) \pm$ twice the standard deviation should be in the admissible region, i.e., should be contained in $\mathcal{S}(\phi)$. This suggests choosing

$$\mathcal{B}(\bar{n}_1, \bar{n}_2) = \left\{ (n_1, n_2) \in \mathbb{N}^2 \;\middle|\; \left(\frac{n_1 - \bar{n}_1}{2\sqrt{\bar{n}_1}}\right)^2 + \left(\frac{n_2 - \bar{n}_2}{2\sqrt{\bar{n}_2}}\right)^2 \leq 1 \right\}.$$

In this example we choose $\bar{n}_1 = 900$ and $\bar{n}_2 = 1600$, which leads to:

$$\begin{aligned} & \mathcal{B}(\bar{n}_1, \bar{n}_2) = \mathcal{B}(900, 1600) \\ = \; & \left\{ (n_1, n_2) \in \mathbb{N}^2 \mid 16(n_1 - 900)^2 + 9(n_2 - 1600)^2 \leq 57600 \right\}. \end{aligned}$$

We suppose that both types of sources correspond to Brownian motions, with $\mu_1 = 0.2$, $\mu_2 = 0.3$, $v_1(t) = 2t$, and $v_2(t) = t$. We rely on explicit results for Brownian motions, as summarized in Section 7.B in the appendix, in particular for the partial derivatives of the $\bar{\Delta}_i(n_1, n_2)$ with respect to the numbers of sources. We choose $C = 1000$, $B_1 = 35$, and $B_2 = 25$.

First suppose the performance targets are $\delta_1 = 9$ and $\delta_2 = 7$ (roughly corresponding to overflow probabilities $1.2 \cdot 10^{-4}$ and $9.1 \cdot 10^{-4}$). Figure 7.2 shows that no weights ϕ exist to meet this target (to guarantee that the overflow probability in queue 1 is small enough, ϕ_1 should be larger than 0.39, but this implies that $K_2(\phi_1) < 5.7 < \delta_2$). Now suppose that $\delta_1 = 8$ and $\delta_2 = 6$. Then an analogous reasoning gives that ϕ_1 should be chosen in the interval $(0.34, 0.37)$.

7.7.3 Admissible region

While above we restricted ourself to just one set of weights, we might allow to switch weights whenever necessary. Clearly, the resulting admissible region can be obtained as the union of the admissible regions for fixed weights.

Example 4

In Figure 7.3 we have computed the admissible region for the same types of sources as in the previous subsection, with performance targets $\delta_1 = 5$, $\delta_2 = 7$.

We have not succeeded in finding explicit expressions for the boundary of the admissible region.

Appendix

7.A Analysis of underload regime with large ϕ_2

This appendix focuses on the underload regime with large ϕ_2. By deriving the counterpart for $J^L(b,x)$ of Theorem 7.4.1, we can prove that for a specific range of ϕ_2 the derived upper and lower bounds match. We first introduce some new notation:

$$\bar{k}_2(x,s,t,u) := \mathbb{E}[A_2(-s,0) + A_1(-s,-u) \mid A_1(-t,0) + A_2(-t,0) = x + b + ct],$$

and

$$\phi_2^{F,L}(x) := \sup_{s\in(0,t^F(x)]} \inf_{u\in[0,s)} \frac{\bar{k}_2(x,s,t^F(x),u) - x + c(u-s)}{cu}.$$

Lemma 7.A.1 *For all $x \geq 0$, it holds that $\phi_2^{F,L}(x) \leq \phi_2^{F,U}(x)$.*

Proof Notice that $\bar{k}_2(x,s,t,u)$ and $k_2(x,s,t)$ coincide for $u = s$. Then the stated follows directly from the definitions of $\phi_2^{F,L}(x)$ and $\phi_2^{F,U}(x)$. □

The counterpart of Theorem 7.4.1 follows directly now.

Lemma 7.A.2 *If $\phi_2 \geq \phi_2^{F,L}(x)$, then*

$$J^L(b,x) = \frac{(x+b+(c-\mu)t^F(x))^2}{2v(t^F(x))}.$$

This leads to the following result.

Theorem 7.A.1 *If*

$$\phi_2 \in \left[\sup_{x\geq 0} \phi_2^{F,u}(x), 1\right],$$

then

$$-\lim_{n\to\infty} \frac{1}{n} \log \mathbb{P}(Q_{1,n}(0) \geq nb) = \inf_{t>0} \frac{(b+(c-\mu)t)^2}{2v(t)}.$$

Proof Due to Lemma 7.A.1, if $\phi_2 \geq \sup_{x\geq 0} \phi_2^{F,U}(x)$, then also $\phi_2 \geq \phi_2^{F,L}(x)$ for all $x \geq 0$. Now the stated follows directly from the fact that, *for all* $x \geq 0$, Theorem 7.4.1 and Lemma 7.A.2 apply. Hence the infima over t and x can be interchanged, and the result follows. □

Our numerical experiments suggest that $\sup_{x\geq 0} \phi_2^{F,U}(x) = \phi_2^{F,U}(0)$. The following shows that this property holds under a sufficient condition that can be verified (relatively) easily. Denote by $s^F(x)$ the optimizing $s \in (0, t^F(x)]$ in (7.15).

Lemma 7.A.3 *If, for all $x \geq 0$,*

$$\frac{v'(t^F(x))}{v(t^F(x))}\,\Gamma_2(s^F(x), t^F(x)) \geq 2\frac{\partial \Gamma_2(s,t)}{\partial t}\bigg|_{s=s^F(x),t=t^F(x)}, \tag{7.21}$$

then $\sup_{x\geq 0}\phi_2^{F,U}(x) = \phi_2^{F,U}(0)$.

Proof We prove that $\phi_2^{F,U}(\cdot)$ is decreasing under condition (7.21). For brevity write

$$\frac{\partial k_2}{\partial s} \equiv \frac{\partial k_2(x,s,t)}{\partial s}\bigg|_{s=s^F(x),t=t^F(x)}, \quad \frac{\partial k_2}{\partial t} \equiv \frac{\partial k_2(x,s,t)}{\partial t}\bigg|_{s=s^F(x),t=t^F(x)}.$$

Notice that $t^F(x)$ and $s^F(x)$ satisfy the following relations

$$\frac{x+b+(c-\mu)t^F(x)}{2(c-\mu)} = \frac{v(t^F(x))}{v'(t^F(x))}, \tag{7.22}$$

and

$$s^F(x)\,\frac{\partial k_2}{\partial s} = k_2(x, s^F(x), t^F(x)) - x. \tag{7.23}$$

It is easy to check that the derivative of $\phi_2^{F,U}(\cdot)$ is non-positive if

$$s^F(x)\left(\frac{\partial k_2}{\partial s}\frac{\mathrm{d}s^F(x)}{\mathrm{d}x} + \frac{\partial k_2}{\partial t}\frac{\mathrm{d}t^F(x)}{\mathrm{d}x} + \frac{\partial k_2}{\partial x} - 1\right)$$
$$-\frac{\mathrm{d}s^F(x)}{\mathrm{d}x}\left(k_2(x, s^F(x), t^F(x)) - x\right) \leq 0.$$

