Stability and Approximation of Queueing Networks

Antonios Dimakis

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2006-131
http://www.eecs.berkeley.edu/Pubs/TechRpts/2006/EECS-2006-131.html

October 16, 2006

Stability and Approximation of Queueing Networks

by

Antonios Dimakis

B.S. (University of Crete) 1996
M.S. (University of Crete) 1999

A dissertation submitted in partial satisfaction
of the requirements for the degree of

Doctor of Philosophy
in

Electrical Engineeing and Computer Sciences

in the

GRADUATE DIVISION

of the

UNIVERSITY OF CALIFORNIA, BERKELEY

Committee in charge:

Professor Jean Walrand, Chair
Professor Pravin Varaiya
Professor Jim Pitman

Fall 2006

The dissertation of Antonios Dimakis is approved.

Chair Date

Date

Date

University of California, Berkeley

Fall 2006

Stability and Approximation of Queueing Networks

Abstract

Stability and Approximation of Queueing Networks

by

Antonios Dimakis

Doctor of Philosophy in Electrical Engineeing and Computer Sciences

University of California, Berkeley

Professor Jean Walrand, Chair

Motivated by applications in communication networks, we consider three different models of resource sharing.

We first study the stability of *Longest Queue First* (LQF), a greedy and low complexity scheduling policy, encountered in packet switching and wireless networks. Contrary to other common policies, the stability of LQF depends on the variance of the arrival processes in addition to their intensities. We identify new sufficient conditions for LQF to be throughput optimal for i.i.d. arrival processes. Deterministic fluid analogs, proved to be powerful in the analysis of stability in queueing networks, do not adequately characterize it in this case. We combine properties of the sample paths over different time-scales to obtain a sharper characterization.

The second part is motivated by the presence of variations in channel quality and traffic demand. For a stable Markov chain whose transition matrix is controlled by a general stationary jump process, we provide a framework for approximating the stationary distribution, similar to the constant and linear terms obtained by singular perturbation analysis (SPA) However, SPA is only applicable to Markovian controlling processes. To extend to the non-Markovian case, we follow a non-algebraic approach based on light-traffic

analysis of stationary point processes. Jackson networks with slowly varying service and arrival rates and fixed routing are dealt with in this framework.

In the third part, we consider the congestion arising from coexistence of data and voice flows in a single internet link. Voice flows are rejected and leave the system if on arrival, the total number of flows is above a certain threshold. We obtain deterministic equations describing the population dynamics for each type of flow. These equations arise as the limit of a large stochastic system. Next, we study the behavior near the voice cut-off threshold. While there is no "hard" boundary as in a purely loss system, admissions and rejections of voice flows result in a similar averaging along a "soft" boundary, where temporary overloads are possible during transient periods.

Professor Jean Walrand
Dissertation Committee Chair

To Aspasia and Spyridakena.

Contents

List of Figures

Acknowledgements

This thesis would not have been possible without the constant support, encourangement, friendship and technical advice of several people. Foremost, I would like to thank my advisor Jean Walrand for his patience, tireless assistance and friendship. I will always cherish the moments we worked together. I tremendously enjoyed discussions with Pravin Varaiya who has been a constant source of inspiration throughout my stay in Berkeley. I also thank Jim Pitman for teaching me probability. I am grateful to the colleagues and friends Linhai He, Rahul Jain, Teresa Tung, Wilson So, Assane Gueye and the members of the smartnets group for their support in both technical and nontechnical matters. Thanks also goes to Ruth Gjerde for bailing me out on many occasions. Last but not least I thank Niki Kourmouli, Ekai Txapartegi, Kehinde Ajao, Stratos Maganas for their friendship and my family for bearing me all these years.

Chapter 1

Introduction

This thesis studies three different models of resource sharing that are motivated by applications in communication networks. The advent of wireless networks and the penetration of the internet, have created new modes of congestion and therefore new problems for the network engineer. Here we consider wireless channel interference, distributed scheduling, channel fading, rate adaptation, and admission control.

Classical queueing models do not account for these new modes, so new models and problems of resource sharing continue to drive intense research in engineering and applied probability. As an example of the inapplicability of old thumb-rules from classical models, we refer to the scheduling problem we study in chapter 2 where average arrival rates are not adequate descriptors of congestion anymore. Ultimately this thesis is a collection of models and methods that we have found useful in tackling problems of congestion; some of these tools are new, others are extensions of old ones.

Section 1.1 studies the stability of a particular scheduling policy encountered in many application contexts and whose stability has remained an open problem. The problem we consider in section 1.2 is more general in flavour, and concerns approximation to the stationary distributions of perturbed Markov chains. In the last section of the Introduction, we consider a simple but informative model of bandwidth sharing the internet. Sections 1.1,1.2, and 1.3, respectively form the introductions for the three chapters to follow.

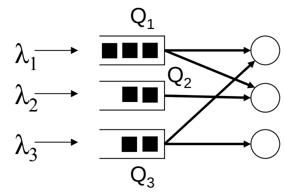

Figure 1.1. Three queues that share resources. A service of each job in a queue requires the exclusive use of the corresponding resources at the right-hand side.

1.1 Stability of Longest Queue First

Consider the queueing system in Figure 1.1. At each discrete time-slot, jobs arrive with respective rates λ_1, λ_2, and λ_3 at queue 1,2, and 3. A job at the head of a queue requires service simultaneously from the resources on the right-hand side associated with that queue. A resource is used by at most one job at a time. Hence, the evolution of queue sizes Q_1, Q_2, Q_3 in time is determined by the choices of the queues served in each time-slot, i.e., the schedule. If the schedule is not chosen wisely, it is possible that the queue lengths increase arbitrarily and the system becomes unstable. Thus, desirable scheduling policies yield schedules that make the system stable for the largest possible set of arrival rates. A policy that achieves the largest such set is called *throughput optimal*. (Of course, this is not the only performance objective that can be considered; e.g. fairness and simplicity in implementation are other important design factors.) In chapter 2 we consider the stability of *Longest Queue First*. This is a greedy policy that gives priority to longer queues. So if the queue lengths are as given in Figure 1.1 then queue 1 will be served. Systems with service constraints such as the one of Figure 1.1 are instances of a much more general model, called the *generalized switch* and described in [39], which includes as special cases models of multiuser data scheduling over a wireless medium, input-queued crossbar switches, and parallel server systems. The problem of throughput optimal scheduling is addressed in [39] where *MaxWeight* scheduling is proposed. MaxWeight serves the subset of queues such that the sum of their lengths is maximum; e.g. in the case of Figure 1.1, queues 2 and 3 will be served, since $Q_2 + Q_3 > Q_1$. MaxWeight is based on the optimal policy of Tassiulas [40], originally proposed in the context of data scheduling in a multihop wireless network. The optimality of MaxWeight, is achieved at the cost of increased complexity. The evaluation of the sums of queue lengths of all possible service combinations is a much more complex operation than the single sorting required in LQF. This is why LQF was first proposed as a lower complexity approximation to MaxWeight by McKeown in [33]. Another reason that makes LQF attractive, is that it lends itself to distributed implementation [22, 23].

The achievable throughput of common scheduling policies in open queueing networks, e.g., first-in first-out, head-of-line processor sharing, or MaxWeight in a generalized switch, is characterized by average arrival intensities, to the effect that the problem of stability is related to that of a deterministic fluid analog. The connection between the stochastic and the fluid system, based on the strong law of large numbers, has been formalized in [30, 14, 38] and applied with great success in multiclass queueing networks [14] and generalized switch models [3, 39, 37, 15]. The premise is that in many cases it is easier to work with a deterministic system than with a stochastic one.

In LQF, on the other hand, variability in stochastic arrivals may affect stability, as will become evident in sections 2.1 and 2.3. In this case, a deterministic analog is not detailed enough to infer stability since deterministic and nondeterministic arrivals with the same intensities can lead to unstable and stable systems, respectively. For this reason, the stability of LQF has remained an open problem. Previous work on LQF [28] considered sufficient conditions for stability by analyzing the associated deterministic systems. This led to stronger than necessary conditions on the average rate of arrivals for certain classes of systems. Moreover, [28] dealt exclusively with constraints arising from cyclic graphs (see section 2.1). To deal with more general graphs, we introduce the concept of *local pooling*, a graph property that is characterized in terms of a linear program. The main result in this part is Theorem 2.3.1 which states that in certain special cases, namely systems satisfying local pooling and a rank condition (see Theorem 2.3.1), LQF is throughput optimal under the assumption of nondeterministic arrivals. Although the systems satisfying the conditions in Theorem 2.3.1 are special, we believe that the techniques used in its proof are interesting on their own right. On the other hand, local pooling is satisfied for many graphs arising in wireless networks, see [11]. A novel feature of our analysis is that we combine combine prop-

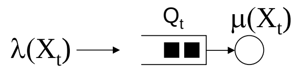

Figure 1.2. A M/M/1 queue with arrival and service rates controlled by process X_t.

erties of the sample paths over different time-scales to obtain a sharper characterization of stability than that obtained by fluid models. Hence we retain the simplicity of dealing with fluid limits but without giving away the key stochastic effects on which stability depends.

1.2 Approximating queues in slowly varying stationary environments

As an example, consider the single queue in Figure 1.2. At time t, arrivals are Poisson with intensity $\lambda(X_t)$, while the service rate is $\mu(X_t)$. Here, $(X_t \in \{0,1\}, t \in \mathbb{R}_+)$ is a stationary piecewise constant process, and we have $P(X_t = 0) = P(X_t = 1) = 1/2$. Assume also that $\rho(x) := \lambda(x)/\mu(x) < 1$, $x = 0, 1$. If $X_t = x$ is constant, we have a normal M/M/1 queue and the invariant distribution is $P(Q = i) = \rho(x)(1 - \rho(x))^i$. In chapter 3 we approximate the stationary distribution when X_t varies slowly. As the rate of jumps, ϵ, of X_t approaches 0, we expect the timescales of X_t and Q_t to separate, giving $P_\epsilon(Q = i) \to .5\rho(0)(1-\rho(0))^i + .5\rho(1)(1-\rho(1))^i =: \pi(i)$. We will show that $P_\epsilon(Q = i) = \pi(i) + \epsilon\phi(i) + o(\epsilon)$ and determine $\phi(\cdot)$.

More generally, we consider approximations to the stationary distribution of a Markov chain whose transition matrix is controlled by a piecewise constant process, henceforth called the environment, only assumed to be stationary. The problem is motivated by the presence of variations in channel quality and traffic demand in communication networks. In general, the computation of this distribution is not tractable and one resorts to simulation. Even in the case of a Markovian environment, where simulation is not necessary, the solution of a large number of equations can be undesirable. (Note that certain simple cases can be solved very efficiently by the matrix-geometric method of Neuts [35].) When environment transitions are rare, Taylor-like expansions around the time-scale separation case are possible under Markovian assumptions. Singular perturbation analysis (see [2, 41] and references therein) is one such approach, where tools there are algebraic. Also, analytic methods have been applied, as in [21]. These do not apply in the non-Markovian case.

Here, we propose a framework inspired by light-traffic approximations [36, 4, 10] and weak convergence instead, in computing the constant and linear term. Key to applying such tools are tightness and uniform convergence properties, which are established by exploiting monotonicity inherent in specific models, and queueing networks in particular. We use monotonicity of Jackson networks, and show that slow variations of service and arrival rates can be dealt within this framework. Since monotonicity properties of queues are well-studied (e.g. see [5]), we expect the approach to carry over in other systems.

As is commonly the case in light-traffic approximations, the first derivative is determined by what happens around a single jump. Hence, the approximations we obtain do not depend on more detailed statistics of the environment other that its transitions. Thus, the approximations are the same as the ones given by a Markov model. It is conceivable that higher-order terms in such expansion can be obtained using the factorial moment expansions of [10]. Also note that, as we do not rely on uniformization, transition rates are not required to be bounded, as assumed in singular perturbation methods (e.g. see [2]).

The main results of this part are Theorems 3.1.1 and 3.2.1. Theorem 3.1.1 in section 3.1 provides an approximation framework for general Markov chains. In section 3.2 we show that Jackson networks belong to this framework (Theorem 3.2.1).

As a side result, we establish existence of a stationary law for varying Jackson networks. In particular, stability is demonstrated under the condition that at every state of the environment, the system is stable; e.g. in terms of the example above, when $\rho(x) < 1, x = 0, 1$. For single queues, weaker stability conditions are known to exist (see [8]). More specifically, it is sufficient that the *average* arrival rate does not exceed the *average* service rate, where the averages are taken over all states of the environment. However, in dealing with approximations when the time-scales of X_t and Q_t separate, we need stability under *any* environment state to avoid trivial cases. This is because the existence of even a single unstable state, will yield $P_\epsilon(Q = i) \to 0$.

1.3 Macroscopic behavior of a lossy resource with real- and non-real time flows

Recent work has focused on the interplay between real-time and non-real-time flows when both they share common resources [16, 32, 13, 31]. Such sharing occurs in data networks, as in the internet, where real-time (e.g. streaming audio/video) and non-real-time (e.g. file transfer, web browsing) services are carried by the same network resources.

Along with the rise in demand of real-time services, rate adaptation mechanisms have been introduced that scale the bandwidth demands of traffic sources in response to congestion signals sent by the network. These mechanisms vary in their level of sophistication, ranging from window-based additive increase/multiplicative decrease, as in TCP and its variants, to highly specialized algorithms that depend on the specific information encoding schemes. We assume that flows adapt their rates in an *ideal* way, such that they converge instantaneously after a new flow arrives or departs.

While real-time and non-real-time applications share the same resources and both employ some form of rate adaptation, their resource usage pattern is fundamentally different in two ways: Firstly, real-time applications use (or intend using) a resource for a period of time that is independent of the other flows. Contrary to this, non-real-time applications (e.g. file transfers) do not leave the system until a certain *amount* of service has been accumulated (e.g. the transfer of an entire file). Secondly, real-time applications are sensitive to their instantaneous bandwidth share since it may cause packets losses, and delays. Thus they have some sort of minimum (instantaneous) bandwidth requirement. User-enforced admission control has been proposed as a mechanism for signalling that this minimum bandwidth share will likely exist [26]. Although such controls have the advantage that they do not need an overlay signaling network, not experiencing congestion after admission of an incoming flow cannot be guaranteed.

An important question is whether this form of control can effectively control overloads and, if not, how do overloads occur. We are also interested in the average time that non-real-time flows spend in the system. For the case of file transfers, this corresponds to file transfer delay. In this part we propose a simple model of congestion at the flow level, and answer these questions at the level of detail captured by the model. We obtain deterministic equations describing the population dynamics for each type of flow. These equations arise as the limit of a large stochastic system, and allow for transient analysis which is difficult to obtain for the stochastic system itself. We note here that a more recent result of Massoulié et al. [31] uses the same deterministic model to tackle the question of stability for the case of a network, under different bandwidth sharing schemes.

In the loss network literature work has focused on identifying macroscopic limits under different control schemes, e.g. see [24, 9, 1, 42]. The limiting regime considered is that where both flow arrival rates and capacity grow large in a fixed proportion, as described by Kelly in [25]. Hunt and Kurtz in [24] establish a weak convergence result for a general loss network and identify the behavior at the boundary of full capacity. At the limit, the system evolves according to a set of deterministic differential equations which arise as a result of a time-scale separation between the number of free circuits and the number of flows currently in the system (normalized with respect to the system size). Although in a purely loss system the boundary restricts the system in a set of acceptable states, i.e. it is a "hard" boundary, admissions and rejections of voice flows result in a similar averaging along a "soft" boundary, where temporary overloads are possible during transient periods. This will permit us, in Section 4.3, to derive a weak convergence result for a sequence of systems that grow large in the regime described above. This is the content of Theorem 4.3.2 and Proposition 4.3.3 which are the main results of chapter 4.

In summary, congestion occurs in three ways.

1. Long-term congestion where, for a period of time of comparable magnitude with the

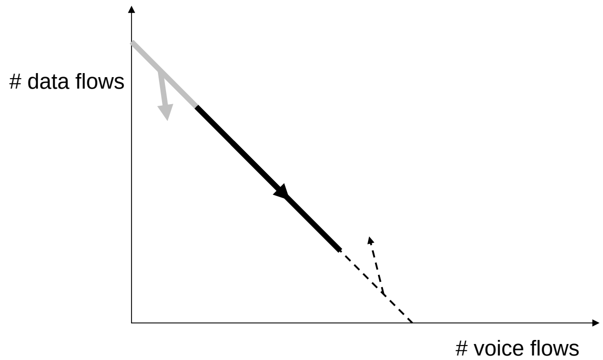

Figure 1.3. Possible trajectories on the voice flow admission boundary.

duration of real-time flows, the bandwidth share for the latter is below the required minimum amount.

2. Short-term congestion, in which the bandwidth share for real-time flows may still fall below the minimum but does this in a timescale much shorter than flow duration.