Notice that because of (7.23) various terms cancel out. Now due to

$$\frac{\partial k_2}{\partial x} = \frac{\Gamma_2(s,t)}{v(t)} \leq \frac{\Gamma_2(s,t)}{v_2(t)} \leq \frac{\Gamma_2(s,t)}{\sqrt{v_2(s)v_2(t)}} \leq 1,$$

(apply Assumption (A2)), and $\mathrm{d}t^F(x)/\mathrm{d}x \geq 0$ (see Lemma 3.1 in [85]), it is left to check that $\partial k_2/\partial t \leq 0$. It is a matter of straightforward calculus, using (7.22), to show that this is equivalent to (7.21). □

7.B Weight setting algorithm: partial derivatives

In this part of the appendix, we determine expressions for the partial derivatives of $\bar{\Delta}_1(n_1, n_2)$ to the numbers of sources, as required in the weight setting algorithm of Section 7.7.2. The Cases (i), (ii), (iii) below correspond to the regimes identified in Section 7.7.1. Recall B_1, C, and the definition of ϕ_2^F in (7.18) as given in Section 7.7.2.

(i) Based on Theorem 6.1.1, in the regime $\phi_2 \in [0, \phi_2^F]$, we can rewrite $\bar{\Delta}_1(n_1, n_2)$ as

$$\bar{\Delta}_1(n_1, n_2) = \inf_{t>0} \sup_{\theta \in \mathbb{R}} \left(\theta(B_1 + (\phi_1 C - n_1\mu_1)t) - \frac{1}{2}\theta^2 n_1 v_1(t) \right).$$

The inner supremum is attained for

$$\theta^\star = \frac{B_1 + (\phi_1 C - n_1\mu_1)t}{n_1 v_1(t)}.$$

Denoting the optimizing t by $t^\star$, we derive

$$\frac{\partial \bar{\Delta}_1(n_1, n_2)}{\partial n_1} = -\theta^\star \mu_1 t^\star - \frac{1}{2}(\theta^\star)^2 v_1(t^\star), \qquad \frac{\partial \bar{\Delta}_1(n_1, n_2)}{\partial n_2} = 0.$$

(ii) Similarly, in the regime $\phi_2 \in [\phi_2^F, n_2\mu_2/C]$, $\bar{\Delta}_1(n_1, n_2)$ can be rewritten as (recall (7.19))

$$\inf_{t>0} \sup_{s\in(0,t]} \sup_{\theta\in\mathbb{R}^2} \left(\theta^{\mathrm{T}} \begin{pmatrix} z_1(t, n_1, n_2) \\ z_2(s, n_1, n_2) \end{pmatrix} - \frac{1}{2}\theta^{\mathrm{T}} \begin{pmatrix} n_1 v_1(t) + n_2 v_2(t) & n_2\Gamma_2(s,t) \\ n_2\Gamma_2(s,t) & n_2 v_2(s) \end{pmatrix} \theta \right).$$

The optimizing θ is given by (observe that θ is a two-dimensional vector)

$$\theta^\star = \begin{pmatrix} n_1 v_1(t) + n_2 v_2(t) & n_2\Gamma_2(s,t) \\ n_2\Gamma_2(s,t) & n_2 v_2(s) \end{pmatrix}^{-1} \begin{pmatrix} z_1(t, n_1, n_2) \\ z_2(s, n_1, n_2) \end{pmatrix}.$$

Straightforward computations give that, with the optimizing s and t denoted by $s^\star$ and $t^\star$,

$$\frac{\partial \bar{\Delta}_1(n_1, n_2)}{\partial n_1} = -\theta_1^\star \mu_1 t^\star - \frac{1}{2}(\theta_1^\star)^2 v_1(t^\star),$$

$$\frac{\partial \bar{\Delta}_1(n_1, n_2)}{\partial n_2} = -\theta_1^\star \mu_2 t^\star - \theta_2^\star \mu_2 s^\star - \frac{1}{2}\theta^{\star\mathrm{T}} \begin{pmatrix} v_2(t^\star) & \Gamma_2(s^\star, t^\star) \\ \Gamma_2(s^\star, t^\star) & v_2(s^\star) \end{pmatrix} \theta^\star.$$

(iii) In the third regime $\phi_2 \in [n_2\mu_2/C, 1]$,

$$\bar{\Delta}_1(n_1, n_2) = \inf_{t>0} \sup_{\theta \in \mathbb{R}} \left(\theta z_1(t, n_1, n_2) - \frac{1}{2}\theta^2 (n_1 v_1(t) + n_2 v_2(t)) \right).$$

The inner supremum is attained for

$$\theta^\star = \frac{z_1(t, n_1, n_2)}{n_1 v_1(t) + n_2 v_2(t)}.$$

Denoting the optimizing t by $t^\star$, we derive

$$\frac{\partial\bar{\Delta}_1(n_1,n_2)}{\partial n_1} = -\theta^\star\mu_1 t^\star - \frac{1}{2}(\theta^\star)^2 v_1(t^\star);$$

$$\frac{\partial\bar{\Delta}_1(n_1,n_2)}{\partial n_2} = -\theta^\star\mu_2 t^\star - \frac{1}{2}(\theta^\star)^2 v_2(t^\star).$$

Now we consider the special case that both types of sources correspond to Brownian motions. We assume $v_1(t) = \lambda_1 t$, $v_2(t) = \lambda_2 t$. We again consider the three regimes separately. We have an explicit formula for the critical value of ϕ_2:

$$\phi_2^F = 1 - \frac{n_1\lambda_1 - n_2\lambda_2}{n_1\lambda_1 + n_2\lambda_2}\left(1 - \frac{n_1\mu_1 + n_2\mu_2}{C}\right) - \frac{n_1\mu_1}{C}.$$

(i) In this case

$$t^\star = \frac{B_1}{\phi_1 C - n_1\mu_1}; \quad \bar{\Delta}_1(n_1,n_2) = 2\,\frac{\phi_1 C - n_1\mu_1}{n_1\lambda_1}\,B_1.$$

This yields:

$$\frac{\partial\bar{\Delta}_1(n_1,n_2)}{\partial n_1} = -2B_1\,\frac{\phi_1 C}{n_1^2\lambda_1}; \quad \frac{\partial\bar{\Delta}_1(n_1,n_2)}{\partial n_2} = 0.$$

(ii) In this case

$$t^\star = B_1 \Bigg/ \sqrt{(\phi_1 C - n_1\mu_1)^2 + (\phi_2 C - n_2\mu_2)^2\frac{n_1\lambda_1}{n_2\lambda_2}};$$

$$\bar{\Delta}_1(n_1,n_2) = \frac{1}{2}\left(\frac{(B_1 + (\phi_1 C - n_1\mu_1)t^\star)^2}{n_1\lambda_1 t^\star} + \frac{(\phi_2 C - n_2\mu_2))^2}{n_2\lambda_2}\,t^\star\right).$$

Also $s^\star = t^\star$. This yields:

$$\frac{\partial\bar{\Delta}_1(n_1,n_2)}{\partial n_1} = -(B_1 + (\phi_1 C - n_1\mu_1)t^\star)\,\frac{\mu_1}{n_1\lambda_1} - \frac{1}{2}\frac{(B_1 + (\phi_1 C - n_1\mu_1)t^\star)^2}{n_1^2\lambda_1 t^\star};$$

$$\frac{\partial\bar{\Delta}_1(n_1,n_2)}{\partial n_2} = -(\phi_2 C - n_2\mu_2)t^\star\,\frac{\mu_2}{n_2\lambda_2} - \frac{1}{2}\frac{(\phi_2 C - n_2\mu_2))^2}{n_2^2\lambda_2}\,t^\star.$$

(iii) In this case

$$t^\star = \frac{B_1}{C - n_1\mu_1 - n_2\mu_2}; \quad \bar{\Delta}_1(n_1,n_2) = 2\,\frac{C - n_1\mu_1 - n_2\mu_2}{n_1\lambda_1 + n_2\lambda_2}\,B_1.$$

This yields:

$$\frac{\partial\bar{\Delta}_1(n_1,n_2)}{\partial n_1} = -\frac{2B_1\mu_1}{n_1\lambda_1 + n_2\lambda_2} - 2B_1\lambda_1\,\frac{C - n_1\mu_1 - n_2\mu_2}{(n_1\lambda_1 + n_2\lambda_2)^2};$$

$$\frac{\partial\bar{\Delta}_1(n_1,n_2)}{\partial n_2} = -\frac{2B_1\mu_2}{n_1\lambda_1 + n_2\lambda_2} - 2B_1\lambda_2\,\frac{C - n_1\mu_1 - n_2\mu_2}{(n_1\lambda_1 + n_2\lambda_2)^2}.$$

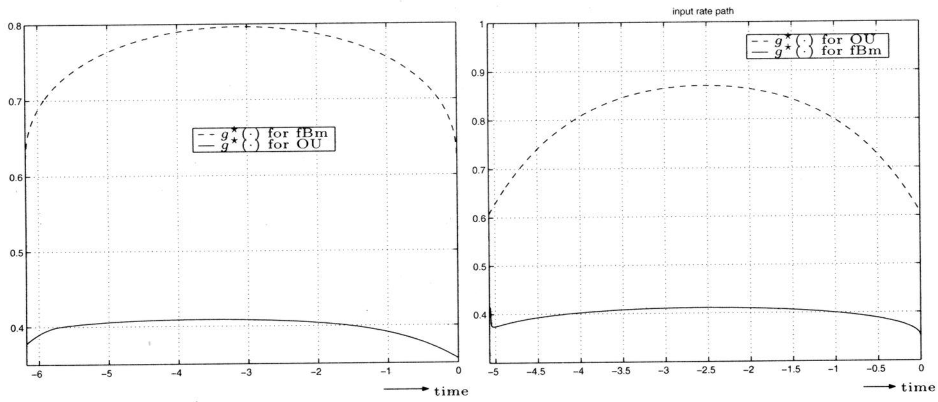

Figure 7.1: Left panel: type 1 corresponds to fBm and type 2 to OU; right panel: type 1 corresponds to OU and type 2 to fBm.

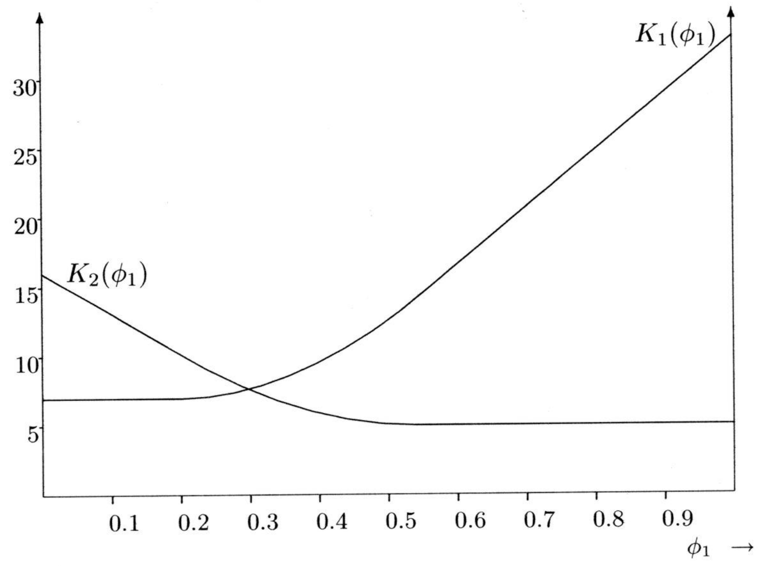

Figure 7.2: The curves $K_i(\phi_1)$ of Example 3.

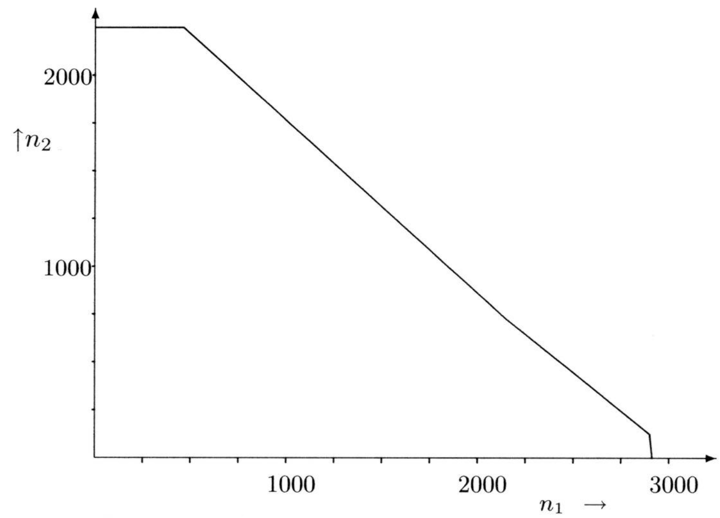

Figure 7.3: Admissible region of Example 4.

References

[1] S. Aalto and W.R.W. Scheinhardt. Tandem fluid queues fed by homogeneous on-off sources. *Operations Research Letters*, 27:73–82, 2000.

[2] J. Abate, G.L. Choudhury, and W. Whitt. Waiting-time tail probabilities in queues with long-tail service-time distributions. *Queueing Systems*, 16:311–338, 1994.

[3] I.J.B.F. Adan and J.A.C. Resing. A two-level traffic shaper for an on-off source. *Performance Evaluation*, 42:279–298, 2000.

[4] R. Addie, P. Mannersalo, and I. Norros. Most probable paths and performance formulae for buffers with Gaussian input traffic. *European Transactions on Telecommunications*, 13:183–196, 2002.

[5] R. Adler. *An Introduction to Continuity, Extrema, and Related Topics for General Gaussian Processes*, volume 12 of *IMS Lecture Notes-Monograph Series*. Institute of Mathematical Statistics, 1990.

[6] R. Agrawal, A.M. Makowski, and P. Nain. On a reduced load equivalence for fluid queues under subexponentiality. *Queueing Systems*, 33:5–42, 1999.

[7] V. Anantharam. How large delays build up in a GI/G/1 queue. *Queueing Systems*, 5:345–368, 1988.