3. Long file transfer delays for non-real-time flows. Coexistence of the two traffic types creates file-transfer delays much larger in magnitude than that in the case where the system carries only non-real-time flows of the same traffic profile.

Although type 1 congestion may appear only during transient periods, only type 2 congestion may occur when the system reaches steady-state. A possible boundary behavior is depicted in Figure 1.3 for the case of voice and data, where the arrows denote directions of trajectories at the respective states. At the gray area flows depart faster than they arrive. Overloads do not occur and the system temporarily moves to states below the boundary. At the dark area, voice flows arrive faster that both voice and data flows depart but not so fast such that the system is driven above the boundary. In this case voice-flow admission control can effectively keep the system from overloading and type 2 congestion occurs. At the dashed area, data flows arrive faster than the rate at which both type of flows depart. Even though no new voice flows enter the system, the ones already in progress do not leave fast enough. This drives the system above the boundary and type 1 congestion occurs.

Chapter 2

Stability of Longest Queue First

This chapter is organized as follows. In section 2.1 we give an informal discussion of our model and provide some intuition for the proofs to be given in section 2.3. In section 2.2 we introduce the model and notation. section 2.3 is devoted to proving the main result of this chapter, Theorem 2.3.1. In section 2.4 we discuss how our results extend to general service rates.

2.1 Discussion and examples

In this section we give an informal description of our results by means of examples. We start by giving a graph representation of our model. See section 2.2 for a rigorous treatment.

A model is specified by a graph (V, E), with V and E being the vertex and edge set, respectively. One can think of the vertices as queues in which arrivals occur in discrete time at some given average rate. We consider independent, identically distributed (i.i.d.) arrival processes, although it is not hard to generalize our results to finite Markov-modulated processes. When a queue is served, a unit of work is removed from that queue, provided that it is nonempty. Not all queues/vertices can be served during the same time slot: if a queue is served then its neighboring queues are not. Hence, at each time, the served queues form an independent set of the graph. LQF scheduling chooses this set iteratively, starting from the longest queue and proceeding in a decreasing order of queue length. Any queue with a neighbor that has already been is not considered in the next iteration step. When two or more queues under consideration have equal backlog, a tie-breaking rule must be specified. This procedure is repeated until no further (nonempty) queue can be included. At each time slot, the served queues form a maximal independent set of the subgraph consisting of the nonempty queues.

2.1.1 Three-queue

As a first example, consider the graph with $V = \{1, 2, 3\}$ and $E = \{\{1, 2\}, \{2, 3\}\}$. Let λ_i be the arrival rate at queue i, for $i = 1, 2, 3$. We are concerned with characterizing the set of vectors $\lambda = (\lambda_1, \lambda_2, \lambda_3)$ for which the queueing process is stable, i.e., positive recurrent, under LQF scheduling. Since either queue 2 or queues 1 and 3 can be served at any given time, we expect that a scheduler can be stable only if $\lambda_1 + \lambda_2 < 1$ and $\lambda_2 + \lambda_3 < 1$. Indeed, the scheduler cannot serve queues 1 and 2 at a total rate larger than 1 since it cannot serve both queues at the same time. The same observation applies to queues 1 and 3.

To see why LQF is expected to be stable when these conditions hold we argue that the longest queue decreases, on average. This is immediate if one queue i is the only longest queue for some interval of time. Indeed, in that case, queue i is constantly served during that interval under LQF and its length tends to decrease because $\lambda_i < 1$. Now, it may happen that a subset L of queues alternate being the longest during some interval of time. For instance, assume that $L = \{1, 2\}$ so that the scheduler selects queue 2 some fraction of time during that interval and queues $\{1, 3\}$ the rest of the interval. During that interval of time, the scheduler always serves either queue 1 or queue 2, so that the total service rate of the queues 1 and 2 is equal to one. The total arrival rate of queues 1 and 2 is $\lambda_1 + \lambda_2 < 1$. Consequently, the length of the longest queue (which is that of queue 1 and also that of queue 2) decreases on average. The same argument can be made for any subset $L \subset \{1, 2, 3\}$[1]. The key property is as follows. For any set of queues L that alternate being the longest, there is a subset F of L that is served at a constant total rate, independently

[1]We use $\subset$ in the sense of non-strict inclusion.

of the selections that the scheduler makes, given that L is the set of longest queues. This constant service rate must be larger than the total arrival rate into the set F, otherwise the system would certainly be unstable. In that case, the longest queue, which is that of any queue in F, decreases on average. Note that this is a topological property of the graph. In our example, if $L = \{1, 2, 3\}$, one can choose $F = \{1, 2\}$. If $L = \{1, 3\}$, one can choose $F = \{1\}$, and so on.

The above condition is captured in our local pooling condition in Definition 2.2.1. A characterization of local pooling in terms of a linear program is given in Proposition 2.2.3. Using the fluid limit technique we are able to use an essentially deterministic argument to show stability under stochastic (and deterministic) arrivals (c.f. subsection 2.3.12).

Summing up, the argument shows that if the set of longest queues satisfies local pooling for some interval of time, then the longest queue tends to decrease during that interval. Consequently, if all sets satisfy local pooling, the system is stable. Thus, the stability conditions for systems that satisfy local pooling are only in terms of average intensities of arrivals. However, for other systems, the stability of LQF may depend on "second-order" properties. In particular, deterministic and nondeterministic arrival processes, with the same rates, may lead to unstable and stable behavior, respectively. We describe one such example next.

2.1.2 Six-Cycle

Consider the system specified by the 6-vertex cycle graph with $V = \{1, \ldots, 6\}$ and $E = \{\{1, 2\}, \{2, 3\}, \ldots, \{6, 1\}\}$ depicted in Figure 2.1.2. One can check that local pooling fails to hold. Indeed, if the set of longest queues is $L = \{1, 2, 3, 4, 5, 6\}$, there is no nonempty subset of L that LQF serves at a constant rate. For instance, LQF serves either 0 or 1 queue from the subset $\{1, 2\}$; it serves either 1 or 2 queue from $\{1, 2, 3\}$ and either 1 or 2 queues from $\{1, 2, 3, 4\}$, and similarly for all other subsets. Consequently, we cannot conclude the stability of this system using the argument in the three-queue example. In fact, the system may not be stable, as we explain next.

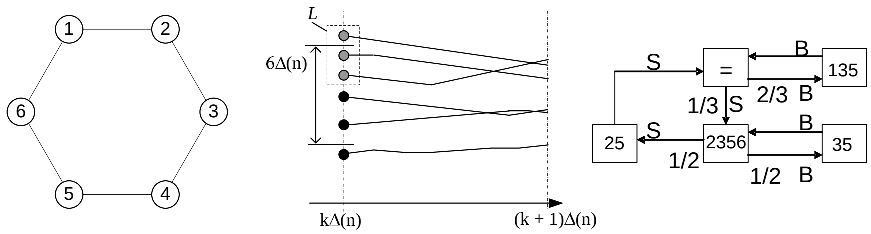

Figure 2.1. Left: A 6-cycle system; neighboring queues cannot be served at the same time. Middle: The separation argument. Right: Possible schedules.

Assume that a constant (deterministic) amount λ of work arrives during each time slot at each queue. Furthermore, assume that the scheduler uses LQF and breaks ties by throwing an unbiased die with as many faces as there are tied queues. We claim that this system is unstable if $\lambda > 4/9$. What makes this fact interesting is that, as we show below, the same system with non-constant i.i.d. arrivals with average value λ is stable whenever $\lambda < 1/2$.

The maximal independent sets of this graph are $M_1 = \{1, 3, 5\}, M_2 = \{2, 4, 6\}, M_3 = \{1, 4\}, M_4 = \{2, 5\}$, and $M_5 = \{3, 6\}$. Assume the system starts with all queues having the same large backlog. Let us call the sets M_1 and M_2 "big matches" and the other sets $M_3 - M_5$ "small matches." There are a number of possible sequences of choices that the LQF scheduler can make. These sequences correspond to choosing a small or big match at each step, whenever the choice is possible. The right part of the figure shows these sequences. Each rectangle represents a representative set of possible longest queues, up to a symmetric renumbering. For instance the queues start all equal; the scheduler chooses a small match, represented by S, with probability 1/3; then it selects a big match B with probability 1/2; then the scheduler is forced to pick a big match again, and so on. By analyzing this Markov chain that describes the choices of the scheduler, one finds that the queues become equal again after a random number of steps. Dividing the average number of times each queue is served by the average number of steps, we find that the scheduler serves each queue at the average rate of 4/9. It follows that the system is unstable if $\lambda > 4/9$.

We now explain why this system is stable when $\lambda < 1/2$ if the arrivals are allowed to vary even slightly. We saw that the instability of the system with deterministic arrivals is caused by the locked step operations that cause the queue lengths to become all equal every so often, which occasionally leads LQF to serve only two queues. The randomness of the arrivals prevents the backlogs from being equal a significant fraction of the time.

The key idea is that the longest queue length tends to decrease whenever it is large. Specifically, for $n \gg 1$, there is some t_0 and some ϵ_* such that

$$\frac{1}{n}E[\max_i Q_i(nt_0) - \max_i Q_i(0)|Q(0)] \leq -\epsilon_* t_0 < 0, \text{on } \{\max_i Q_i(0) > nh\}. \tag{2.1}$$

It follows that $\max_i Q_i(nt)$ is a Lyapunov function for the Markov chain, which implies its positive recurrence. To show (2.1), one considers a sequence of processes Q^n that behave as the original system but with initial conditions bounded by $O(n)$. One shows that $Q^n(nt)/n$ converges along a subsequence to $\bar{Q}(t)$, a deterministic system (called the "fluid limit") that satisfies a set of differential equations. Moreover, one proves that

$$\bar{Q}(t) = 0 \text{ for } t \geq t_0. \tag{2.2}$$

This fact, together with the convergence, implies that

$$\frac{1}{n}\max_i Q_i(nt) \to 0, \text{ for } t \geq t_0.$$

With a uniform integrability argument, one concludes that $E[\max_i Q_i(nt)/n] \to 0$ for $t \geq t_0$, and this implies (2.1). This line of argument is standard. The novelty in this work is in the argument to show (2.2). The main step is to show that if $\max_i \bar{Q}_i(t) > 0$, then

$$\frac{d}{dt}\max_i \bar{Q}_i(t) \leq -\epsilon_* < 0.$$

If the set of longest queues satisfies the local pooling condition, then one obtains the inequality above as in the three-queue example. If it does not, then the argument is qualitatively different.

For the 6-cycle example, the only set of queues that does not satisfy the local pooling condition is the set of all queues. Assume then that all the queues are equally large at some time t, in this limit $\bar{Q}(t)$. This implies that all the queue lengths are of order $n\delta$ for the system Q^n at time nt. In that case, all the queues will remain nonempty for some interval of time of order $n\delta$, since there is at most one departure from any given queue in each step.

One then shows that the difference between the maximum and the minimum queue lengths is at least of order $6\Delta(n)$ for most of the times $s \in I := \{nt, nt + \Delta(n), nt + 2\Delta(n), \ldots, nt + n\delta - \Delta(n)\}$, where $\Delta(n) = n^{1/6}$ and $t > 0$ is fixed. We explain the main steps of that argument below. Moreover, if that condition holds for some time $s = k\Delta(n) \in I$, there is a proper subset L of the queues that remain strictly longer than the others during the interval $I(s) := \{s, s+1, \ldots, s + \Delta(n) - 1\}$. (See Figure 2.1.2.) In addition, any proper subset of the queues satisfies the local pooling conditions. According to the argument we used for the three-queue example, it follows that the longest queue decreases, on average, during $I(s)$. Consequently, the longest queue decreases on average during most of the intervals $I(s)$ and, therefore, most of the time. This argument implies that

$$\frac{1}{n} E[\max_i Q_i(nt + n\delta) - \max_i Q_i(nt) | Q(nt)] \leq -\delta\epsilon_* < 0.$$

This fact implies the stability of the system.

We now explain why the queues remain separated for most of the times $s \in I$. The idea is that the difference between the maximum and the minimum queue lengths is lower-bounded by differences between cumulative arrivals. We approximate these differences by Gaussian random variables, after proper scaling. Using the Gaussian distribution, we find that these random variables cannot be small most of the time.

First, observe that if all the queues remain nonempty the number of queues served among queues $\{2, 3\}$ is always the same as among queues $\{5, 6\}$ under any service set M_i. Designate by $Q_i(t)$ the length of queue i at time t. If no queue empties, since the departures cancel out,

$$B(t) := Q_2(t) + Q_3(t) - Q_5(t) - Q_6(t) = B(0) + A_2(t) + A_3(t) - A_5(t) - A_6(t) =: B(0) + Z(t)$$

where $A_i(t)$ is the cumulative number of arrivals in queue i up to time t.

Second, note that $D(t) := \max_i Q_i(t) - \min_i Q_i(t) \geq |B(t)|/2$, so that

$$\frac{D(nt)}{\sqrt{n}} \geq \frac{1}{2\sqrt{n}} |B(0) + Z(nt)|.$$

Finally, since the arrivals are i.i.d., the random variable $Z(nt)/\sqrt{n}$ is approximately Gaussian and zero-mean, so that

$$E\sum_{m=1}^{n\delta} 1\{D(nt+m) \leq 6\Delta(n)\} \approx n\delta P\left(\frac{D(nt)}{\sqrt{n}} \leq \frac{6\Delta(n)}{\sqrt{n}}\right)$$
$$\lessapprox n\delta P\left(\frac{|Z(nt)|}{2\sqrt{n}} \leq \frac{6\Delta(n)}{\sqrt{n}}\right) \lessapprox O(\sqrt{n}\Delta(n)\delta),$$

as $n \to \infty$. Choosing $\Delta(n)$ so that $\sqrt{n}\Delta(n) = o(n/\Delta(n))$ – for instance, $\Delta(n) = n^{1/6}$ – we conclude that the fraction of the values of $s \in I = \{nt, nt+\Delta(n), nt+2\Delta(n), \ldots, nt+n\delta-\Delta(n)\}$ such that $D(s) \leq 6\Delta(n)$ is negligible, because $|I| = n\delta/\Delta(n)$. This step completes the argument.

Looking back, here are the main ideas of the argument. When the set of longest queues satisfies local pooling, the longest queue tends to decrease. Also, the set of all queues – for which local pooling fails – cannot remain longest for any significant fraction of time: the scheduler cannot fully compensate the fluctuations in the arrivals because the set of service vectors has only rank 4 in the 6-dimensional space of queue lengths. Summing up, the key stability condition is that when the set of longest queues does not satisfy local pooling, the rank of the corresponding service vectors is small (at most the number of queues minus two is sufficient, as will become evident in the proof of Lemma 1 in Appendix B).

These ideas form the basis of Theorem 2.3.1, where we show that LQF is stable for any nondeterministic i.i.d. arrival processes, under a weakening of the local pooling condition.

When (the weakened) local pooling condition fails to hold, we have simulation evidence suggesting the system may be unstable, even for Bernoulli arrivals. One such example is a system specified by an 8-vertex cyclic graph driven by Bernoulli arrivals at rate $\lambda_i = 0.4984$ for $i = 1, \ldots, 8$; ties were broken as in the 6-vertex example. In Figure 2.1.2, we plot the queue length evolution of this system.

Sufficient stability conditions for LQF in cyclic graphs, are considered in [28]. There, the analysis is based on the stability of the associated fluid system. It is found that if the total arrival rate of three consecutive queues on the cycle is strictly below 1, then the system is stable. Intuitively, if all queues alternate being the longest for some considerable

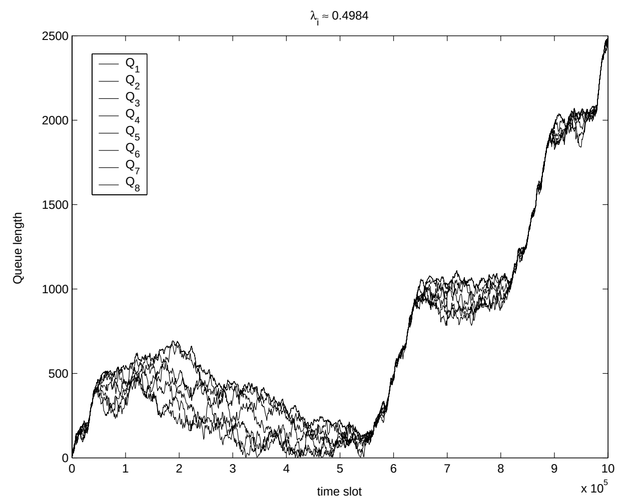

Figure 2.2. A case of possible instability in a 8-cycle system with $\lambda_1 = \ldots = \lambda_8 = 0.4984$.

amount of time, then the three queues satisfying the condition above must decrease because at least one of them is served at each timeslot. On the other hand, under nondeterministic arrivals, this condition on the rates is not necessary for the 6-cycle as described above. For systems not of the special form as the 6-cycle, i.e., those satisfying the rank condition in Theorem 2.3.1, it is possible to still derive stability by imposing stricter than feasibility, conditions on the arrival rates. (See the remarks following Theorem 2.3.1.)

2.2 Model and notation

We consider a discrete-time model of a set, K, of queues. Let $Q_k(t)$ denote the backlog of queue $k \in K$ at time slot $t \geq 0$, and let $A_k(t)$ and $D_k(t)$ respectively denote the cumulative arrivals and departures of work at queue k up to time t. We assume that arrivals are mutually independent i.i.d. processes with $E[A_k(1)] = \lambda_k < \infty$ for every k. Additionally, the processes $A_k(\cdot)$ are independent for each k.

The scheduler cannot serve all nonempty queues at any given time. The allowable schedules form a finite set $S[K] \subset \{0,1\}^{|K|}$. At any time t, the scheduler selects a vector $m \in S[K]$ and removes $m_k \in \{0,1\}$ units of work from each queue k, unless queue k empties first. (We relax the '$\{0,1\}$-assumption' in section 2.4.)

By $M[K]$ we denote the maximal elements of $S[K]$, and treat it as a matrix whose columns are its elements. For $L \subset K$, we use $S[L]$ to denote the restriction of $S[K]$ to L and $M[L]$ to denote the maximal elements of $S[L]$. This model is a special case of the generalized switch of [39] and is similar to the model of the input-buffered packet switch of [15].

Given $M[K]$, the evolution of queue backlogs is completely determined by the arrival process, $A(\cdot)$, and the scheduling policy. For $m \in M[K]$, $T_m(t)$ records the number of time slots, up to and including t, during which the scheduling policy chooses m.

Thus, $(Q(\cdot), T(\cdot))$ satisfies the following equations:

$$Q(t) = Q(0) + A(t) - D(t), \tag{2.3}$$

$$D_k(t+1) - D_k(t) = \begin{cases} \sum_{m \in M[K]} m_k(T_m(t+1) - T_m(t)) & \text{if } Q_k(t) > 0\,, \\ 0 & \text{otherwise}\,, \end{cases} \tag{2.4}$$

$$\sum_{m \in M[K]} T_m(t) = t, \quad \forall t \geq 0, \tag{2.5}$$

$$Q_k(t) \geq 0, \quad \forall k \in K \text{ and } t \geq 0\,, \tag{2.6}$$

$$D_k(t), T_m(t) \geq 0, \text{ nondecreasing } \forall k \in K, m \in M[K], t \geq 0. \tag{2.7}$$

We consider all the above processes to be defined for noninteger times as well: for noninteger u, $D(u) = D(\lfloor u \rfloor)$, and the other processes are extended similarly.

LQF is a stationary policy that at any time t selects a vector $n \in M[K]$ as follows. First, the queues are sorted in decreasing order according to their backlogs. Given such an ordering $k(1), \ldots, k(|K|)$ (where $Q_{k(1)}(t) \geq \ldots \geq Q_{k(|K|)}$), n is constructed iteratively by considering the "longest queues first": In the first iteration, we start with $M_1 := \{m \in M[K] : m_{k(1)} = 1\}$, i.e. all service vectors activating the longest queue. If $M[K]$ is nondenegerate then $M_1 \neq \emptyset$. Moreover, we set $n_{k(1)} = 1$. In the i-th iteration (for $i > 1$), the search is confined to vectors $M_i \subset M_{i-1}$ that activate queue $k(i)$ if possible; i.e. $M_i := \{m \in M_{i-1} : m_{k(i)} = 1\}$ when $M_i \neq \emptyset$ (in which case we set $n_{k(i)} = 1$). If no such vectors exist (i.e. $m_{k(i)} = 0$ for all $m \in M_{i-1}$), then we set $n_{k(i)} = 0$, and $M_i := M_{i-1}$. Having fixed $n_{k(i)}$, we continue with the $(i+1)$-th iteration until we have enumerated all queues in K and chosen all components of the service vector n. To completely describe such a policy one has to specify (stationary) rules for breaking ties in the queue backlogs. Once such rules are fixed, the resulting process $(Q(t), t = 0, 1, 2, \ldots)$ becomes a Markov chain.

Consider a vector of arrival rates $\lambda \in \mathbb{R}_+^K$. If for this λ the process $Q(\cdot)$ has a stationary distribution, then the time-averaged service rates must exceed the arrival rates. Thus, by the ergodic theorem, $\lambda \leq \phi$, i.e. $\lambda_k \leq \phi_k$ for all $k \in K$, for some $\phi \in \mathrm{Co}(M[K])$. (For more details on this, see [3]. For a matrix M, $\mathrm{Co}(M)$ denotes the convex hull of its columns.) The vector $\lambda \in \mathbb{R}_+^K$ is called *feasible* if there exists a $\phi \in \mathrm{Co}(M[K])$ such that $\lambda < \phi$.

For a vector v we let v' denote its transpose, and by e denote vectors whose coordinates are equal one. If $v \in \mathbb{R}^K$ and $L \subset K$ then we write v_L for the vector $(v_i, i \in L) \in \mathbb{R}^L$. We use $\mathbb{Q}$ to denote the rationals and $\mathbb{Q}_+$ to denote the nonnegative rationals. Finally, we let $a \vee b = \max(a, b)$.

Definition 2.2.1. *We say that* $L \subset K$ satisfies local pooling *if there exists a nonzero vector* $\alpha \in \mathbb{R}_+^K$ *such that* $\alpha' \phi$ *is a positive constant for all* $\phi \in \mathrm{Co}(M[L])$. *We say that* local pooling *is satisfied if every* $L \subset K$ *satisfies local pooling.*

Note that if L satisfies local pooling, no vector in $\mathrm{Co}(M[L])$ strictly dominates (in all coordinates) any other vector in $\mathrm{Co}(M[L])$.

Remark 2.2.2. If L satisfies local pooling then, for any feasible vector λ and $\phi \in \mathrm{Co}(M[L])$, $\lambda_k < \phi_k$ for some $k \in L$.

For a fixed, feasible λ we define

$$\epsilon_* := \inf\{\max_{k \in L}(\phi_k - \lambda_k) \mid L \subset K, L \text{ satisfies local pooling}, \phi \in \mathrm{Co}(M[L])\}. \tag{2.8}$$

By Remark 2.2.2, $\epsilon_* > 0$. We interpret ϵ_* as the excess of the service rate over the arrival rate that must exist at some queue whenever the set of longest queues satisfies local pooling.

The observation above yields an equivalent characterization of local pooling in terms of a linear program.

Proposition 2.2.3. *Consider the LP:*

$$\max_{c,\mu,\nu} c \tag{2.9}$$

$$s.t.\;\; M[L]\mu \geq M[L]\nu + ce \tag{2.10}$$

$$e'\mu = 1 \tag{2.11}$$

$$e'\nu = 1 \tag{2.12}$$

$$\mu, \nu \in \mathbb{R}^r_+, c \in \mathbb{R}\,, \tag{2.13}$$

where r is the number of columns of $M[L]$.

The set L satisfies local pooling if and only if the optimal value is $c^* = 0$.

Proof. Let $M = M[L]$ and first assume $c^* = 0$. Then, for any $M\mu^* \in \mathrm{Co}(M[L])$ (with $e'\mu^* = 1$), $(c, \mu, \nu) = (0, \mu^*, \mu^*)$ must be a solution since (2.10) is satisfied with equality for $\nu = \mu$ and $c = 0$. Hence, complementary slackness and dual feasibility must hold for some dual variables α^*, β^*, and γ^*, corresponding to (2.10),(2.11),and (2.12). By complementary slackness, we have

$$(\alpha^{*\prime} M - \beta^* e')\mu^* = -(\alpha^{*\prime} M + \gamma^* e')\mu^* = 0\,.$$

This fact and the feasibility of $(0, \mu^*, \mu^*)$, implies $\beta^* = -\gamma^*$. Dual feasibility yields $\alpha^* \geq 0$, $\alpha^{*\prime} e = 1$, and

$$-\gamma^* e' \leq \alpha^{*\prime} M \leq \beta^* e'\,.$$

Since $\beta^* = -\gamma^*$, the last display implies that $\alpha^{*\prime} M = \beta^* e'$. Since $\alpha^{*\prime} e = 1$, $\alpha^* \geq 0$ is nonzero, and $\beta^* > 0$ by the nondegeneracy of M. This implies local poofling holds for L, according to Definition 2.2.1 with α^* as the required nonzero vector.

To prove the converse, it suffices to note that under local pooling for L, no vector can dominate another in $\mathrm{Co}(M[L])$.

□

In the proof of Theorem 2.3.1 we make use of a large deviation bound on $A(\cdot)$. Henceforth, we assume that for each $k \in K$ and $\epsilon > 0$,

$$P\left[\left|\frac{A_k(n)}{n} - \lambda_k\right| > \epsilon\right] \leq \beta \exp(-n\gamma(\epsilon)), \quad \text{for all } n \geq 1, \tag{2.14}$$

for some $\gamma(\epsilon) > 0$ and $\beta > 0$.

2.3 Stability

In this section, we first state the main result and outline its proof. We then establish the main properties we need and proceed with the formal proof. We conclude the section with a remark about stability when local pooling holds.

2.3.1 Main result

Theorem 2.3.1. *Consider a system $M[K]$ such that every $L \subset K$ with $\operatorname{rank} M[L] \geq |L| - 1$ satisfies local pooling. Assume that the independent arrival processes $A_k(\cdot)$, one for each queue $k \in K$, are i.i.d., satisfy* (2.14)*, and have nonzero variance.*

The system is then stable under LQF for all feasible vectors λ.

According to the condition of the theorem, if a subset L does not satisfy local pooling then the rank off the service vectors is too small for the scheduler to be able to prevent the queue lengths from diverging.

In applications, it may be useful to infer stability for some given feasible λ rather than for all feasible rates as in Theorem 2.3.1. In these cases, Theorem 2.3.1 implies stability, under the same conditions on the arrival processes, when queues L that do *not* satisfy local pooling and do *not* have $\operatorname{rank} M[L] \geq |L| - 1$ are instead such that the optimum of the following linear program is negative,

$$\max_{c,\nu} c$$

$$\text{s.t. } \lambda_L \geq M[L]\nu + ce$$

$$e'\nu = 1$$

$$\nu \in \mathbb{R}^r_+, c \in \mathbb{R}\,.$$

(In the above, r is the number of columns of $M[L]$.) This guarantees (according to Lemma 2.3.10, below) that when queues in L are the longest, they must decrease in length at a certain rate even though local pooling does not hold for L.

2.3.2 Proof outline

The strategy of the proof is as follows.

1. We consider $Q^n(nt)/n$ where $Q^n(\cdot)$ is the vector of queue lengths in a system with initial backlogs not exceeding n, i.e., $\sum_k Q^n_k(0) \leq n$.

2. We show that $E(\sum_k Q^n_k(nt)/n) \to 0$ for $t \geq t_0$. This result provides a Lyapunov function for the stability of Q^n. Indeed, the limit implies that the expected queue lengths tend to decrease.

3. To prove the previous result, we use uniform integrability and the fact that $Q^n(nt)/n \to 0$ for $t \geq t_0$. To prove this claim, we show that $Q^n(nt)/n$ converges along a subsequence to $\bar{Q}(t)$, a fluid system described by a system of differential equations, and that $\bar{Q}(t) = 0$ for $t \geq t_0$, for some t_0.

4. To prove that $\bar{Q}(t) = 0$ for $t \geq t_0$, we show that as long as $\max_i \bar{Q}_i(t)$ is positive it must decrease at least at a given rate. Specifically, we show that if $\max_i \bar{Q}_i(t) > 0$ and t is a regular time, then $d \max_i \bar{Q}_i(t)/dt \leq -\epsilon_* < 0$.

5. To derive the rate of decrease, we use the observation that while the set of longest queues is fixed and satisfies local pooling, the queue lengths decrease at least at a given rate.

6. Finally, we show that most of the time, a set of queues that satisfies local pooling dominates the other queues. We demonstrate this by proving that if L does not satisfy local pooling, then a proper subset of L must dominate the other queues most of the time, which argument has two parts:

   (a) We show that the queues must be separated (i.e. differ in length) by some amount $\beta\Delta(n)|L|$ at most time steps that are multiples of Δ, as we did in the six-cycle argument.

   (b) We then use a uniform bound on the arrivals to conclude the domination by a subset of L during the subsequent $\Delta(n) - 1$ steps.

The first four ideas of the proof are identical to the general programme followed in [14] and [37]. However, the subsequent steps are quite different since, as explained in the examples, they involve second-order properties of the processes.

This section is organized as follows.

1. We first define the sequence of systems Q^n in Definition 2.3.2 and state the convergence to the fluid limit $\bar{Q}$ in Proposition 2.3.3.

2. We prove the separation of the queues at multiples of Δ in Lemma 2.3.4.

3. We show that queues remain separated for the subsequent $\Delta(n) - 1$ steps in Lemma 2.3.8.

4. We derive the rate of decrease when the longest queues satisfy longest pooling in Lemma 2.3.10.

5. We show that the fluid limits of the queue lengths decrease with an upper-bounded rate in Lemma 2.3.11.

6. Finally, we conclude the proof of Theorem 2.3.1 in section 2.3.8.

2.3.3 Fluid limit

We define a sequence of systems scaled in space and time by a factor n. Proposition 2.3.3 shows that they converge to a fluid limit.

Definition 2.3.2 (Sequence of scaled systems). *On the same probability space implicitly used above, we define a sequence of systems with the same $M[K]$ and indexed by $n = 1, 2, \ldots$. The initial values, $Q^n(0)$, for the n-th system satisfy $\sum_{i \in K} Q_i^n(0)/n \leq 1$. The arrival process is the same for all elements of the sequence, i.e. $A^n(\cdot) = A(\cdot)$. Let $Q^n(\cdot)$, $T^n(\cdot)$, $D^n(\cdot)$ respectively be the resulting queueing, service and departure processes under LQF. (All systems use the same LQF scheduling policy.)*

We use the fluid limits $(\bar{Q}(\cdot), \bar{T}(\cdot), \bar{D}(\cdot))$ to study the stability of the original (prelimit) processes. The existence and properties of these limits are stated in the following proposition, whose proof, as it is similar to that of Theorem 4.1 of [14] and that of Theorem 3 of [15], is omitted.

Proposition 2.3.3 (Fluid limit). *A limit $(\bar{Q}(\cdot), \bar{T}(\cdot), \bar{D}(\cdot))$ of $(Q^n(\cdot), T^n(\cdot), D^n(\cdot))$ as $n \to \infty$ exists a.s., in the topology of uniform convergence over compact sets, along some subsequence, and satisfies the following properties:*

$$\bar{Q}(t) = \bar{Q}(0) + \lambda t - \bar{D}(t), \quad \forall t \geq 0;$$

$$\sum_{m \in M[K]} \dot{\bar{T}}_m(t) = 1, \quad \forall t \geq 0;$$

$$\dot{\bar{D}}_k(t) = \sum_{m \in M[K]} m_k \dot{\bar{T}}_m(t), \text{ if } \bar{Q}_k(t) > 0, \quad \forall k \in K, t \geq 0;$$

$$\bar{Q}_k(t) \geq 0, \quad \forall k \in K, t \geq 0;$$

$$\sum_{k \in K} \bar{Q}_k(0) \leq 1;$$

$$\bar{D}_k(t), \bar{T}_m(t) \geq 0, \text{ nondecreasing } \forall k \in K, m \in M[K], t \geq 0.$$

Moreover, $\bar{Q}(\cdot), \bar{T}(\cdot), \bar{D}(\cdot)$ are absolutely continuous. Times $t \geq 0$ for which the derivatives of $\bar{Q}(t), \bar{T}(t), \bar{D}(t)$ exist will be called regular *times.*

2.3.4 Separation of the queues at multiples of $\Delta(n)$

Let $\Delta(n) = n^{1/6}$. The following lemma states that the queue lengths remain separated at most of the multiples of $\Delta(n)$.

Lemma 2.3.4 (Queue separation)**.** *Consider any nonempty set $L \subset K$ that does not satisfy local pooling. For all $1 \geq \delta > \delta_0 > 0$, and $\alpha > 0$, we have*

$$\frac{1\{\min_{i\in L} Q_i^n(nt) > n\delta\}}{n^{\frac{1}{2}+\alpha}\Delta(n)} \sum_{m=n\delta_0}^{n\delta} 1\{\max_{i\in L} Q_i^n(nt+m) - \min_{i\in L} Q_i^n(nt+m) \leq \gamma\Delta(n)\} \to 0 \quad (2.15)$$

in probability, as $n \to \infty$ with $\gamma := 1.1\,|L|\,(\max_{i\in K}\lambda_i \vee 1)$.