[8] A. Arvidsson and P. Karlsson. On traffic models for TCP/IP. In P. Key and D. Smith, editors, *Proceedings ITC-16*, pages 457–466, Edinburgh, UK, 1999. North-Holland, Amsterdam.

[9] R. Bahadur and S. Zabell. Large deviations of the sample mean in general vector spaces. *Annals of Probability*, 7:587–621, 1979.

[10] J.C.R. Bennett and H. Zhang. Hierarchical packet fair queueing algorithms. *IEEE/ACM Transactions on Networking*, 5:675–689, 1997.

[11] A. Berger and W. Whitt. Effective bandwidths with priorities. *IEEE/ACM Transactions on Networking*, 6:447–460, 1998.

[12] D. Bertsimas, I.C. Paschalidis, and J.N. Tsitsiklis. Large deviations analysis of the Generalized Processor Sharing policy. *Queueing Systems*, 32:319–349, 1999.

[13] N.H. Bingham and R.A. Doney. Asymptotic properties of supercritical branching processes I: The Galton-Watson process. *Advances in Applied Probability*, 6:711–731, 1974.

[14] N.H. Bingham, C.M. Goldie, and J.L. Teugels. *Regular Variation*. Cambridge University Press, Cambridge, 1987.

[15] N.K. Boots. *Rare Event Simulation in Models with Heavy-Tailed Random Variables*. PhD thesis, Tinbergen Institute, Vrije Universiteit Amsterdam, Amsterdam, The Netherlands, 2002.

[16] S.C. Borst, O.J. Boxma, and P.R. Jelenković. Generalized processor sharing with long-tailed traffic sources. In P. Key and D. Smith, editors, *Proceedings ITC-16*, pages 345–354, Edinburgh, UK, 1999. North-Holland, Amsterdam.

[17] S.C. Borst, O.J. Boxma, and P.R. Jelenković. Induced burstiness in generalized processor sharing queues with long-tailed traffic flows. In *Proceedings 37th Annual Allerton Conference on Communication, Control, and Computing*, pages 316–325, Allerton, IL, USA, 1999.

[18] S.C. Borst, O.J. Boxma, and P.R. Jelenković. Asymptotic behavior of generalized processor sharing with long-tailed traffic sources. In *Proceedings Infocom 2000*, pages 912–921, Tel-Aviv, Israel, 2000.

[19] S.C. Borst, O.J. Boxma, and P.R. Jelenković. Coupled processors with regularly varying service times. In *Proceedings Infocom 2000*, pages 157–164, Tel-Aviv, Israel, 2000.

[20] S.C. Borst, O.J. Boxma, and P.R. Jelenković. Reduced-load equivalence and induced burstiness in GPS queues with long-tailed traffic flows. Technical Report PNA-R0016, CWI, The Netherlands, 2000.

[21] S.C. Borst, O.J. Boxma, and P.R. Jelenković. Reduced-load equivalence and induced burstiness in GPS queues with long-tailed traffic flows. *Queueing Systems*, 43:273–306, 2003.

[22] S.C. Borst, O.J. Boxma, and R. Núñez-Queija. Heavy tails: the effect of the service discipline. In T. Field, P.G. Harrison, J. Bradley, and U. Harder, editors, *Computer Performance Evaluation – Modelling Techniques and Tools. Proceedings TOOLS 2002*, pages 1–30, London, UK, 2002. Springer Verlag, Berlin.

[23] S.C. Borst, O.J. Boxma, and M.J.G. van Uitert. Two coupled queues with heterogeneous traffic. In J. Moreira de Souza, N.L.S. da Fonseca, and E.A. de Souza e Silva, editors, *Proceedings ITC-17*, pages 1003–1014, Salvador da Bahia, Brazil, 2001. North-Holland, Amsterdam.

[24] S.C. Borst, O.J. Boxma, and M.J.G. van Uitert. The asymptotic workload behavior of two coupled queues. *Queueing Systems*, 43:81–102, 2003.

[25] S.C. Borst, K. Dębicki, and A.P. Zwart. Subexponential asymptotics of hybrid fluid and ruin models. Technical Report SPOR-Report 2003-10, Eindhoven University of Technology, The Netherlands, 2003.

[26] S.C. Borst, M. Mandjes, and M.J.G. van Uitert. Generalized Processor Sharing with heterogeneous traffic classes. *Performance Evaluation Review*, 29:40–42, 2001. Proceedings MAMA 2001 Workshop, Cambridge, MA, USA.

[27] S.C. Borst, M. Mandjes, and M.J.G. van Uitert. GPS queues with heterogeneous traffic classes. In *Proceedings Infocom 2002*, pages 74–83, New York, NY, USA, 2002.

[28] S.C. Borst, M. Mandjes, and M.J.G. van Uitert. Generalized Processor Sharing queues with heterogeneous traffic classes. *Advances in Applied Probability*, 35:806–845, 2003.

[29] S.C. Borst, M. Mandjes, and M.J.G. van Uitert. Generalized Processor Sharing with light-tailed and heavy-tailed input. To appear in *IEEE/ACM Transactions on Networking*, 11, 2003.

[30] S.C. Borst, R. Núñez-Queija, and M.J.G. van Uitert. User-level performance of elastic traffic in a differentiated-services environment. *Performance Evaluation*, 49:507–519, 2002.

[31] S.C. Borst and A.P. Zwart. A reduced-peak equivalence for queues with a mixture of light-tailed and heavy-tailed input flows. *Advances in Applied Probability*, 35:793–805, 2003.

[32] D. Botvich and N. Duffield. Large deviations, the shape of the loss curve and economies of scale in large multiplexers. *Queueing Systems*, 20:293–320, 1995.

[33] O.J. Boxma, Q. Deng, and A.P. Zwart. Waiting-time asymptotics for the M/G/2 queue with heterogeneous servers. *Queueing Systems*, 40:5–31, 2002.

[34] O.J. Boxma and V. Dumas. The busy period in the fluid queue. *Performance Evaluation Review*, 26:100–110, 1998.

[35] O.J. Boxma and V. Dumas. Fluid queues with heavy-tailed activity period distributions. *Computer Communications*, 21:1509–1529, 1998.

[36] O.J. Boxma and I.A. Kurkova. The M/M/1 queue in a heavy-tailed random environment. *Statistica Neerlandica*, 54:221–236, 2000.

[37] J. Cao and K. Ramanan. A Poisson limit for buffer overflow probabilities. In *Proceedings Infocom 2002*, pages 994–1003, New York, NY, USA, 2002.

[38] G.L. Choudhury, D.M. Lucantoni, and W. Whitt. Squeezing the most out of ATM. *IEEE Transactions on Communications*, 44:203–217, 1996.

[39] J.W. Cohen. Some results on regular variation for distributions in queueing and fluctuation theory. *Journal of Applied Probability*, 10:343–353, 1973.

[40] J.W. Cohen. *The Single Server Queue*. North-Holland Publishing Company, Amsterdam, revised edition, 1982.

[41] J.W. Cohen. On a functional relation in three complex variables; three coupled processors. Technical Report 359, Mathematical Institute, University of Utrecht, The Netherlands, 1984.

[42] J.W. Cohen and O.J. Boxma. *Boundary Value Problems in Queueing System Analysis*. North-Holland, Amsterdam, 1983.

[43] C. Courcoubetis and R. Weber. Buffer overflow asymptotics for a buffer handling many traffic sources. *Journal of Applied Probability*, 33:886–903, 1996.

[44] M.H.A. Davis. Piecewise-deterministic Markov processes: a general class of non-diffusion stochastic models. *Journal of the Royal Statistical Society B*, 46:353–376, 1984.

[45] K. Dębicki and M. Mandjes. Traffic with an fBm limit: convergence of the workload process. Technical Report PNA-R0220, CWI, The Netherlands, 2002.

[46] K. Dębicki and M. Mandjes. Exact overflow asymptotics for queues with many Gaussian inputs. *Journal of Applied Probability*, 40:704–720, 2003.

[47] A. de Meyer and J.L. Teugels. On the asymptotic behaviour of the distributions of the busy period and service time in M/G/1. *Journal of Applied Probability*, 17:802–813, 1980.

[48] G. de Veciana and G. Kesidis. Bandwidth allocation for multiple qualities of service using Generalized Processor Sharing. *IEEE Transactions on Information Theory*, 42:268–272, 1996.

[49] S. Delas, R.R. Mazumdar, and C.P. Rosenberg. Tail asymptotics for HOL priority queues handling a large number of independent stationary sources. *Queueing Systems*, 40:183–204, 2002.

[50] A. Dembo and O. Zeitouni. *Large Deviations Techniques and Applications*. Jones and Bartlett Publishers, Boston, 1993.

[51] A. Demers, S. Keshav, and S. Shenker. Analysis and simulation of a fair queueing algorithm. *Internetworking: Research and Experience*, 1:3–26, 1990.

[52] Q. Deng. *Queues with Regular Variation.* PhD thesis, Eindhoven University of Technology, Eindhoven, The Netherlands, 2001.

[53] J.-D. Deuschel and D. Stroock. *Large Deviations.* Academic Press, London, 1989.

[54] N. Duffield. Queueing at large resources driven by long-tailed $M/G/\infty$-modulated processes. *Queueing Systems*, 28:245–266, 1998.

[55] V. Dumas and A. Simonian. Asymptotic bounds for the fluid queue fed by subexponential on/off sources. *Advances in Applied Probability*, 32:244–255, 2000.

[56] P. Dupuis and K. Ramanan. A Skorokhod problem formulation and large deviation analysis of a processor sharing model. *Queueing Systems*, 28:109–124, 1998.

[57] A. Elwalid, D. Heyman, T.V. Lakshman, D. Mitra, and A. Weiss. Fundamental bounds and approximations for ATM multiplexers wit applications to video teleconferencing. *IEEE Journal on Selected Areas in Communications*, 13:1004–1016, 1995.

[58] A.I. Elwalid and D. Mitra. Analysis and design of rate-based congestion control of high speed networks, I: stochastic fluid models, access regulation. *Queueing Systems*, 9:29–64, 1991.

[59] A.I. Elwalid and D. Mitra. Effective bandwidth of general Markovian traffic sources and admission control of high speed networks. *IEEE/ACM Transactions on Networking*, 1:329–343, 1993.

[60] A.I. Elwalid and D. Mitra. Design of Generalized Processor Sharing schedulers which statistically multiplex heterogeneous QoS classes. In *Proceedings Infocom 1999*, pages 1220–1230, New York, NY, USA, 1999.

[61] G. Fayolle and R. Iasnogorodski. Two coupled processors: the reduction to a Riemann-Hilbert problem. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 47:325–351, 1979.

[62] G. Fayolle, I. Mitrani, and R. Iasnogorodski. Sharing a processor among many job classes. *Journal of the Association for Computing Machinery*, 27:519–532, 1980.

[63] P.W. Glynn and W. Whitt. Logarithmic asymptotics for steady-state tail probabilities in a single-server queue. *Journal of Applied Probability*, 31A:131–159, 1994.

[64] M. Grossglauser and J.-C. Bolot. On the relevance of long-range dependence in network traffic. *IEEE/ACM Transactions on Networking*, 7:629–640, 1999.

[65] F. Guillemin, R.R. Mazumdar, A. Dupuis, and J. Boyer. Analysis of the fluid weighted fair queueing system. *Journal of Applied Probability*, 40:180–199, 2003.

[66] D. Heath, S.I. Resnick, and G. Samorodnitsky. Heavy tails and long range dependence in on-off processes and associated fluid models. *Mathematics of Operations Research*, 23:145–165, 1998.

[67] P.R. Jelenković and A.A. Lazar. Asymptotic results for multiplexing subexponential on-off processes. *Advances in Applied Probability*, 31:394–421, 1999.

[68] P.R. Jelenković and P. Momčilović. Finite buffer queue with Generalized Processor Sharing and heavy-tailed input processes. *Computer Networks*, 40:433–443, 2002.

[69] P.R. Jelenković and P. Momčilović. Asymptotic loss probability in a finite buffer fluid queue with heterogeneous heavy-tailed on-off processes. *Annals of Applied Probability*, 13:576–603, 2003.

[70] P.R. Jelenković, P. Momčilović, and A.P. Zwart. Reduced-load equivalence and subexponentiality. Technical Report RR-4444, INRIA Rocquencourt, France, 2002. Accepted for publication in Queueing Systems.

[71] O. Kella and W. Whitt. A tandem fluid network with Lévy input. In I. Basawa and U. Bhat, editors, *Queues and Related Models*, pages 112–128. Oxford University Press, Oxford, 1992.

[72] G. Kesidis, J. Walrand, and C.-S. Chang. Effective bandwidths for multiclass Markov fluids and other ATM sources. *IEEE/ACM Transactions on Networking*, 1:424–428, 1993.

[73] C. Klüppelberg. Subexponential distributions and integrated tails. *Journal of Applied Probability*, 25:132–141, 1988.

[74] C. Knessl and J.A. Morrison. Heavy traffic analysis of two coupled processors. *Queueing Systems*, 43:173–220, 2003.

[75] A.G. Konheim, I. Meilijson, and A. Melkman. Processor sharing of two parallel lines. *Journal of Applied Probability*, 18:952–956, 1981.

[76] L. Kosten. Liquid models for a type of information buffer problem. *Delft Progress Report*, 11:71–86, 1986.

[77] C. Kotopoulos, N. Likhanov, and R.R. Mazumdar. Asymptotic analysis of GPS systems fed by heterogeneous long-tailed sources. In *Proceedings Infocom 2001*, pages 299–308, Anchorage, AK, USA, 2001.

[78] C. Kotopoulos and R.R. Mazumdar. Buffer occupancy and delay asymptotics in multi-buffered systems with Generalized Processor Sharing handling a large number of independent traffic streams. Preprint, 2002.

[79] C. Kotopoulos and R.R. Mazumdar. Many sources asymptotics for a 2-buffer system with Generalized Processor Sharing. Preprint, 2002.