Therefore, with probability 1, for any subsequence (n_k^0) there is a further subsequence, (n_k^1), along which the convergence in (2.15) *holds simultaneously for all sets L that do not satisfy local pooling, for all $t \in \mathbb{Q}_+$, and for all $\delta, \delta_0 \in \mathbb{Q}$ with $1 \geq \delta > \delta_0 > 0$.*

Before we proceed to the proof, we need to prove Lemma 2.3.5 first.

Let S_m^n be the partial sums of a triangular array, $((X_m^n)_{m=1,\ldots,n})_{n\geq 1}$, of i.i.d. r.v's with nonzero variance and finite third moment. (That is, $S_m^n = X_1^n + \ldots + X_m^n$.) Let (Y_n) be a sequence of r.v's such that for each n, Y_n is independent of $(S_m^n, m = 1,\ldots,n)$.

Lemma 2.3.5. *For all $1 \geq \delta > \delta_0 > 0, \alpha > 0$,*

$$\frac{1}{n^{\frac{1}{2}+\alpha}\Delta(n)} \sum_{m=n\delta_0}^{n\delta} 1\{|Y_n + S_m^n| \leq \Delta(n)\} \to 0$$

in probability as $n \to \infty$.

Proof. Let Z be a standard normal random variable. Then, by the Berry-Esseen theorem (see, e.g. Theorem 2.4.9 of Durrett [17]), for all sufficiently large n we have

$$\begin{aligned}\frac{P\left[\left|S_m^n - mEX_1^1\right| \leq \Delta(n)\right]}{\sup_y P\left[|y + S_m^n| \leq \Delta(n)\right]} &\geq \frac{P\left[|Z| \leq \Delta(n)/\sqrt{m}\right] - c_1/\sqrt{m}}{P\left[|Z| \leq \Delta(n)/\sqrt{m}\right] + c_1/\sqrt{m}} \\ &\geq 1 - \frac{2}{c_2\Delta(m) + 1}\,,\end{aligned}$$

for some constants $c_1, c_2 > 0$ and an m such that $n\delta_0 \leq m \leq n\delta$.

Hence,

$$\sum_{m=n\delta_0}^{n\delta} P\left[|Y_n + S_m^n| \le \Delta(n)\right] \le c_3 \sum_{m=n\delta_0}^{n\delta} P\left[\left|\frac{S_m^n - mEX_1^1}{\sqrt{m}}\right| \le \frac{\Delta(n)}{\sqrt{n\delta_0}}\right]$$
$$\le c_3 n P\left[|Z| \le \frac{\Delta(n)}{\sqrt{n\delta_0}}\right] + nO\left(\frac{1}{\sqrt{n}}\right)$$
$$= O\left(\sqrt{n}\Delta(n)\right) \ , \quad \text{as } n \to \infty \, ,$$

for some constant $c_3 > 0$. The result follows by Markov's inequality. □

Proof of Lemma 2.3.4. By the rank assumption of the Theorem 2.3.1, there must exist some nonzero vector $v^L \in \mathbb{R}^{|L|}$ such that $v^L e = 0, v^L M[L] = 0$, and $\max_{i\in L} |v_i^L| = 1$. Now, provided that $\min_{i\in L} Q_i^n(nt) > n\delta$, queues in L do not empty during $\{nt, \cdots, n(t+\delta)\}$; thus,

$$v^L Q_L^n(nt+m) = v^L Q_L^n(nt) + v^L\left[A_L(nt+m) - A_L(nt)\right] \, .$$

We can apply Lemma 2.3.5 to

$$S_m^n = \gamma^{-1} v^L\left[A_L(nt+m) - A_L(nt)\right] \, ,$$

with $\gamma := 1.1\,|L|\,(\max_{i\in K}\lambda_i \vee 1)$ and $Y_n = \gamma^{-1} v^L Q_L^n(nt)$ for $n = 1, 2, \ldots$, to conclude that for all δ and δ_0, $1 \ge \delta > \delta_0 > 0$, and $\alpha > 0$,

$$\frac{1\{\min_{i\in L} Q_i^n(nt) > n\delta\}}{n^{\frac{1}{2}+\alpha}\Delta(n)} \sum_{m=n\delta_0}^{n\delta} 1\{\left|v^L Q_L^n(nt+m)\right| \le \gamma\Delta(n)\} \to 0 \, ,$$

in probability as $n \to \infty$.

Moreover, we can find some constant $\kappa > 0$, which depends on v^L, such that

$$\max_{i\in L} Q_i^n(nt+m) - \min_{i\in L} Q_i^n(nt+m) \ge \kappa \left|v^L Q_L^n(nt+m)\right| \ , \quad m \ge 0.$$

Consequently, for all δ and δ_0, $1 \ge \delta > \delta_0 > 0$, and $\alpha > 0$,

$$\frac{1\{\min_{i\in L} Q_i^n(nt) > n\delta\}}{n^{\frac{1}{2}+\alpha}\Delta(n)} \sum_{m=n\delta_0}^{n\delta} 1\{\max_{i\in L} Q_i^n(nt+m) - \min_{i\in L} Q_i^n(nt+m) \le \gamma\Delta(n)\} \to 0 \, ,$$

in probability as $n \to \infty$.

□

2.3.5 Separation of the queues during intervals

To show that a set of queues that satisfies local pooling dominates the other queues most of the time, we use a bound on the fluctuations of the queue lengths. This bound relies on the following property of the arrivals.

Lemma 2.3.6 (Property of arrivals). *The following event has probability 1: For all $T \in \mathbb{Q} \cap (0, +\infty)$,*

$$\limsup_{n} \max_{0 \le i \le n^{6/7}T, k \in K} \left| \frac{A_k((i+1)n^{1/7}) - A_k(n^{1/7})}{n^{1/7}} - \lambda_k \right| = 0. \tag{2.16}$$

Proof. Fix any $T, \epsilon > 0$ and note that

$$\sum_{n=1}^{\infty} P\left[\max_{0 \le i \le n^{6/7}T, k \in K} \left| \frac{A_k((i+1)n^{1/7}) - A_k(n^{1/7})}{n^{1/7}} - \lambda_k \right| > \epsilon \right]$$
$$\le \beta T \sum_{n=1}^{\infty} n^{6/7} \exp(-n^{1/7}\gamma(\epsilon)) < \infty .$$

The result now follows by the Borel-Cantelli lemma. □

Definition 2.3.7 (Domination). *Let $I \subset \{0, 1, \ldots\}$ and $L, L' \subset K$. We say that L is not dominated by L' during I for the system Q^n if there exist $u \in I, i \in L$, and $j \in L'$ such that $Q_i^n(u) \ge Q_j^n(u)$.*

Fix any sequence $(n_k^0) \subset \mathbb{Z}_+$ with $n_k^0 \to \infty$ as $k \to \infty$. Pick a subsequence (n_k) of (n_k^0) such that the convergence in Lemma 2.3.4 is achieved almost surely. From now on we fix a sample point, ω, such that that limit is attained and (2.16) holds.

We have the following result.

Lemma 2.3.8. *For $n \ge 1$, let $g(n) = n^{\frac{1}{2}+\alpha}\Delta(n)$ for $\alpha > 0$, such that $g(n)\Delta(n) = o(n)$ as $n \to \infty$. Assume that there exist $\delta \in \mathbb{Q} \cap (0, 1)$ and $K' \subset K$ such that K' dominates $K \setminus K'$ during $\{n_k t, \cdots, n_k(t+\delta)\}$, for all sufficiently large k. Also, assume that $Q_i^{n_k}(n_k t) > n_k \delta$ for all $i \in K'$ and fix a $\delta_0 \in \mathbb{Q}$ with $0 < \delta_0 < \delta$.*

Divide $[n_k t, n_k(t+\delta)]$ into intervals of size $\Delta(n_k)$, i.e.

$$I_j^{n_k} = [n_k t + (j-1)\Delta(n_k), n_k t + j\Delta(n_k)]$$

with $j \in \{1, 2, \ldots, \delta n_k/\Delta(n_k)\}$.

The number of intervals not dominated by some set L *that satisfies local pooling for the system* Q^{n_k} *does not exceed* $g(n_k) + \delta_0 n_k/\Delta(n_k)$ *as* $k \to \infty$.

Proof. If K satisfies local pooling then there is nothing to prove, so consider a nonempty set $L \subset K$ that does not satisfy local pooling. From Lemma 2.3.4, the number of intervals $I_j^{n_k}$ in which

$$\max_{i \in L} Q_i^{n_k}(t) - \min_{i \in L} Q_i^{n_k}(t) \leq \gamma\Delta(n_k)$$

at the left-hand endpoint, t, of the interval is $o(g(n_k))$ and, consequently, does not exceed $g(n_k) + \delta_0 n_k/\Delta(n_k)$ as $k \to \infty$.

For sufficiently large k large, the arrivals at each queue j in all the intervals are bounded by $1.1\lambda_j$, since ω satisfies (2.16). Moreover, there is at most one departure from each queue in each time step. It is easy to check that if the queue lengths are separated by more than $\gamma\Delta(n_k)$ at the left-hand endpoint of an interval $I_j^{n_k}$ of size $\Delta(n_k)$, there must be a proper subset, L', of L that dominates the other queues during that interval. (This fact is illustrated in the right part of Figure 2.1.2.)

Consequently, the number of intervals in which any set L that does not satisfy local pooling is not dominated by some $L' \subsetneq L$ does not exceed $g(n_k) + \delta_0 n_k/\Delta(n_k)$ as $k \to \infty$. Since all subsets of L with fewer than four elements satisfy local pooling, the number of intervals not dominated by some L that does satisfy local pooling does not exceed $g(n_k) + \delta_0 n_k/\Delta(n_k)$ in the limit.

□

2.3.6 Rate of decrease under Local Pooling

Here our main objective is to prove Lemma 2.3.10, which concerns the rate of decrease of a local fluid limit. The local fluid limit idea comes from [37], and we will use it for proving Lemma 5.

Fix $L \subset K$. Assume that $(r_k), (n_k) \subset \mathbb{Z}_+$ are sequences for which $r_k/n_k \to r < \infty$, and

$$Q_i^{n_k}(r_k + u\Delta(n_k)) > Q_j^{n_k}(r_k + u\Delta(n_k)) ,$$

for all $u \in [0,1], i \in L, j \notin L$, and that the ordering π in L, given by $(Q_i^{n_k}(r_k))_{i \in L}$ is kept fixed for all k. Let i_0 be a maximum element of L according to π. We define

$$\hat{Q}_i^{n_k}(u) = \frac{Q_i^{n_k}(r_k + u\Delta(n_k)) - Q_{i_0}^{n_k}(r_k)}{\Delta(n_k)} \quad , u \in [0,1], i \in L .$$

For $u \in [0,1]$ we also define $\hat{F}^{n_k}(u) = (F^{n_k}(r_k + u\Delta(n_k)) - F^{n_k}(r_k))/\Delta(n_k)$ for $F^{n_k} \in \{A^{n_k}, D^{n_k}, T^{n_k}\}$. Limits will be taken with respect to the topology of uniform convergence on $[0,1]$.

Proposition 2.3.9. *The limit* $(\hat{Q}_L^\infty(\cdot), \hat{A}_L^\infty(\cdot), \hat{D}_L^\infty(\cdot), \hat{T}^\infty(\cdot))$ *of* $(\hat{Q}_L^{n_k}(\cdot), \hat{A}_L^{n_k}(\cdot), \hat{D}_L^{n_k}(\cdot), \hat{T}^{n_k}(\cdot))$ *as* $k \to \infty$*, exists over some subsequence and satisfies*

$$\hat{Q}_i^\infty(u) = \hat{Q}_i^\infty(0) + \lambda_i u - \hat{D}_i^\infty(u) \quad \forall u \in [0,1] \tag{2.17}$$

$$\hat{D}_i^\infty(\cdot) \quad \textit{is nonnegative and nondecreasing,} \quad \hat{D}^\infty(0) = 0 , \tag{2.18}$$

$$u^{-1}\hat{D}_L^\infty(u) \in \mathrm{Co}(M[L]) , \forall u \in (0,1] . \tag{2.19}$$

Proof. Note that $\hat{Q}_i^{n_k}(0) \le 0$ for all $i \in L$ and large k, so $\hat{Q}^\infty(0)$ exists in $[-\infty, 0]$, over a subsequence. We approximate $T^{n_k}(\cdot)$ by continuous processes such that the uniform error in $[r_k, r_k + \Delta(n_k)]$ is no more than n_k^{-1}. (For example, at noninteger times approximate by the linear interpolation of the values at the two neighboring integer times.) These processes are Lipschitz continuous so, by the Arzela-Ascoli theorem, the limit $\hat{T}^\infty(\cdot)$, of $\hat{T}^{n_k}$ as $k \to \infty$, exists over some subsequence of (n_k).

Properties (2.17) and (2.18) follow by uniform convergence of $\hat{A}^{n_k}(\cdot)$ (which in turn follows by (2.16)), continuity, and (2.3) and (2.7).

Since queues in L are given higher priority by LQF than those in $K \setminus L$ during $[r_k, r_k + \Delta(n_k)]$, and no queue in L is empty during the same interval, from (2.4) and (2.5) we obtain

$$\frac{D_L^{n_k}(r_k + u\Delta(n_k)) - D_L^{n_k}(r_k)}{u\Delta(n_k)} \in \mathrm{Co}(M[L]) .$$

□

Lemma 2.3.10. *If local pooling holds for L then, for any ordering, π, of the elements of L, any local limit $\hat{Q}_L^\infty(\cdot)$ of $\hat{Q}_L^{n_k}(\cdot)$ satisfies*

$$\max_{i\in L}\hat{Q}_i^\infty(1) - \max_{i\in L}\hat{Q}_i^\infty(0) \le -\epsilon_* .$$

Proof. Without loss of generality, consider a regular time $t \in (0,1)$ at which $\hat{Q}_i^\infty(t) = \hat{Q}_j^\infty(t) > -\infty$ and $\dot{\hat{Q}}_i^\infty(t) = \dot{\hat{Q}}_j^\infty(t)$ for all $i,j \in \arg\max_{l\in L}\hat{Q}_l =: \hat{L}$. Choose a $\delta > 0$ sufficiently small that $\hat{Q}_i^\infty(u) > \hat{Q}_j^\infty(u)$ for all $i \in \hat{L}, j \in L\setminus\hat{L}, u \in [t, t+\delta]$.

Fix an element of the probability space where, by Proposition 2.3.9, the fluid limit exists. Let (n_k) be any sequence along which $(\hat{Q}_L^{n_k}(\cdot), \hat{A}_L^{n_k}(\cdot), \hat{D}^{n_k}(\cdot), \hat{T}^{n_k}(\cdot))$ converges to a fluid limit uniformly on compact sets. For sufficiently large k, $\hat{Q}_i^{n_k}(u) > \hat{Q}_j^{n_k}(u)$ for all $i \in \hat{L}, j \in L\setminus\hat{L}$, and $u \in [n_k t, n_k(t+\delta)]$. Since queues in L are given higher priority by LQF than are those in $K\setminus\hat{L}$ during $[n_k t, n_k(t+\delta)]$, from (2.4) and (2.5) we obtain

$$\frac{\hat{D}_{\hat{L}}^{n_k}(t+\delta) - \hat{D}_{\hat{L}}^{n_k}(t)}{n_k\delta} \in \mathrm{Co}(M[\hat{L}]) .$$

By letting $k \to \infty$ and then $\delta \downarrow 0$, we see that $\dot{\hat{D}}_{\hat{L}}^\infty(t) \in \mathrm{Co}(M[\hat{L}])$.

Now, since local pooling holds for $\hat{L}$ and λ is feasible, according to Remark 2.2.2 there exists $i_0 \in \hat{L}$ with $\lambda_{i_0} - \dot{\hat{D}}_{i_0}^\infty(t) \le -\epsilon_* < 0$. However, $\dot{\hat{Q}}_i^\infty(t) = \dot{\hat{Q}}_{i_0}^\infty(t)$ for all $i \in \hat{L}$, so $d\max_j \hat{Q}_j^\infty(t)/dt \le -\epsilon_* < 0$ for all $t \in (0,1)$ except some of Lebesgue measure 0. □

2.3.7 Rate of decrease of fluid limit

Pick a subsequence, (n_k^1) of (n_k^0) such that the convergence in Lemma 2.3.4 holds almost surely. From now on we fix ω such that that limit holds and such that (2.16) holds. Now consider a subsequence (n_k) of (n_k^1) along which the fluid limit $(\bar{Q}(\cdot), \bar{T}(\cdot), \bar{D}(\cdot))$ exists. Since $\bar{Q}(\cdot)$ is absolutely continuous, consider any regular time $t \in \mathbb{Q}_+$ for which $\bar{Q}(t) \neq 0$ and the derivative of $\max_i \bar{Q}_i(\cdot)$ exists.

The following result holds.