[80] K. Kumaran, G.E. Margrave, D. Mitra, and K.R. Stanley. Novel techniques for the design and control of Generalized Processor Sharing schedulers for multiple QoS classes. In *Proceedings Infocom 2000*, pages 932–941, Tel Aviv, Israel, 2000.

[81] R.J. La and V. Anantharam. Adaptive modification of Generalized Processor Sharing. Submitted, 2003.

[82] 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*, 2:1–15, 1994.

[83] M. Lelarge. Asymptotic behavior of Generalized Processor Sharing queues under subexponential hypothesis. Technical Report RR-4339, INRIA Rocquencourt, France, 2001.

[84] N. Likhanov and R.R. Mazumdar. Cell loss asymptotics for buffers fed with a large number of independent stationary sources. *Journal of Applied Probability*, 36:86–96, 1999.

[85] M. Mandjes. A note on the benefits of buffering. Submitted, 2003.

[86] M. Mandjes and J.H. Kim. Large deviations for small buffers: an insensitivity result. *Queueing Systems*, 37:349–362, 2001.

[87] M. Mandjes, P. Mannersalo, I. Norros, and M.J.G. van Uitert. Most probable busy period paths in Gaussian queues. Temporary document, COST-279, 2003.

[88] M. Mandjes and A. Ridder. Finding the conjugate of Markov fluid processes. *Probability in the Engineering and Informational Sciences*, 9:297–315, 1995.

[89] M. Mandjes, I. Saniee, and A. Stolyar. Load characterization, overload prediction, and load anomaly detection for voice over IP traffic. In *Proceedings 38th Annual Allerton Conference on Communication, Control, and Computing*, pages 567–576, Allerton, IL, USA, 2000.

[90] M. Mandjes and M.J.G. van Uitert. Sample-path large deviations for tandem and priority queues with Gaussian inputs. Technical Report PNA-R0221, CWI, The Netherlands, 2002.

[91] M. Mandjes and M.J.G. van Uitert. Large deviations for complex buffer architectures: the short-range dependent case. Work in progress, 2003.

[92] M. Mandjes and M.J.G. van Uitert. Sample-path large deviations for Generalized Processor Sharing queues with Gaussian inputs. Technical Report PNA-R0308, CWI, The Netherlands, 2003.

[93] M. Mandjes and M.J.G. van Uitert. Sample-path large deviations for tandem queues with Gaussian inputs. In J. Charzinski, R. Lehnert, and P. Tran-Gia, editors, *Proceedings ITC-18*, pages 521–530, Berlin, Germany, 2003. Elsevier, Amsterdam.

[94] P. Mannersalo. *Gaussian and Multifractal Processes in Teletraffic Theory.* PhD thesis, Helsinki University of Technology, Finland, 2003.

[95] P. Mannersalo and I. Norros. Approximate formulae for Gaussian priority queues. In J. Moreira de Souza, N.L.S. da Fonseca, and E.A. de Souza e Silva, editors, *Proceedings ITC-17*, pages 991–1002, Salvador da Bahia, Brazil, 2001. North-Holland, Amsterdam.

[96] P. Mannersalo and I. Norros. GPS schedulers and Gaussian traffic. In *Proceedings Infocom 2002*, pages 1660–1667, New York, NY, USA, 2002.

[97] P. Mannersalo and I. Norros. A most probable path approach to queueing systems with general Gaussian input. *Computer Networks*, 40:399–412, 2002.

[98] L. Massoulié. Large deviations estimates for polling and Weighted Fair Queueing service systems. *Advances in Performance Analysis*, 2:103–127, 1999.

[99] T. Mikosch, S. Resnick, H. Rootzén, and A. Stegeman. Is network traffic approximated by stable Lévy motion or fractional Brownian motion? *Annals of Applied Probability*, 12:23–68, 2002.

[100] I. Norros. A storage model with self-similar input. *Queueing Systems*, 16:387–396, 1994.

[101] I. Norros. Busy periods of fractional Brownian storage: a large deviations approach. *Advances in Performance Analysis*, 2:1–19, 1999.

[102] I. Norros. Most probable paths in Gaussian priority queues. Technical Report 16, COST-257, 1999.

[103] I. Norros. Most probable path techniques for Gaussian queueing systems. In *2nd International IFIP-TC6 Networking Conference*, pages 86–104, Pisa, Italy, 2002. Springer Verlag, Berlin.

[104] I. Norros and J. Kilpi. Gaussian traffic modelling for differentiated services. In *Proceedings 15th Nordic Teletraffic Seminar*, Lund, Sweden, 2000.

[105] R. Núñez-Queija. *Processor-Sharing Models for Integrated-Services Networks.* PhD thesis, Eindhoven University of Technology, Eindhoven, The Netherlands, 2000.

[106] A.G. Pakes. On the tails of the waiting-time distributions. *Journal of Applied Probability*, 12:555–564, 1975.

[107] A.K. Parekh and R.G. Gallager. A Generalized Processor Sharing approach to flow control in integrated services networks: the single-node case. *IEEE/ACM Transactions on Networking*, 1:344–357, 1993.

[108] 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*, 2:137–150, 1994.

[109] K. Park and W. Willinger. Self-similar network traffic: an overview. In K. Park and W. Willinger, editors, *Self-Similar Network Traffic and Performance Evaluation*, pages 1–38. Wiley, New York, 2000.

[110] M. Parulekar and A.M. Makowski. $M/G/\infty$ input processes: A versatile class of models for network traffic. In *Proceedings Infocom 1997*, pages 419–426, Kobe, Japan, 1997.

[111] I.C. Paschalidis. Class-specific quality of service guarantees in multimedia communication networks. *Automatica*, 35:1951–1968, 1999.

[112] A. Paxson and S. Floyd. Wide area traffic: the failure of Poisson modeling. *IEEE/ACM Transactions on Networking*, 3:226–244, 1995.

[113] F.M. Pereira, N.L.S. Fonseca, and D.S. Arantes. On the performance of Generalized Processor Sharing servers under long-range dependent traffic. *Computer Networks*, 40:413–431, 2002.

[114] K. Ramanan and P. Dupuis. Large deviation properties of data streams that share a buffer. *Annals of Applied Probability*, 8:1070–1129, 1998.

[115] K.M. Rege and B. Sengupta. A decomposition theorem and related results for the discriminatory processor sharing queue. *Queueing Systems*, 18:333–351, 1994.

[116] K.M. Rege and B. Sengupta. Queue-length distribution for the discriminatory processor-sharing queue. *Operations Research*, 44:653–657, 1996.

[117] E. Reich. On the integrodifferential equation of Takács I. *Annals of Mathematical Statistics*, 29:563–570, 1958.

[118] S. Resnick and G. Samorodnitsky. Activity periods of an infinite server queue and performance of certain heavy tailed fluid queues. *Queueing Systems*, 33:43–71, 1999.

[119] J. Roberts, U. Mocci, and J. Virtamo. *Broadband network teletraffic*. Springer, Berlin, 1996. Final report of action COST 242.

[120] I. Rubin. Message path delays in packet-switching networks. *IEEE Transactions on Communications*, 23:186–192, 1975.