Lemma 2.3.11 (Rate of decrease of fluid limit)**.** *If $\bar{Q}(t) \neq 0$ then*

$$\frac{d}{dt}\max_{i\in K}\bar{Q}_i(t) \le -\epsilon_* ,$$

where ϵ_* *is as defined in* (2.8).

Proof. We proceed using contradiction. Assume that

$$\frac{d}{dt}\max_{i\in K}\bar{Q}_i(t) > -\epsilon_* \ .$$

Then, for some $\delta \in \mathbb{Q}\cap(0,1)$ and all sufficiently large k,

$$\max_{i\in K} Q_i^{n_k}(n_k(t+\delta)) - \max_{i\in K} Q_i^{n_k}(n_k t) > -\delta\epsilon_* n_k \ . \tag{2.20}$$

Divide $[n_k t, n_k(t+\delta)]$ into intervals of size $\Delta(n_k)$, i.e.

$$I_j^{n_k} = [n_k t + (j-1)\Delta(n_k), n_k t + j\Delta(n_k)]$$

with $j \in \{1, 2, \ldots, \delta n_k/\Delta(n_k)\}$.

Since $\bar{Q}(t) \neq 0$, there exists nonempty set $K' \subset K$ such that K' dominates $K \setminus K'$ during $\{n_k t, \cdots, n_k(t+\delta)\}$ and $Q_i^{n_k}(n_k t) > n_k\delta$ for all $i \in K'$ by an appropriate choice of δ, noting that ω satisfies (2.16) and that the departure rate is bounded for each queue. From Lemma 2.3.8, the number of intervals not dominated by any subset of queues that satisfies local pooling does not exceed $g(n_k) + \delta_0 n_k/\Delta(n_k)$ as $k \to \infty$.

Hence, by (2.20), there exists an interval $I_{j_k}^{n_k}$ dominated by some L that satisfies local pooling, and

$$\max_{i\in L} Q_i^{n_k}(n_k t + j_k\Delta(n_k)) - \max_{i\in L} Q_i^{n_k}(n_k t + (j_k-1)\Delta(n_k)) > -\Delta(n_k)\epsilon_* \ ,$$

for all sufficiently large k. Otherwise, for arbitrarily large k,

$$\begin{aligned}\max_{i\in L} Q_i^{n_k}(n_k(t+\delta)) - \max_{i\in L} Q_i^{n_k}(n_k t) \leq \\ -\epsilon_*\Delta(n_k)\left[\frac{n_k\delta}{\Delta(n_k)} - \left(g(n_k) + \frac{\delta_0 n_k}{\Delta(n_k)}\right)\right] + \left(g(n_k) + \frac{\delta_0 n_k}{\Delta(n_k)}\right)\gamma\Delta(n_k) \ ,\end{aligned}$$

for some fixed $\gamma > 0$, which contradicts (2.20) for sufficiently small δ_0.

Since $I_{j_k}^{n_k} \subset [n_k t, n_k(t+\delta)]$ for all large k, there exist further subsequences (n_k') of (n_k) and (j_k') of (j_k) such that $(n_k' t + (j_k'-1)\Delta(n_k'))/n_k' \to \infty$, the order of elements of L implied by the ordering of coordinates in $Q_L^{n'_k}(n'_k t + (j'_k - 1)\Delta(n'_k))$ is fixed for all sufficiently large k, and the limit of $\hat{Q}^{n'_k}(\cdot)$ in $[0,1]$ as $k \to \infty$ exists and satisfies

$$\max_{i \in L} \hat{Q}_i^{\infty}(1) - \max_{i \in L} \hat{Q}_i^{\infty}(0) > -\epsilon_* \,.$$

Since L satisfies local pooling, by Lemma 2.3.10 we arrive at a contradiction, proving Lemma 2.3.11. □

2.3.8 Proof of Theorem 2.3.1

Fix any sequence $(n_k^0) \subset \mathbb{Z}_+$ with $n_k^0 \to \infty$ as $k \to \infty$.

From Lemma 2.3.11 and the condition on $\bar{Q}(0)$ in Proposition 2, there exists $t_0 \geq 0$ such that $\bar{Q}(t) = 0$ for all $t \geq t_0$, and since (n_k) is an arbitrary subsequence of (n_k^1) so long as it satisfies the prescribed properties, $Q^{n_k^1}(n_k^1 t_0)/n_k^1 \to 0$ almost surely as $k \to \infty$. After some uniform integrability analysis for $(Q^{n_k^1}(n_k^1 t_0)/n_k^1)_k$ (see, e.g. [14]), we find that

$$\frac{1}{n_k} E\left[\sum_{i \in K} Q_i^{n_k^1}(n_k^1 t_0)\right] \to 0 \text{ as } k \to \infty \,.$$

Since (n_k^1) could have been any arbitrary subsequence of (n_k^0) with the prescribed properties, the limit over (n_k^0) is also 0.

Remark 2.3.12. The same proof can be used (virtually unchanged) to show that a system that satisfies local pooling is stable under i.i.d. (possibly constant) arrivals. For systems $M[K]$ that form trees, in the graph notation of section 2.1 it is easy to see that local pooling holds: at all times either a leaf or its parent vertex is served.

2.4 Extension to general service rates

The model can be generalized to all nonzero nonnegative matrices $M[K]$ in a straighforward manner. We have not considered that level of generality, overcomplicating the notation.

For nonintegral $M[K]$ we must redefine LQF appropriately, since the definition given above depends on the fact that elements of $M[K]$ can take values in $\{0, 1\}$ only. A stationary scheduling policy assigns service vectors, i.e. columns of $M[K]$, to queue lengths, i.e. vectors in $\mathbb{R}_+^K$. LQF will do this by "prioritizing" queues according to their backlog size. The allocation of service rates (more general than 0 or 1) should be based on these priorities. Such considerations motivate the following definition: For a service matrix $M[K]$, let $\mathcal{S}$ be the set of possible service vectors, i.e. columns of $M[K]$. A stationary policy $\Phi : \mathbb{R}_+^K \to \mathcal{S}$ is an LQF policy for $M[K]$ if $\Phi = \Phi_2 \circ \Phi_1$. Here Φ_1 is a (possibly random) mapping to the space of permutations of K that reorders the coordinates of its argument in decreasing order: if $\Phi_1(q) = \pi$, then $q_{\pi(1)} \geq q_{\pi(2)} \geq \ldots$. For a permutation π, $\Phi_2(\pi)$ is obtained as follows. First π reorders the rows of $M[K]$. Let $M^{\pi\cdot}$ be the resulting matrix. The column $j(\pi)$ indexed by Φ_2 is the (unique) element of $\mathcal{S}$ such that the vector $(M^{\pi\cdot}_{1j(\pi)}, M^{\pi\cdot}_{2j(\pi)}, \ldots, M^{\pi\cdot}_{kj(\pi)})$ is maximal across all columns, for all k with $1 \leq k \leq K$.

Our results carry over to this more general rule (although it is not included in the proofs), because the essential property of LQF we make use of is that shorter queues do not affect longer ones; in each time slot the service rate of a (nonzero) queue does not depend on the queue length of any shorter queues.

Chapter 3

Approximating Queues in Slowly Varying Stationary Environments

The main results of this chapter are Theorems 3.1.1 and 3.2.1. Theorem 3.1.1 in section 3.1 provides an approximation framework for general Markov chains. In section 3.2 we consider Jackson networks with varying arrival and service rates and fixed routing, and apply the framework developed in 3.1 This is demonstrated by Theorem 3.2.1.

3.1 Approximation of Markov chains

Consider a family of continuous-time Markov chains with countable state space $\mathcal{Y}$, given by the transition matrices $\{Q^x, x \in \mathcal{X}\}$, where $\mathcal{X}$ is a countable set. Assume each Q^x is positive recurrent and let π^x be the corresponding invariant distribution.

Assume a family of probability spaces $((\Omega, \mathcal{F}, P_\epsilon), \epsilon > 0)$ exists, equipped with a measure preserving flow $\theta = (\theta^t, t \in \mathbb{R})$.

On $(\Omega, \mathcal{F})$, the environment process $X = (X_t, t \in \mathbb{R})$ is given, compatible with flow θ, taking values in $\mathcal{X}$, with right-continuous piecewise-constant paths P_ϵ-a.s. for all $\epsilon > 0$. (For more details on flow compatibility and the stationary framework for point processes, the reader is referred to [5].) Assume the jump times $-\infty < \ldots < T_{-1} < T_0 \leq 0 < T_1 <$

$\ldots < +\infty$ of X_t, form a simple point process. Under P_1, these jumps have finite intensity λ, and under P_ϵ, the jump times are dilated by ϵ^{-1}, i.e.,

$$P_\epsilon((T_{-n}, \ldots, T_n) \in A_{-n} \times \ldots \times A_n) = P_1((\epsilon^{-1}T_{-n}, \ldots, \epsilon^{-1}T_n) \in A_{-n} \times \ldots \times A_n) ,$$

for each $\epsilon > 0, n = 0, 1, \ldots$, and Borel sets $(A_m : m = -n, \ldots, n)$. The sequence of values right after the jumps is not affected by ϵ, i.e., $P_\epsilon(X_{T_m} = x_m, m = -n, \ldots, n)$ does not depend on ϵ. Let P_ϵ^0 be the Palm distribution (w.r.t. the jumps of X_t) corresponding to P_ϵ. Also, by $\mathcal{F}^X$ we denote the σ-field generated by X.

On the same space, a process $Y = (Y_t, t \in \mathbb{R})$ is given, compatible with θ, for which under P_ϵ, conditionally on X it is a Markov chain with generator Q^{X_t}. That is, for any $f : \mathcal{Y} \to \mathbb{R}$ with $E_\epsilon|f(Y_u)| < \infty$ for all u,

$$t \mapsto f(Y_t) - f(Y_s) - \int_s^t Q^{X_u} f(Y_{u-}) du$$

is a martingale for all $-\infty < s \leq t < \infty$, w.r.t. the filtration generated by $\mathcal{F}^X$ and $(Y_u : u \leq t)$ for all t, and w.r.t. P_ϵ.

We will use multiple versions of Q^{X_t} Markov chains started from different initial states. These couple when they hit the same state and evolve according to $g : \mathcal{X} \times \mathcal{Y} \times \mathbb{R}_+ \times \Omega \to \mathcal{Y}$, defined such that

1. $g(x, y, 0, \omega) = y\ , \quad g(x, g(x, y, t, \omega), s, \theta^t\omega) = g(x, y, t + s, \omega)\ , \quad \forall s, t \geq 0,$
2. $(t, \omega) \mapsto g(x, y, t, \omega)$ is a measurable stochastic process for all x, y.
3. $\mathcal{F}^X$ is independent of $\mathcal{G}$, the σ-field generated by $(g(x, y, t, \omega), x \in \mathcal{X}, y \in \mathcal{Y}, t \geq 0)$. (Intuitively, ω is "noise" independent of X.)
4. For all x, y, the process $(t, \omega) \mapsto g(x, y, t, \omega)$ is a Markov chain with generator Q^x.

We will compare Y with another process Z which corresponds to "$\epsilon = 0$". Before we give the definion of Z, let us informally explain the idea behind this since it plays an important role. At the jumps of X, e.g. at T_0, Z is independently started according to the invariant distribution that corresponds to the new state of the environment, i.e. $\pi^{X_{T_0}}$. Between times

T_0 and T_1, Z evolves according to the transition matrix $Q^{X_{T_0}}$. Hence, Z behaves as Y with the difference that equilibrium is reached "instantaneously" by the former. Notice that the marginal distribution of Z at stationarity, is the one which Y is expected to converge to, as $\epsilon \downarrow 0$. If Z and Y are defined in an appropriate way then they may couple before T_1. Thus, the difference $E_\delta(f(Y_0) - f(Z_0))$ measures the "distance" of the marginal distributions of Y, between cases $\epsilon = \delta$ and $\epsilon = 0$.

We now define Z. Let $V_t = V_0 \circ \theta^t \in \mathcal{Y}$ be a right-continuous piecewise-constant process whose jumps are those of X, where the r.v's $V_{T_i}, i \in \mathbb{Z}$ are conditionally independent and each distributed according to $\pi^{X_{T_i}}$, given $\mathcal{F}^X$. Let $\mathcal{F}^V$ be the σ-field generated by $V = (V_t, t \in \mathbb{R})$. Define,

$$Z_t(\omega) = g(X_{T_0}, V_{T_0}, -T_0, \omega) \circ \theta^t \,, \quad t \in \mathbb{R} \,.$$

Now, consider the following assumptions:

(A1) $(P^0_\epsilon(Y_0 \in \cdot), \epsilon > 0)$ is tight.

(A2) $\inf\{t > 0 | g(X_0, Y_0, t, \omega) = g(X_0, V_0, t, \omega)\} < \infty \,, \quad P^0_\epsilon$-a.s., $\forall \epsilon > 0$.

(A3) $E_\epsilon |f(Y_0)| < \infty$, $E_\epsilon |f(V_0)| < \infty$.

(A4) There exists r.v. B for which

$$\int_0^\infty |f(g(X_0, Y_0, t, \omega)) - f(g(X_0, V_0, t, \omega))| dt \le B \quad \text{and} \ \limsup_{\epsilon \downarrow 0} E^0_\epsilon(B) < \infty \,.$$

(A5)

$$E^0_1 \left[\int_0^\infty |f(g(X_0, W, t, \omega)) - f(g(X_0, V_0, t, \omega))| dt \right] < \infty \,,$$

where, conditionally on $X_{T_{-1}}$, $W \stackrel{d}{=} \pi^{X_{T_{-1}}}$ and independent of $\mathcal{F}^X \vee \mathcal{F}^V \vee \mathcal{G}$.

We will show the following:

Theorem 3.1.1. *Under (A1)-(A5),*

$$E_\epsilon(f(Y_0)) = f^{(0)} + \epsilon f^{(1)} + o(\epsilon)$$

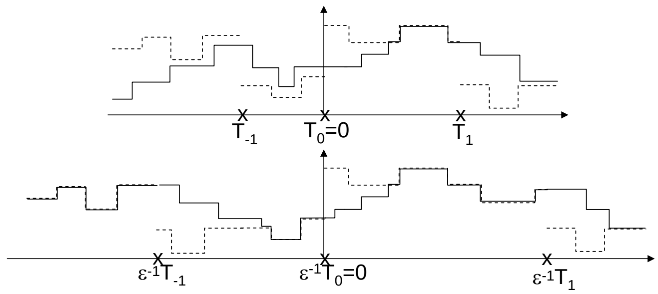

Figure 3.1. The argument behind the approximation in Theorem 3.1.1 based on comparison of sample paths of Y (solid) and Z (dashed line). In the lower part, the jump times of X are dilated by ϵ^{-1}.

where

$$f^{(0)} = E_1(f(Z_0)) = \sum_{x\in\mathcal{X}} P_1(X=x) \sum_{y\in\mathcal{Y}} \pi^x(y) f(y)$$
$$f^{(1)} = \lambda E_1^0 \left(\int_0^\infty [f(g(X,W,t,\omega)) - f(g(X,V,t,\omega))]dt \right) .$$

Remark 3.1.2. The term $f^{(0)}$ can be computed in terms of the deviation matrix of a Markov chain [12], which is given as a solution to a matrix equation. In certain cases, most notably the M/M/1 queue, explicit formulas exist [27] (see the example following Theorem 3.2.1).

Remark 3.1.3. Functions $f : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$ can be handled as well, without changing the proofs.

Remark 3.1.4. In the case that (X, Y) is a Markov chain, (A1) is equivalent to tightness of the time-stationary marginals. This fact is shown in section 3.3.

Remark 3.1.5. The idea behind the approximation in Theorem 3.1.1 is the following. Assume that (X, Y) are ergodic processes even though this is not needed in the proof. Compute the "finite difference" $(E_\epsilon(f(Y_0)) - E_\epsilon(f(Z_0)))/\epsilon$ using the mean-cycle formula over the jumps of X. The difference is equal to the difference accumulated during a typical cycle, i.e.

$$\frac{E_\epsilon(f(Y_0)) - E_\epsilon(f(Z_0))}{\epsilon} = \lambda E_\epsilon^0 \left(\int_0^{T_1} [f(Y_t) - f(Z_t)]dt \right) , \tag{3.1}$$

This is a function of the sample paths of Y and Z between times T_0 and T_1. Now, consider dilations of the jump times of X by ϵ^{-1}. This is depicted in Figure 3.1. The assumption on tightness, (A1), ensures that Y_t and Z_t are not far apart at the beginning of these cycles. So, in Figure 3.1 at T_{-1}, the difference of Y_t and Z_t does not grow as ϵ is decreased. If $\epsilon > 0$ is sufficiently small then Y_t, and Z_t have ample time to couple by time 0, as it happens in the lower part of the figure. Indeed, (A2) will guarantee that coupling occurs during sufficiently long cycles. Hence Y_0 is approximately distributed as $\pi^{X_{0-}}$. Now, since $\epsilon^{-1}T_1 \to +\infty$ P_1^0-a.s., we expect the RHS of (3.1) to converge to $f^{(1)}$. This is valid as long as the limit can be exchanged with the expectation. This is justified by dominated convergence, which applies because of assumptions (A4) and (A5).