[121] B. Ryu and A.I. Elwalid. The importance of long-range dependence of VBR video traffic in ATM traffic engineering: myths and realities. *Computer Communication Review*, 26:3–14, 1996.

[122] G. Samorodnitsky and M.S. Taqqu. *Stable Non-Gaussian Random Processes*. Chapman & Hall, New York, 1994.

[123] W.R.W. Scheinhardt. *Markov Modulated and Feedback Fluid Queues*. PhD thesis, University of Twente, Enschede, The Netherlands, 1998.

[124] S. Shakkottai and R. Srikant. Many-sources delay asymptotics with applications to priority queues. *Queueing Systems*, 39:183–200, 2001.

[125] M. Shalmon and M.A. Kaplan. A tandem network of queues with deterministic service and intermediate arrivals. *Operations Research*, 32:753–773, 1984.

[126] A. Shwartz and A. Weiss. *Large Deviations for Performance Analysis*. Chapman & Hall, London, 1995.

[127] K. Sigman. A primer on heavy-tailed distributions. *Queueing Systems*, 33:261–275, 1999.

[128] A. Stegeman. Heavy tails versus long-range dependence in self-similar network traffic. *Statistica Neerlandica*, 54:293–314, 2000.

[129] D. Stiliadis. *Traffic Scheduling in Packet Switched Networks: Analysis, Design and Implementation*. PhD thesis, University of California, Santa Cruz, California, USA, 1996.

[130] W.G.L. Sutton. The asymptotic expansion of a function whose operational equivalent is known. *Journal of the London Mathematical Society*, 9:131–137, 1934.

[131] M.S. Taqqu, W. Willinger, and R. Sherman. Proof of a fundamental result in self-similar traffic modeling. *Computer Communication Review*, 27:5–23, 1997.

[132] F. Toutain. Decoupled Generalized Processor Sharing: a fair queueing principle for adaptive multimedia applications. In *Proceedings Infocom 1998*, pages 291–298, San Francisco, California, USA, 1998.

[133] M.J.G. van Uitert and S.C. Borst. Generalised Processor Sharing networks fed by heavy-tailed traffic flows. In *Proceedings Infocom 2001*, pages 269–278, Anchorage, AK, USA, 2001.

[134] M.J.G. van Uitert and S.C. Borst. A reduced-load equivalence for Generalised Processor Sharing networks with long-tailed input flows. *Queueing Systems*, 41:123–163, 2002.

[135] A. Weiss. An introduction to large deviations for communication networks. *IEEE Journal on Selected Areas in Communications*, 13:938–952, 1995.

[136] W. Whitt. *Stochastic-Process Limits*. Springer Verlag, New York, 2002.

[137] E. Willekens and J.L. Teugels. Asymptotic expansions for waiting time probabilities in an M/G/1 queue with long-tailed service time. *Queueing Systems*, 10:295–311, 1992.

[138] W. Willinger, M.S. Taqqu, R. Sherman, and D.V. Wilson. Self-similarity through high-variability: statistical analysis of Ethernet LAN traffic at the source level. *IEEE/ACM Transactions on Networking*, 5:71–86, 1997.

[139] D. Wischik. Sample path large deviations for queues with many inputs. *Annals of Applied Probability*, 11:379–404, 2001.

[140] O. Yaron and M. Sidi. Performance and stability of communication networks via robust exponential bounds. *IEEE/ACM Transactions on Networking*, 1:372–385, 1993.

[141] O. Yaron and M. Sidi. Generalized processor sharing networks with exponentially bounded burstiness arrivals. In *Proceedings Infocom 1994*, pages 628–634, Toronto, Canada, 1994.

[142] O. Yaron and M. Sidi. Generalized Processor Sharing networks with exponentially bounded burstiness arrivals. *Journal of High Speed Networks*, 3:375–387, 1994.

[143] S.F. Yashkov. Processor-sharing queues: Some progress in analysis. *Queueing Systems*, 2:1–17, 1987.

[144] S.F. Yashkov. Mathematical problems in the theory of processor-sharing queueing systems. *Journal of Soviet Mathematics*, 58:101–147, 1992.

[145] Z.-L. Zhang. *End-to-End Support for Statistical Quality-of-Service Guarantees in Multimedia Networks*. PhD thesis, University of Massachusetts Amherst, MA, USA, 1997.

[146] Z.-L. Zhang. Large deviations and the Generalized Processor Sharing scheduling for a two-queue system. *Queueing Systems*, 26:229–245, 1997.

[147] Z.-L. Zhang. Large deviations and the Generalized Processor Sharing scheduling for a multiple-queue system. *Queueing Systems*, 28:349–376, 1998.

[148] Z.-L. Zhang, Z. Liu, J. Kurose, and D. Towsley. Call admission control schemes under Generalized Processor Sharing scheduling. *Telecommunication Systems*, 7:125–152, 1997.

[149] Z.-L. Zhang, D. Towsley, and J. Kurose. Statistical analysis of the Generalized Processor Sharing scheduling discipline. *IEEE Journal of Selected Areas in Communications*, 13:1071–1080, 1995.

[150] A.P. Zwart. *Queueing Systems with Heavy Tails.* PhD thesis, Eindhoven University of Technology, Eindhoven, The Netherlands, 2001.

[151] A.P. Zwart. Tail asymptotics for the busy period in the GI/G/1 queue. *Mathematics of Operations Research*, 26:485–493, 2001.

[152] A.P. Zwart, S.C. Borst, and M. Mandjes. Exact asymptotics for fluid queues fed by multiple heavy-tailed on-off flows. To appear in *Annals of Applied Probability*, 13, 2003.

Samenvatting (Summary)

In dit proefschrift staat de prestatie-analyse van Generalized Processor Sharing (GPS) centraal. GPS ligt ten grondslag aan mechanismes die capaciteit verdelen tussen gebruikers in communicatienetwerken. Daartoe worden gebruikers ingedeeld in klassen en wordt aan elke klasse een gewicht toegewezen. Dit gewicht geeft de fractie van de capaciteit aan, welke gegarandeerd beschikbaar is voor de betreffende klasse. Wanneer voor sommige klassen meer capaciteit beschikbaar is dan nodig, wordt het teveel aan capaciteit verdeeld over de andere klassen volgens dezelfde gewichten.

Naast het bestuderen van modellen met GPS (Hoofdstuk 3, 4 en 7), bekijken we in dit proefschrift ook andere modellen die ofwel nauw verwant zijn aan GPS (Hoofdstuk 5), ofwel inzicht verschaffen in de wijze waarop de analyse van modellen met GPS kan worden voortgezet (Hoofdstuk 6).

We richten ons in dit proefschrift op het afleiden van het staartgedrag van de kans dat de bufferinhoud van een bepaalde klasse een zekere waarde overschrijdt. Hierbij bekijken we twee asymptotische regimes. In Hoofdstuk 3, 4 en 5 laten we de drempelwaarde van de bufferinhoud groot worden en bepalen het staartgedrag van de kansverdeling met behulp van de *sample-path* methode (Hoofdstuk 3 en 4) of met behulp van complexe functietheorie (Hoofdstuk 5). Vervolgens laten we in Hoofdstuk 6 en 7 het aantal gebruikers binnen een klasse groot worden en bepalen het staartgedrag van de kansverdeling met *sample-path large-deviations* theorie. Hoofdstuk 1 en 2 geven informatie over de praktische achtergrond en beschrijven de benodigde wiskundige theorie.