Proof. By (A3), the Palm inversion formula (see [5]) gives,

$$\frac{E_\epsilon(f(Y_0) - f(Z_0))}{\epsilon} = \lambda E_\epsilon^0 \left[\int_0^{T_1} (f(Y_t) - f(Z_t))dt \right] .$$

Since $E_\epsilon(f(Z_0)) = f^{(0)}$, to prove the theorem it suffices to show that the limit, as $\epsilon \downarrow 0$, of the RHS is $f^{(1)}$.

First we show $P_\epsilon^0[Y_0 \in \cdot | X_{0-}, X_0] \to \pi^{X_{0-}}(\cdot)$ as $\epsilon \downarrow 0$, P_1^0-a.s. Since $P_\epsilon^0 \circ (X_{-0}, X_0)^{-1}$ does not depend on ϵ, tightness of $(P_\epsilon^0[(Y_0, T_1) \in \cdot | X_{-0} = x, X_0 = x'], \epsilon > 0)$ and $(P_\epsilon^0[(Y_0, T_1) \in \cdot | X_0 = x, X_{T_1} = x'], \epsilon > 0)$, for each x, x' with $P_\epsilon^0(X_0 = x, X_{T_1} = x') > 0$, follows. (These are distributions on $\mathbb{Z}_+ \times (\mathbb{R}_+ \cup \{\infty\})$, where $\mathbb{R}_+ \cup \{\infty\}$ is equiped with the Borel σ-field of the one-point compactification.) Since $\mathcal{X}$ is countable, by diagonalization we can extract a subsequence $\epsilon_k \downarrow 0$ such that $P_{\epsilon_k}^0[(Y_0, T_1) \in \cdot | X_{0-}, X_0] \to \nu^{X_{0-}, X_0}(\cdot) \times \delta_\infty$, as $k \to \infty$, P_1^0-a.s., where δ_∞ is the point-mass at ∞. We will show that $\nu^{X_{0-}, X_0} = \pi^{X_{0-}}$. Take any subsequence of (ϵ_k) (which we denote it again by (ϵ)) for which the limit $\mu^{X_0, X_{T_1}} \times \delta_\infty(\cdot)$ of $P_\epsilon^0[(Y_0, T_1) \in \cdot | X_0, X_{T_1}]$ converges P_1^0-a.s. . Now, take r.v's $((Y^\epsilon, T_1^\epsilon), \epsilon > 0)$ on P_1 (perhaps when enlarged), satisfying $P_1^0[(Y^\epsilon, T_1^\epsilon) \in \cdot | X_0, X_{T_1}] = P_\epsilon^0[(Y_0, T_1) \in \cdot | X_0, X_{T_1}]$ P_1^0-a.s., and whose P_1^0-a.s. limit (Y^0, ∞), as $\epsilon \downarrow 0$, is distributed according to $\mu^{X_0, X_{T_1}} \times \delta_\infty$. Moreover, $(Y^\epsilon, T_1^\epsilon)$ is taken to be conditionally independent of $\mathcal{F}^V \vee \mathcal{G}$ given X_0. Hence,

$$P_\epsilon^0[g(X_0, Y_0, T_1, \omega) \in \cdot | X_0, X_{T_1}] = P_1^0[g(X_0, Y^\epsilon, T_1^\epsilon, \omega) \in \cdot | X_0, X_{T_1}] . \tag{3.2}$$

Now, from (A2), $g(X_0, Y^0, T_1^\epsilon, \omega) = g(X_0, V_0, T_1^\epsilon, \omega)$, for small enough ϵ. Since $Y^\epsilon \to Y^0$, we have $g(X_0, Y^\epsilon, T_1^\epsilon, \omega) = g(X_0, Y^0, T_1^\epsilon, \omega)$ for small $\epsilon > 0$. Combining with (3.2) this yields,

$$P_\epsilon^0[g(X_0, Y_0, T_1, \omega) \in \cdot | X_0, X_{T_1}] - P_1^0[g(X_0, V_0, T_1^\epsilon, \omega) | X_0, X_{T_1}] \to 0\,, \quad P_1^0\text{-a.s.}\,,$$

so, the limit of the first term on the LHS is π^{X_0} by the conditional independence of V_0 and $\mathcal{G}$ given X_0. Notice that

$$P_\epsilon^0[g(X_0, Y_0, T_1, \omega) \in \cdot | X_0, X_{T_1}] \circ \theta^{T_1} = P_\epsilon^0[Y_0 \in \cdot | X_{0-}, X_0]\,,$$

so (3.2) implies $\nu^{X_{0-}, X_0} = \pi^{X_{0-}}$.

Redefine $(Y^\epsilon, T_1^\epsilon)$ such that it again has P_1^0-a.s. limit (Y^0, ∞), and $P_1^0[(Y^\epsilon, T_1^\epsilon) \in \cdot | X_{0-}, X_0] = P_\epsilon^0[(Y_0, T_1) \in \cdot | X_{0-}, X_0]$, and $(Y^\epsilon, T_1^\epsilon)$ is conditionally independent of $\mathcal{F}^V \vee \mathcal{G}$ given X_{0-}. Now by (A2),

$$\int_0^{T_1^\epsilon} [f(g(X_0, Y^\epsilon, t, \omega)) - f(g(X_0, V_0, t, \omega))]dt \underset{\epsilon \downarrow 0}{\longrightarrow} \int_0^\infty [f(g(X_0, Y^0, t, \omega)) - f(g(X_0, V_0, t, \omega))]dt < \infty\,, \quad P_1^0\text{-a.s.}\,.$$

(A4) implies that the sequence of r.v's on the LHS is uniformly integrable, so together with (A5) imply that the limit under E_1^0 exists and equals $f^{(1)}/\lambda$. □

3.2 Application to queues

Here, we show that Jackson networks can be put in the framework of the previous section.

Let $J \geq 1$ and X be as in the previous section. On state x of the environment, $Y = (Y_t(i), i = 1, \ldots, J, t \in \mathbb{R})$ is a J-station open Jackson network with external arrival rates $(\lambda(x)R_{01}, \ldots, \lambda(X)R_{0J})^T$ for some $(R_{0i}, i = 1, \ldots, J) =: R_0$. with $R_{0i} \geq 0, \sum_{i=1}^J R_{0i} = 1$, service rates $\mu(x) := (\mu_1(x), \ldots, \mu_J(x))$, and a $J \times J$ routing matrix $R = (R_{ij} | i, j \in$

$\{1,\ldots,J\}$). In other words, for any $f:\mathbb{Z}_+^J \to \mathbb{R}$, the generator Q^x is given by

$$
\begin{aligned}
Q^x f(y) = & \sum_{i=1}^{J} \lambda(x) R_{0i}(f(y+e_i) - f(y)) \\
& + \sum_{i=1}^{J}\sum_{j=1}^{J} \mu_i(x) 1(y_i > 0) R_{ij}(f(y - e_i + e_j) - f(y)) \\
& + \sum_{i=1}^{J} 1(y_i > 0)\mu_i(x) R_{i0}(f(y-e_i) - f(y)) \,,
\end{aligned} \tag{3.3}
$$

where R is a stochastic matrix such that $I - R$ is non-singular, and e_i is the vector with zero elements except a 1 in the i-th coordinate. The vector of arrival rates at each node is given by $a(x) = \lambda(x)(I - R^T)^{-1} R_{0\cdot}$.

We now define the processes $g(x, y, t, \omega)$. For any $i \in \{1,\ldots,J\}, j \in \{0,\ldots,J\}$, let N_{0i}^x, N_{ij}^x be Poisson processes with intensities $\lambda(x)R_{0i}$ and $\mu_i(x)R_{ij}$, respectively, such that $(N_{0i}^x, N_{ij}^x, i \in \{1,\ldots,J\}, j \in \{0,\ldots,J\})$ are independent. Define $g(x, y, t, \omega)$ to be the solution of the SDE

$$
U_i(t) = y_i + N_{0i}^x(t) + \sum_{j=1}^{J} \int_0^t 1\{U_j(s-) > 0\} N_{ji}^x(ds) - \sum_{j=0}^{J} \int_0^t 1\{U_i(s-) > 0\} N_{ij}^x(ds) \,,
$$

$i = 1,\ldots,J$, w.r.t. $(U(t), t \geq 0)$.

We will consider the assumptions

(B1) $\tilde{\rho}_i := \sup_x a_i(x)/\mu_i(x) < 1$ for all i, x,

(B2) $\inf_x \mu_i(x) > 0$ for all i, and $\inf_x \lambda(x) > 0$,

and show the following

Theorem 3.2.1. *Suppose (B1)-(B2) hold. Then,*

(i) (X, Y) possesses a stationary law, for each $\epsilon > 0$.

(ii) (A1)-(A5) are satisfied, and therefore $P_\epsilon(Y_0 = y)$ can be approximated using Theorem 3.1.1.

(Notice that in part (ii), we invoke Theorem 3.1.1 for $f(Y_0) = 1\{Y_0 = y\}$, where $y \in \mathcal{Y}$.)

Remark 3.2.2. Assumption (B2) regarding the service rates can be weakened by putting restrictions on X. What is necessary is that equations (3.4)-(3.6) below possess a unique solution.

Example 3.2.3 (M/M/1 queue). Here, $J = 1$ and $\pi^x(y) = \rho(x)(1-\rho(x))^y$, $y = 0, 1, \ldots$. Using the formula in [27] for the deviation matrix of an M/M/1 queue, we have

$$P_\epsilon(Y_0 = y) = \sum_x P_1(X_0 = x)\pi^x(y) + \epsilon\lambda \sum_x P_1^0(X_{0-} = x, X_0 = x') \cdot \sum_{i=1}^{\infty} \pi^x(i) \frac{\rho(x')^{(y-i)^+} - (i+y-1)\pi^{x'}(y)}{\mu(x') - \lambda(x')} + o(\epsilon)\,, \quad \text{as } \epsilon \downarrow 0\,.$$

In sections 3.2.2 and 3.2.3, parts (i) and (ii) are proved, respectively. Specifically, in section 3.2.2 we construct a stationary version of (X, Y). To do this we compare Y with the queue lengths $\tilde{Y}$ of another Jackson network with service speeds $\tilde{\mu}(x) \le \mu(x)$ for all x. $\tilde{Y}$ is constructed by applying a time-change to a non-varying Jackson network. The existence of a stationary version for the latter is established in [7]. In section 3.2.3, again key is the comparison with $\tilde{Y}$. The fact that the marginal distribution of $\tilde{Y}$ does not depend on ϵ and the monotonicity of $g(x, y, t, \omega)$ in $y \in \mathcal{Y}$, reduce the proofs of (A4) and (A5) into showing that the mean residual busy period of some non-varying Jackson network is finite. This fact follows from the fact that in Jackson networks, recurrence times possess finite moments of all orders (see [19] and Theorem 15.0.1 in [34]).

In the next section, we review the concept of an Euler network from Bacceli and Foss [7] modified to allow for variable service speeds. This is used for establishing the monotonicity properties in sections 3.2.2 and 3.2.3.

3.2.1 Preliminaries

We follow the construction in [7] allowing for variable service speeds. A tuple $\Sigma = (N, t, \sigma, R, \mu)$ is an *Euler network* when

- $t = (t(1), \ldots, t(N)), N \ge 1$ is a sequence of non-decreasing real numbers.
- $\sigma = (\sigma_i(n) \in \mathbb{R}_+,\ i = 1, \ldots, J, n = 0, \ldots, d_i)$, for some $d_i \ge 0$. We call σ, the *service*

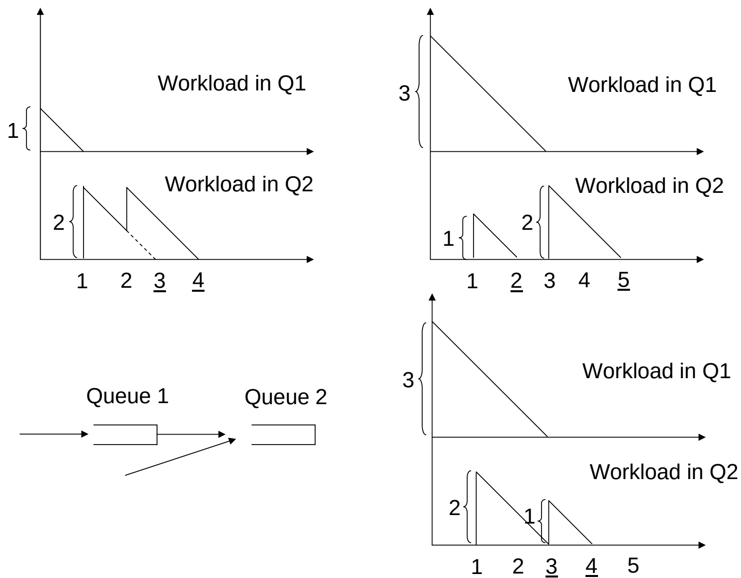

Figure 3.2. An example motivating the association of service sizes to stations instead of customers. Doing so, yields monotonicity properties not present in the customer-centric case.

sequence, and $\sigma_i(n)$ is interpreted as the job size of the n-th customer served by station $i \geq 1$.

- $r = (r_i(n) \in \{0, \ldots, J\},\ i = 0, \ldots, J, n = 0, \ldots, d_i)$ for the d_i's above. r is called the *switching sequence*, and $r_i(n)$ is interpreted as the station that the n-th serviced customer from station i jumps to.

- $\mu = (\mu_i(t) \in \mathbb{R}_+,\ i \in \{1, \ldots, J\}, t \geq 0)$ with $\inf_t \mu_i(t) > 0$. $\mu_i(t)$ is interpreted as the instantaneous server speed at the i-th station.

- Under the above interpretations of t, σ, r, μ, the number of arrivals and departures from station i both equal d_i.

Example 3.2.4. To motivate this construction, consider two tandem queues as depicted in the lower left part of Figure 3.2.1. There are two customers ($N = 2$) arriving in the system. The first customer arrives at $t(1) = 0$ at queue 1 and the second at $t(2) = 2$ at queue 2.

Hence, $d_1 = 1$, $d_2 = 2$, $r_0(1) = 1$, $r_0(2) = r_1(1) = 2$, and $r_2(1) = r_2(2) = 0$. The service sequence is given by $\sigma_1(1) = \sigma_2(2) = 1$ and $\sigma_2(1) = 2$. Let μ_i be the speed of the server at queue i and first consider $\mu_1 = \mu_2 = 1$. At the top left part of Figure 3.2.1, we plot the workload (i.e. the time until emptiness) for each queue as a function of time. On the time-axis, the times of departures from the system are underlined.

In this and the following sections we seek to compare the total number of customers in systems which differ only in their service speeds, e.g. see Lemma 3.2.5 . It turns out that the number of customers in faster systems is dominated by that in slower ones. This is not evident under a more "standard" construction where service times are associated to customers and not to stations as above. Since the first customer is also the first customer served by station 2, we should have $\sigma_1(1) = 1$, and $\sigma_2(1) = 2$ associated with this customer. For the same reason, we associate $\sigma_2(2) = 1$ with the second customer (i.e. the one arriving directly at queue 2). At the top rightmost part of Figure 3.2.1, we plot the resulting workloads when $\mu_1 = 1/3$, and $\mu_2 = 1$. Notice that between times 2 and 3, the number of customers in this system is *dominated* by that in the system with service speeds $\mu_1 = 1$, and $\mu_2 = 1$ (plot on the left hand side).

However, if service sizes are associated to stations as in the definition of an Euler network, the number of customers is never below that in the faster case. This is depicted in the leftmost bottom part of the figure.

By Theorem 8 in [7], for such Σ, the number of arrivals and departures from station i is determined by r, N alone.

Given Σ, the set of departure times $D = (D_j(n), j = 1, \ldots, J, n = 1, \ldots, d_j)$ is determined by the equations

$$D_0(n) = t_n \,, \quad 1 \le n \le N \,, \tag{3.4}$$

$$\int_{S_j(n)}^{D_j(n)} \mu_j(t) ds = \sigma_j(n) \,, \quad j = 1, \ldots, J \,, \tag{3.5}$$

$$S_j(n) := \max(D_j(n-1), \min_{n_0+\ldots+n_k=n} (\max_{i=0,\ldots,J} D_i(\eta_{i,j}(n_i))) \,, \quad j = 0, \ldots, J \,, \tag{3.6}$$

where

$$\eta_{i,j}(n) = \inf\{d_i \geq m \geq 1 | \sum_{p=1}^{m} 1\{r_i(p) = j\} = n\}\ , \quad i, j = 0, \ldots, J\ ,$$

is the minimum number of service completions at station i such that n of them are routed to j. The assumption $\inf_t \mu_i(t) > 0$ was made so that (3.4)-(3.6) always have a unique solution. (For more on evaluating (3.4)-(3.6) inductively, see [6].) We will sometimes write $d_j(\Sigma), D(\Sigma) = (D_j(\Sigma, n), j = 1, \ldots, J, n = 1, \ldots, d_j(\Sigma)), S_j(\Sigma, n)$ to emphasize the dependence on Σ.

Define the time to empty after the last arrival $\max_{i=1,\ldots,J} \max_{n=1,\ldots,d_i} D_i(n) - t_N =: Z(\Sigma)$, and the size of queue j at time s by

$$Q_j(\Sigma, s) = \sum_{i=0}^{J} \sum_{n=1}^{d_i} 1\{D_i(n) < s\ ,\ r_i(n) = j\} - \sum_{n=1}^{d_j} 1\{D_j(n) < s\}\ .$$

Lemma 3.2.5. *Let* $\Sigma^k = (N, t^k, \sigma, r, \mu^k), k = 1, 2$ *be two Euler networks with* $\mu^1 \geq \mu^2$, $t^1 \leq t^2$. *Then,* $D(\Sigma^1) \leq D(\Sigma^2)$.

Proof. Follows directly by applying induction to (3.4)-(3.6). □

Before obtaining the important monotonicity property of Lemma 3.2.6, we need the concept of composition from [7]. The composition $\Sigma = (N, t, \sigma, r, \mu) = \Sigma^1 + \Sigma^2$, where $\Sigma^k = (N^k, t^k, \sigma^k, r^k, \mu), k = 1, 2$ satisfying $t^1(N^1) \leq t^2(1)$, is defined as follows:

$$N = N^1 + N^2$$

$$t = (t^1(1), \ldots, t^1(N^1), t^2(1), \ldots, t^2(N^2))$$

$$r_i(n) = \begin{cases} r_i^1(n) & 1 \leq n \leq d_i(\Sigma^1) \\ r_i^2(n - d_i(\Sigma^1)) & d_i(\Sigma^1) < i \leq d_i(\Sigma^1) + d_i(\Sigma^2) \end{cases}$$

$$\sigma_i(n) = \begin{cases} \sigma_i^1(n) & 1 \leq n \leq d_i(\Sigma^1) \\ \sigma_i^2(n - d_i(\Sigma^1)) & d_i(\Sigma^1) < i \leq d_i(\Sigma^1) + d_i(\Sigma^2) \end{cases}$$

By Theorem 9 in [7], $d_i(\Sigma) = d_i(\Sigma^1) + d_i(\Sigma^2)$ which follows from the fact that the number of arrival and departures equals $d_i(\Sigma^\cdot)$, regardless of the timing information in t, σ. This fact also implies that composition is associative.

Lemma 3.2.6. *$Z(\Sigma^1 + \Sigma^2) \geq Z(\Sigma^2)$.*

Proof. Define $\Sigma^1(\Delta) = (N^1, (t^1(n) - \Delta, n = 1, \ldots, N^1), \sigma^1, r^1, \mu)$, for $\Delta > 0$. By (3.5),

$$D_i(\Sigma^1(\Delta), n) \leq \frac{\sigma_i^1(n)}{\inf_t \mu_i(t)} + S_i(\Sigma^1(\Delta), n) ,$$

for n, i such that all quantities are defined. Since no more than $\sum_{k=1}^J d_k(\Sigma^1)$ induction steps of (3.5)-(3.6) are needed in evaluating $S_i(\Sigma^1(\Delta), n)$, we get

$$D_i(\Sigma^1(\Delta), n) - (t^1(N^1) - \Delta) \leq \sum_{k=1}^J d_k(\Sigma^1) \max_{j=1,\ldots,J} \frac{\sigma_j^1(n)}{\inf_t \mu_j(t)} ,$$

with the LHS not depending on Δ. Since $\inf_t \mu_i(t) > 0$ for all i, we have $Z(\Sigma^1(\Delta)) + t^1(N^1) - \Delta < t^2(1)$ for sufficiently large Δ. For such Δ, $Z(\Sigma^1(\Delta) + \Sigma^2) = Z(\Sigma^2)$, so by Lemma 3.2.5, $Z(\Sigma^2) \leq Z(\Sigma^1 + \Sigma^2)$ □

We will need one more definition; that of a time-changed Euler network. Let $\tau : \mathbb{R} \to \mathbb{R}$ be a strictly increasing bijection. For an Euler network $\Sigma = (N, t, \sigma, r, \mu)$, define $\tau\Sigma := (N, \tau(t), \sigma, r, \nu)$, where $\nu(s) = \mu(\tau^{-1}(s))/\tau'(\tau^{-1}(s))$ for all $s \in \mathbb{R}$. (Notice that ν is defined Lebesgue-a.e. .) It is immediate from (3.4)-(3.6) that the following holds.

Lemma 3.2.7. *$D(\tau\Sigma) = \tau(D(\Sigma))$, $Q(\Sigma, s) = Q(\tau\Sigma, \tau(s))$ for all $s \in \mathbb{R}$.*

3.2.2 Proof of Theorem 3.2.1(i)

Let $(t(k), k \in \mathbb{Z})$ be a Poisson process with stochastic intensity $\lambda(X_t)$ conditional on $\mathcal{F}^X$. For each $k \in \mathbb{Z}$, let r^k be a switching sequence generated by a random path in the Jackson network, with probabilities given by the routing matrix (R_{ij}). Moreover, for each such r^k, let σ^k be the associated service sequence, such that the $\sigma_j^k(n)$'s are i.i.d. exponential r.v's with mean one and independent of everything else. Let $\Sigma^k := (1, \{t(k)\}, \sigma^k, r^k, \mu(X.))$, $\tilde{\Sigma}^k := (1, \{t(k)\}, \sigma^k, r^k, \tilde{\mu}(X.))$ for $k \in \mathbb{Z}$. For $k, m \in \mathbb{Z} \cup \{-\infty, +\infty\}$, with $m \geq k$, we write $\Sigma_k^m := \Sigma^k + \ldots + \Sigma^m$ and similarly for $\tilde{\Sigma}_k^m$. Also, consider the time-change

$$\mathbb{R} \ni s \mapsto \tau(s) = \inf \left\{ v \in \mathbb{R} | s = \int_0^v 1/\lambda(X_u) du \right\} .$$

By (B2) this is bijective, so Lemma 3.2.7 can be applied.

Since $\tilde{\mu}(X_s) \le \mu(X_s)$ for all $s \in \mathbb{R}$, by Lemma 3.2.5 we have $D(\tilde{\Sigma}^0_k) \ge D(\Sigma^0_k)$ for all $k \le 0$. This and Lemma 3.2.6 yields $Z(\tilde{\Sigma}^0_{-\infty}) \ge Z(\Sigma^0_{-\infty})$. Now, let $N_s, s \in \mathbb{R}$ be the counting process of points $(t(k), k \in \mathbb{Z})$, so e.g. $N_s = \sum_{k\in\mathbb{Z}} 1\{s \le t(k) < 0\}$ for $s \le 0$. Then, say if $s_2, s_1 \in \mathbb{R}$, $s_2 \le s_1 \le 0$,

$$E[N_{\tau^{-1}(s_2)} | \mathcal{F}^N_{\tau^{-1}(s_1)} \vee \mathcal{F}^X] = E\left[\int_{s_1}^{\tau^{-1}(s_2)} \lambda(X_u)dt \middle| \mathcal{F}^N_{\tau^{-1}(s_1)} \vee \mathcal{F}^X \right] ,$$

where $\mathcal{F}^N_u$ is the filtration generated by histories of N from time $u \le 0$ to $+\infty$. But by the definition of τ,

$$\int_{s_1}^{\tau^{-1}(s_2)} \lambda(X_u)du = s_2 - s_1 , \quad \forall s_1, s_2, s_2 \le s_1 \le 0 .$$

Therefore, $(\tau(t(k)), k \in \mathbb{Z})$ are the jump times of a Poisson process with unit intensity. Notice also, that $\tilde{\mu}_i(\tau^{-1}(s))/\tau'(\tau^{-1}(s)) = \tilde{\rho}_i^{-1} e_i^T (I - R^T)^{-1} R^T_{0\cdot}$. Hence, $\tau\tilde{\Sigma}$ corresponds to a Jackson network with constant arrival and service rates. Notice that the arrival rate $e_i^T(I-R^T)^{-1}R^T_{0\cdot}$ at queue i, is less than the service rate, so $Z(\tau\tilde{\Sigma}^m_{-\infty}) \le \tau(t(m+1)) - \tau(t(m))$ for some $m < 0$ (see [7]). By Lemma 3.2.7 and 3.2.6, this implies $Z(\tilde{\Sigma}^m_{-\infty}) \le t(m+1) - t(m)$, i.e., the network empties before the $m+1$-st arrival. Therefore, $Z(\Sigma^m_{-\infty}) \le t(m+1) - t(m)$ and $Q(\Sigma^0_{-\infty}, 0) := Q(\Sigma^0_{m+1}, 0)$ is well-defined. Now, $Y(t) := Q(\Sigma^0_{-\infty}, 0) \circ \theta^t, t \in \mathbb{R}$ defines a (P, θ)-compatible process, for which it is straightforward to show that conditionally on $\mathcal{F}^X$, it is a Markov chain with generator given by (3.3) on $\{X_t = x\}$.

3.2.3 Proof of Theorem 3.2.1(ii)

Let m be as in the proof of part (i). Since $D(\Sigma^0_{m+1}) \le D(\tilde{\Sigma}^0_{m+1})$, we have

$$\sum_{i=1}^J Y_0(i) \le \sum_{i=1}^J Q_i(\tilde{\Sigma}^0_{m+1}, 0) . \tag{3.7}$$

But $Q_i(\tilde{\Sigma}^0_{m+1}, 0) = Q_i(\tau\tilde{\Sigma}^0_{m+1}, 0)$ is a geometric r.v. with mean $1/(1 - \tilde{\rho}_i)$ as the invariant queue lengths in a Jackson network. This implies that $(P^0_\epsilon(Y_0(i) \in \cdot), \epsilon > 0)$ is tight. Also (A3) is satisfied.

Let $y^2 := (y^2_1, \dots, y^2_J) \ge (y^1_1, \dots, y^1_J) =: y^1$ be vectors with nonnegative components, and consider the processes $t \mapsto (g(x, y^k, t, \omega), k = 1, 2)$ defined in the beginning of this section. Say $g(x, y^2, s, \omega) \geq g(x, y^1, s, \omega)$ for all $0 \leq s < t$, and at time t one of the two processes jumps. Then, if $g_j(x, y^2, t, \omega)$ jumps down, then so does $g_j(x, y^1, t, \omega)$ provided it is nonzero. If $g_j(x, y^1, t, \omega)$ jumps upwards, then so does $g_j(x, y^2, t, \omega)$ since, in the case of a non-external arrival, the upstream queue i will have $g_i(x, y^2, t-, \omega) \geq g_i(x, y^1, t-, \omega) > 0$. Hence, $g(x, y^2, t, \omega) \geq g(x, y^1, t, \omega)$ for all $t \geq 0$. This implies that (A2) holds.

It remains to show (A4), (A5). By the monotonicity above,

$$\int_0^\infty |1\{g(X_0, Y_0, t, \omega) = y\} - 1\{g(X_0, V_0, t, \omega) = y\}| dt$$
$$\leq \inf\{t > 0 | g(X_0, Y_0, t, \omega) = g(X_0, V_0, t, \omega) = 0\}$$
$$\leq \inf\{t > 0 | g(X_0, Y_0 \vee V_0, t, \omega) = 0\} .$$

By (3.7), and the definition of V_0, $\sum_i (Y_0(i) \vee V_0(i)) \leq \sum_{i=1}^J Q_i^k$, $k = 1, 2$ for some i.i.d. Q^1, Q^2 each distributed as the RHS in (3.7). Thus, the integral above is bounded by

$$\inf \left\{ t > 0 | g(X_0, \left(\sum_{i=1}^J Q_i^1 \vee Q_i^2 \right) \sum_{j=1}^J e_j, t, \omega) = 0 \right\} =: B .$$

Now, on the same probability space we can define independent geometric r.v.'s $(Z_i, i = 1, \ldots, J)$, each with mean $1/(1 - \hat{\rho}_i) > 0$, such that

$$\sum_i Q_i^1 \vee Q_i^2 \leq \wedge_{i=1}^J Z_i , \quad P_1^0\text{-a.s.} \tag{3.8}$$

on a set $\{\sum_i Q_i^1 \vee Q_i^2 > M\}$, where M is a finite constant. This is because the LHS of (3.8) has a geometric tail which can be dominated by the tail of the RHS for $\hat{\rho}_1, \ldots, \hat{\rho}_J$ close enough to 1. Thus,

$$E_1^0(B) \leq E_1^0(\inf\{t > 0 | g(X_0, M \sum_{j=1}^J e_j, t, \omega) = 0\})$$
$$+ E_1^0(\inf\{t > 0 | g(X_0, (Z_1, \ldots, Z_J)^T, t, \omega) = 0\}) .$$

To show that this is finite, it suffices to do it for the 2nd term which we compare with the time to empty for an alternative Jackson network, when started with the invariant distribution.

Let $\hat{\mu}_i(x) = a_i(x)/\hat{\rho}_i$ for all i, and $\hat{\Sigma}(k) = (1, t(k), \sigma^k, R^k, \hat{\mu}), k \in \mathbb{Z}$. Now $\tau\hat{\Sigma}$ corresponds to a Jackson network with unit arrival rate and service rate $\hat{\rho}_i e_i^T (I - R^T)^{-1} R_{0\cdot}^T$ at station i, so $Q(\hat{\Sigma}_{-\infty}^0, 0) = Q(\tau\hat{\Sigma}_{-\infty}^0, 0) \stackrel{d}{=} (Z_1, \ldots, Z_J)^T$. Let $n(t) = \max\{m \in \mathbb{Z} | m \leq t\}$. Now,

$$
\begin{aligned}
E_1^0(\inf\{t > 0 | g(X_0, (Z_1, \ldots, Z_J)^T, t, \omega) = 0\}) \\
&= E_1^0(\inf\{t > 0 | Q(\hat{\Sigma}_{-\infty}^0 + \Sigma_1^{n(t)}) = 0\}) \\
&\leq E_1^0(\inf\{t > 0 | Q(\hat{\Sigma}_{-\infty}^{n(t)}, t) = 0\}) \\
&\leq \frac{1}{\inf_x \lambda(x)} E_1^0(\tau(\inf\{t > 0 | Q(\hat{\Sigma}_{-\infty}^{n(t)}, t) = 0\})) \\
&\leq \frac{1}{\inf_x \lambda(x)} E_1^0(\inf\{t > 0 | Q(\tau\hat{\Sigma}_{-\infty}^{n(t)}, t) = 0\})) < \infty \,,
\end{aligned}
$$

since this is the mean time to hit the empty state, starting from the invariant distribution. This shows that (A4) is satified. Since $W \leq_{\text{st}} Q^1$, where W is as in Theorem 3.1.1, the same arguments show that (A5) holds as well.

3.3 An equivalent condition to (A1) for Markovian environments

We show the following statement: If X is a Markov chain, then (A1) is equivalent to tightness of $(P_\epsilon(Y_0 \in \cdot), \epsilon > 0)$.

Proof. Let $q(X_t)$ be the instantaneous rate of jumps for X at time t. Since the marginal distribution of X does not depend on the parameter ϵ, if $(P_\epsilon(Y_0 \in \cdot), \epsilon > 0)$ is tight, it follows that $(P_\epsilon[Y_0 \in \cdot | X_0 = x], \epsilon > 0)$ is tight.

For fixed $\alpha > 0$ and any x, let $F_x \subset \mathcal{X}$ be finite and such that $\inf_{\epsilon > 0} P_\epsilon[Y_0 \in F_x | X_0 = x] > 1 - \alpha$. Now, for fixed arbitrary $\beta > 0$ let D be a finite set in $\mathcal{X}$ such that

$$E_\epsilon[q(X_0) 1\{X_0 \in D\}] > (1 - \beta) E_\epsilon(q(X_0)) \,, \quad \text{for all } \epsilon > 0$$

(because of the assumption that the marginal distribution of X does not depend on $\epsilon > 0$), and define $F := \cup_{x \in D} F_x$. Then by Papangelou's formula (see [5]),

$$
\begin{aligned}
P^0_\epsilon(Y_0 \in F) &= \frac{E_\epsilon[1\{Y_0 \in F\} q(X_{0-})]}{E_\epsilon(q(X_{0-}))} \\
&\geq \frac{E_\epsilon[P_\epsilon[Y_0 \in F | X_{0-}] q(X_{0-}) 1\{X_{0-} \in D\}]}{E_\epsilon(q(X_0))} \\
&\geq (1-\alpha)(1-\beta)\,, \quad \text{for any } \epsilon > 0\,,
\end{aligned}
$$

where we have used the observation that X, Y have no simultaneous jumps (P_ϵ and P^0_ϵ a.s.).