In Hoofdstuk 3 analyseren we een netwerk waarbij ieder knooppunt opereert volgens GPS. We zijn geïnteresseerd in de staartverdeling van de bufferinhoud van een klasse die alleen uit bron i bestaat. Bron i genereert verkeer volgens een aan-/uitproces met zwaarstaartig verdeelde aantijden en exponentieel verdeelde uittijden, of genereert met exponentieel verdeelde tussentijden instantaan een hoeveelheid verkeer met een zwaarstaartige verdeling. We maken de volgende aannames. Ten eerste, de capaciteit die de klassen gegarandeerd krijgen, is groter dan de gemiddelde intensiteit waarmee ze verkeer genereren. Ten tweede, de capaciteit die beschikbaar is voor bron i, na aftrek van de gemiddelde intensiteit waarmee de andere bronnen sturen, wordt steeds kleiner bij elk volgend knooppunt dat bron i doorkruist. Onder deze aannames bepalen we de staartverdeling van de bufferinhoud bij het laatste knooppunt dat bron i aandoet. In het bijzonder laten we zien dat deze (asymptotisch) alleen afhangt van de gemiddelde intenstiteit waarop de andere bronnen sturen.

In Hoofdstuk 4 bekijken we een model met een knooppunt dat wederom opereert volgens GPS. Er zijn twee klassen van bronnen die verkeer genereren. De bronnen in klasse 1 genereren verkeer volgens een *Markov-modulated fluid* proces, dat lichtstaartig is. De bronnen in klasse 2 gedragen zich zoals bron i in Hoofdstuk 3. We maken de aanname dat de gewichten van beide klassen, ϕ_1 en ϕ_2, groter zijn dan de gemiddelde intensiteit waarmee verkeer gestuurd wordt. We bewijzen dat de kans dat de bufferinhoud van klasse 1 een grote waarde overschrijdt zich asymptotisch gedraagt als een produkt van twee kansen, waarbij de één exponentieel, en de ander polynomiaal daalt in het bufferniveau. De dominante, exponentieel dalende, term is de kans dat de bufferinhoud van klasse 1 een grote waarde overschrijdt in een geïsoleerd systeem met capaciteit ϕ_1. De andere, polynomiaal dalende, term is de kans dat klasse 2 gedurende een voldoende lange periode volledig beslag legt op ϕ_2.

In Hoofdstuk 5 beschouwen we een model dat nauw verwant is aan het model van Hoofdstuk 4, het gekoppelde-processorenmodel. De capaciteit wordt in dit model als volgt verdeeld tussen de twee klassen. Als beide klassen werk hebben krijgen ze capaciteit 1, maar als één van beide klassen geen capaciteit nodig heeft gaat de capaciteit van de andere klasse omhoog. De resultaten zijn kwalitatief gezien hetzelfde als die verkregen in Hoofdstuk 4, maar de gehanteerde methode is van een essentieel ander karakter, zoals eerder aangegeven.

In Hoofdstuk 6 richten we onze aandacht op twee andere modellen, een tandem netwerk met twee knooppunten, capaciteit c_1 en c_2, en een prioriteitssysteem met twee klassen. Deze modellen blijken nauw verwant te zijn aan elkaar, en vormen een opstap voor de analyse in Hoofdstuk 7. Voor het tandem netwerk nemen we aan dat een groot aantal bronnen verkeer genereert volgens een Gaussisch proces. We bestuderen het staartgedrag (logaritmisch) van de kans dat de bufferinhoud bij het tweede knooppunt een bepaalde waarde overschrijdt. In het bijzonder laten we zien dat (asymptotisch) het eerste knooppunt geen rol speelt als c_1 groter is dan een kritische waarde. We verkrijgen ook het meest waarschijnlijke scenario om de betreffende waarde van de bufferinhoud te overschrijden. In het prioriteitssysteem leiden we soortgelijke resultaten af voor de klasse met de lage prioriteit, onder de aanname dat beide klassen bestaan uit een groot aantal Gaussische bronnen.

In Hoofdstuk 7 bekijken we een GPS model met twee klassen die bestaan uit een groot aantal Gaussische bronnen. Met behulp van dezelfde methode als in Hoofdstuk 6, leiden we expliciete onder- en bovengrenzen af voor het staartgedrag (logaritmisch) van de kans dat de bufferinhoud van één van de klassen een bepaalde waarde overschrijdt. In het geval dat de bronnen zich gedragen als Brownse bewegingen laten we zien dat de grenzen met elkaar overeenkomen. Met behulp van numerieke resultaten en heuristische argumenten beargumenteren we dat de gevonden grenzen ook voor andersoortige Gaussische bronnen aan elkaar gelijk zijn. In het bijzonder komen we tot een onderverdeling van de resultaten in drie regimes, afhankelijk van de waarde van de gewichten. Tevens beschrijven we twee methoden om tot een geschikte keuze van de gewichten te komen.

About the author / Over de auteur

Miranda (Maria Johanna Gerarda) van Uitert was born in Kerkdriel, The Netherlands, on August 13, 1976. She completed Grammar School at the Jeroen Boschcollege, Den Bosch, in June 1994. On April 16, 1999, she received her master's degree in econometrics (cum laude) from the Katholieke Universiteit Brabant (currently Universiteit van Tilburg), Tilburg. She also participated in the Graduate Program in Management at CentER (Universiteit van Tilburg). She started as a PhD student at CWI (Center for Mathematics and Computer Science), Amsterdam, in May 1999. During her PhD project she worked for two years part-time at KPN Research (currently TNO Telecom), Leidschendam. Since September 1, 2003, she has a teaching position at the Vrije Universiteit, Amsterdam. At the same time she started a study in Jungian psychology. She defends this PhD thesis at the Technische Universiteit Eindhoven on November 24, 2003.

Miranda (Maria Johanna Gerarda) van Uitert werd geboren op 13 augustus 1976 in Kerkdriel. In juni 1994 behaalde zij haar gymnasium diploma aan het Jeroen Boschcollege te Den Bosch, waarna zij econometrie ging studeren aan de Katholieke Universiteit Brabant te Tilburg (tegenwoordig Universiteit van Tilburg). Tijdens haar studie nam zij ook deel aan het Graduate Program in Management bij het CentER. Op 16 april 1999 behaalde zij het doctoraal diploma Econometrie (cum laude). Hierna werd zij onderzoeker-in-opleiding bij het CWI (Centrum voor Wiskunde en Informatica) te Amsterdam. Daarnaast werkte zij gedurende twee jaar part-time bij KPN Research (tegenwoordig TNO Telecom) te Leidschendam. Sinds 1 september 2003 is zij werkzaam als docent aan de Vrije Universiteit te Amsterdam. Tegelijkertijd is zij begonnen met een studie Jungiaanse psychologie. Op 24 november 2003 verdedigt zij dit proefschrift aan de Technische Universiteit Eindhoven.