Since α, β were arbitrary, tightness of the LHS follows. □

Chapter 4

Macroscopic Behavior of a Lossy Resource Shared by Real- and Non-real-time Flows

In section 4.1 we describe a stochastic model for real-time and non-real-time flows. In section 4.2 we establish that the model is stable, as long as the same system but comprised only of non-real-time flows, is stable. In section 4.3 we obtain a weak convergence result for a sequence of systems that grow large, and derive an ODE that characterizes the limit.

4.1 Markov model

We consider the case of a single processor sharing resource of capacity C which is shared by real- and non-real-time flows. If at any time instant, there are n flows in the system, each one receives an equal bandwidth share of C/n. Flow dynamics are modeled as follows (see Figure 4.1):

1. Real-time flows arrive as a Poisson process of rate λ_1. Each flow has a minimum bandwidth requirement equal to 1. Upon arrival, such a flow is accepted if the in-

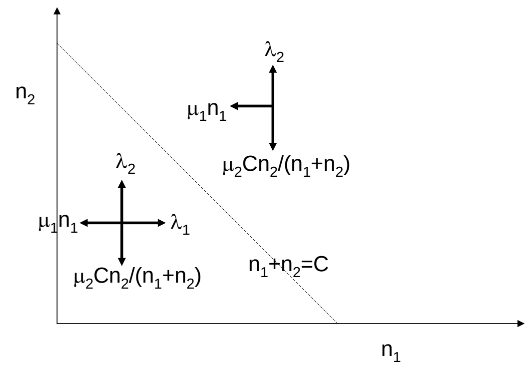

Figure 4.1. A Markov model for real-time, n_1, and non-real-time flows, n_2.

stantaneous bandwidth share is greater than 1; otherwise it is lost. Once accepted, it stays in the system for an exponential time with mean μ_1^{-1}.

2. Non-real-time flows arrive as a Poisson process of rate λ_2, independent of real-time ones. Each flow is associated with the transmission of a "file" of a random size, so that if served by one unit of bandwidth, its transfer will complete after an exponentially distributed time with mean μ_2^{-1}.

Let $n_1(t), n_2(t)$ denote the number of real- and non-real-time flows respectively, at time $t \geq 0$. Then $((n_1(t), n_2(t)); t \geq 0)$ is a Markov process describing the evolution in the number of flows. Admission decisions depend on the state only through $n_1(t) + n_2(t)$. In particular, we assume that a real-time flow is accepted if $n_1(t) + n_2(t) < C$ and rejected otherwise. Non-real-time flows are always accepted.

According to the above description, if the system is in state (n_1, n_2), transitions occur to state

$(n_1 + 1, n_2)$ with rate $\lambda_1 I_{\{n_1+n_2<C\}}$

$(n_1 - 1, n_2)$ with rate $\mu_1 n_1$

$(n_1, n_2 + 1)$ with rate λ_2

$(n_1, n_2 - 1)$ with rate $\mu_2 C \dfrac{n_2}{n_1 + n_2}$, with the convention $0/0 = 0$.

These transitions define an irreducible Markov process.

Observe that if $\lambda_1 = 0, \lambda_2 > 0$, the system is described by a M/M/1-PS queue. (PS stands for processor-sharing.) On the other hand if $\lambda_1 > 0, \lambda_2 = 0$, the system is a M/M/C queue. Thus, when $\lambda_1 > 0, \lambda_2 > 0$ one expects the system to share properties of both queueing, and loss systems. Accordingly, in the next section we study stability and in section 4.3 we focus on the behavior near the boundary $n_1 + n_2 = C$.

4.2 Stability

In this section we show that the system is stable, as long as $\lambda_2 < \mu_2 C$. This is a tight characterization of stability, since if there were no real-time flows, the system is just a M/M/1-PS queue with arrival rate λ_2 and service rate $\mu_2 C$.

Also, the condition above does not depend on the real-time flows. This is to be expected, since as the number of non-real-time flows grows large, they tend to consume almost the entire capacity, making the effect of real-time flows irrelevant. This observation is the key in showing stability:

Proposition 4.2.1. *The Markov process $(n_1(\cdot), n_2(\cdot))$ is positive recurrent if and only if $\lambda_2 < \mu_2 C$.*

Proof. We will use the Foster-Lyapunov criterion for stability of Markov chains. Define $f : \mathbb{Z}_+^2 \to \mathbb{R}$ by $f(n_1, n_2) = n_1 + n_2$. The generator A of $(n_1(\cdot), n_2(\cdot))$ is given by

$$Af(n_1, n_2) = \lambda_1 I_{\{n_1+n_2<C\}} + \lambda_2 - \mu_1 n_1 - \mu_2 C \frac{n_2}{n_1 + n_2} \, .$$

Fix $c > 0$ such that $c < \mu_2 C - \lambda_2$. Now, $Af(n_1, n_2) < -c$ for all $n_2 \geq 0$ and $n_1 > M_1$ for some $M_1 > 0$. If $n_1 \leq M_1$ then for every $n_2 > M_2$, for some $M_2 > C$,

$$Af(n_1, n_2) \leq \lambda_2 - \mu_2 C \frac{n_2}{M_1 + n_2} < -c \, .$$

Outside the bounded rectangle $\{(n_1, n_2) \in \mathbb{Z}_+^2 : n_1 \leq M_1, n_2 \leq M_2\}$, we have $Af(n_1, n_2) < -c$, so the process is positive recurrent. If $\lambda_2 \geq \mu_2 C$ then by we compare with another system comprised only of non-real time flows. This is coupled with the original, such that arrivals and file sizes are identical for non-real time flows. This system is an M/M/1 queue with arrival rate no less than the service rate. Starting from a state with the same number of non-real time flows, the flows in the second system will always be less than $n_2(t)$. But, an M/M/1 queue is unstable if $\lambda_2 \geq \mu_2 C$, hence the original system is unstable too. □

4.3 Fluid model

In this section we study the transient behavior of the system. We use $\xrightarrow{d}$ to denote convergence in distribution.

Consider a sequence of systems indexed by $N = 1, 2, \ldots$ such that $C^N = NC, \lambda_1^N = N\lambda_1, \lambda_2^N = N\lambda_2$, and denote by $(n_1^N(t), n_2^N(t))$ the state of the corresponding Markov process at time t. We study the behavior of these processes as $N \to \infty$.

In the case with no non-real-time flows, the limiting dynamics can be derived using the results of Hunt and Kurtz in [24] for general admission controls.

They establish that the limit $\{x_1(t) : t \geq 0\}$ of any weakly convergent subsequence of $t \mapsto \{n_1^N(t)/N; t \geq 0\}$, as $N \to \infty$, satisfies the equation

$$x_1(t) = x_1(0) + \int_0^t \left(\lambda_1 f(x_1(u)) - \mu_1 x_1(u)\right) du \,, \qquad (4.1)$$

for some function $f(\cdot)$ that depends on the specific controls, provided $N^{-1}n_1(0) \xrightarrow{d} x_1(0)$. (The convergence above is weak convergence of distributions on the space $D_{\mathbb{R}}[0, \infty)$, the space of right-continuous real functions on $[0, \infty)$ having left-limits, equipped with the Skorokhod topology.) Depending on the control, (4.1) might have multiple solutions.

In the next section we show that a similar limiting regime holds when real- and non-real-time flows coexist. Moreover the limit is unique.

4.3.1 Convergence

The analysis in this section is close to that in [24]; the differences being first that there is a discontinuity in the interior of the state space, instead of only on the boundary, and second, that departure rates are nonlinear. The discontinuity is dealt by starting the system away from that point; since this point is repelling it does not play any significant role. Nonlinear rates pose no problem, because these concern how fast the system empties. Thus, key properties as tigntness, are not affected. Another difference with [24] is that the system can temporarily exceed capacity.

Define the independent unit rate Poisson processes A_1, A_2, D_1, D_2, on the same probability space. Let $n_1^N(t), n_2^N(t)$ denote the number of real- and non-real-time flows respectively and let $x_i^N := n_i^N(t)/N$, $i = 1, 2$. Also define $m^N(t)$, the number of "free circuits", as

$$m^N(t) = C^N - n_1^N(t) - n_2^N(t)\,.$$

The process $m^N(\cdot)$ takes values in $E := \mathbb{Z} \cup \{+\infty\} \cup \{-\infty\}$, endowed with the usual two-point compactification.

As $N \to \infty$, in a $O(1)$ time interval, the change of (x_1^N, x_2^N) is $O(1)$. During the same interval, m^N makes $O(N)$ transitions, so it moves on a faster timescale than (x_1^N, x_2^N). In the limit there is a complete separation of timescales [24], where the dynamics of the "free circuits" will depend only on the limit $(x_1(t), x_2(t))$ of $(x_1^N(t), x_2^N(t))$.

The evolution of the N-th system is given by the equations

$$\begin{aligned}
x_1^N(t) &= x_1^N(0) + \frac{1}{N} A_1 \left(\int_0^t \lambda_1^N I_{\{m^N(u-) \in \mathcal{A}\}} du \right) \\
&\quad - \frac{1}{N} D_1 \left(\int_0^t \mu_1 n_1^N(u) du \right) \\
x_2^N(t) &= x_2^N(0) + \frac{1}{N} A_2 \left(\int_0^t \lambda_2^N du \right) \\
&\quad - \frac{1}{N} D_2 \left(\int_0^t \mu_2 C^N \frac{x_2^N(u)}{x_1^N(u) + x_2^N(u)} du \right) \,,
\end{aligned} \tag{4.2}$$

for all $t \geq 0$, where $\mathcal{A} = \mathbb{N} \cup \{+\infty\}$ is the *acceptance region*. Notice that since $\{x^N\}$ has no explosions, these equations define a unique solution, for each N.

Let $\mathcal{L}(E)$ denote the space of measures γ on $[0,\infty)\times E$ such that $\gamma([0,t]\times E)=t$ for each t which we endow with the topology of weak convergence of the measures restricted in $[0,t]\times E$ for each t (see [29]). Define the random measure ν^N by

$$\nu^N((0,t)\times\Gamma)=\int_0^t I_{\{m^N(u)\in\Gamma\}}du$$

for all $t\in[0,\infty)$, $\Gamma\in\mathcal{B}(E)$, so for every $\omega\in\Omega$, ν^N takes values in $\mathcal{L}(E)$.

Before we proceed, we need the following

Lemma 4.3.1. *For any $\kappa\in\mathcal{L}(E)$, we have*

$$\kappa([0,t]\times\Gamma)=\int_0^t p_u(\Gamma)du\,,\quad \Gamma\in\mathcal{B}(E)\,, \tag{4.3}$$

where $u\mapsto p_u(\Gamma)$ is Borel measurable for all $\Gamma\in\mathcal{B}(E)$, and $\Gamma\mapsto p_u(\Gamma)$ is a probability measure. (That is, $(u,\Gamma)\mapsto p_u(\Gamma)$ is a regular conditional distribution.)

Proof. Fix $t\geq 0$. Since $\kappa([0,t]\times E)/t=1$, i.e. a probability measure on $[0,t]\times E$, and the marginal is the uniform distribution on $[0,1]$, we have

$$\kappa(C\times\Gamma)=\int_C Q(u,\Gamma)du\,,\quad C\in\mathcal{B}([0,1]),\Gamma\in\mathcal{B}(E)\,,$$

for some regular conditional distribution $(u,\Gamma)\mapsto Q(u,\Gamma)/t$. Now, define $p_u(\Gamma)=Q(u,\Gamma)/n$ for $n-1\leq u<n, n\geq 1$. It is easy to check that (4.3) holds.

□

Now we are ready to state the convergence result.

Theorem 4.3.2. *(i) Suppose $(x_1^N(0),x_2^N(0))\xrightarrow{d}(x_1(0),x_2(0))$, where $x_i(0)>0, i=1,2$. The sequence (x_1^N,x_2^N,ν^N) is relatively compact in $D_{\mathbb{R}^2}[0,\infty)\times\mathcal{L}(E)$, and the limit of any convergent subsequence satisfies,*

$$\begin{aligned} x_1(t)&=x_1(0)+\lambda_1\nu([0,t]\times\mathcal{A})-\int_0^t\mu_1x_1(u-)du\,,\\ x_2(t)&=x_2(0)+\lambda_2t-\int_0^t\mu_2C\frac{x_2(u-)}{x_1(u-)+x_2(u-)}du\,, \end{aligned} \tag{4.4}$$

for all $t\geq 0$.

(ii) *For all* $t \geq 0$,

$$\nu([0,t] \times \mathcal{A}) = \int_0^t \pi_u(\mathcal{A}) du \,, \tag{4.5}$$

where π_u *is a probability measure on* E. *In particular,* $\pi_u(\mathcal{A}) = 1$ *if* $x_1(u)+x_2(u) < C$, $\pi_u(\mathcal{A}) = 0$ *if* $x_1(u) + x_2(u) > C$.

Proof. For the first part, it suffices to establish relative compactness for each $\{x^N\}$ and $\{\nu^N\}$ separately, by Proposition 3.2.4 in Ethier-Kurtz [18]. Since E is compact, $\mathcal{L}(E)$ is compact by Prohorov's theorem. $\{\nu^N\}$ is a sequence of random variables taking values in the compact space $\mathcal{L}(E)$, so using again Prohorov's theorem yields that $\{\nu^N\}$ is tight. Now, note that $x_i^N(t) \geq 0$ for all $t \geq 0$, and

$$x_1^N(t) \leq x_1^N(0) + \frac{1}{N}\tilde{A}_1\left(\int_0^t \lambda_1^N I_{\{m^N(u-)\in\mathcal{A}\}} du\right) + \frac{1}{N}\int_0^t \lambda_1^N I_{\{m^N(u-)\in\mathcal{A}\}} du$$
$$x_2^N(t) \leq x_2^N(0) + \frac{1}{N}\tilde{A}_2\left(\int_0^t \lambda_2^N du\right) + \frac{1}{N}\int_0^t \lambda_2^N du$$

for $\tilde{A}_i(u) := A_i(u) - u, i = 1,2$. But $\tilde{A}_i$ for $i = 1,2$ are square integrable martingales so by Doob's inequality, as $N \to \infty$ the second terms of the left-hand-sides converge to zero uniformly over any bounded time interval, in probability. Using this fact and Corrolary 3.7.4 of Ethier-Kurtz [18], relative compactness of $\{x_1^N, x_2^N\}$ follows.

To show (4.4) we plan to invoke the continuous mapping theorem for the departure terms in (4.2). For this to apply we have to ensure that for each $t \geq$, the limit $(x_1(\cdot), x_2(\cdot))$ is not a point of discontinuity of the map

$$D_{\mathbb{R}_+^2}[0,+\infty) \ni y \mapsto \int_0^t \frac{y_2(u)}{y_1(u)+y_2(u)} du$$

with probability 1. To prove this, we show that this is the case first for some other process $(\tilde{x}_1(\cdot), \tilde{x}_2(\cdot))$. This process is the weak limit of $(\tilde{x}_1^N(\cdot), \tilde{x}_2^N(\cdot))$ defined below, and is such that the continuous mapping theorem directly applies. The construction is such that $\tilde{x}_i(\cdot) = x_i(\cdot), i = 1,2$ if $(x_1(0), x_2(0)) \neq (0,0)$.

Define $(\tilde{x}_1^N(\cdot), \tilde{x}_2^N(\cdot))$ in the following way. Fix $\epsilon > 0$, and let $f : \mathbb{R}_+^2 \to \mathbb{R}$ be such that f is continuous,

$$f(x_1,x_2) = \frac{x_2}{x_1+x_2} \quad \text{on } \{(y_1,y_2) \in \mathbb{R}_+^2 | y_1+y_2 \geq \epsilon\}\,, \quad \text{and } 0 \leq f(x_1,x_2) \leq 1\,.$$

Let $\tilde{m}^N(\cdot) = C^N - \tilde{x}_1^N(\cdot) - \tilde{x}_2^N(\cdot)$ and

$$\tilde{\nu}^N(t) = \int_0^t 1\{\tilde{m}(u-) \in \mathcal{A}\}du \,.$$

Now, let $\tilde{x}_i^N$ be given as the solutions of the equations,

$$\tilde{x}_1^N(t) = x_1^N(0) + \frac{1}{N}A_1\left(\int_0^t \lambda_1^N I_{\{\tilde{m}^N(u-)\in\mathcal{A}\}}du\right) - \frac{1}{N}D_1\left(\int_0^t \mu_1\tilde{x}_1^N(u)Ndu\right)$$
$$\tilde{x}_2^N(t) = x_2^N(0) + \frac{1}{N}A_2\left(\int_0^t \lambda_2^N du\right) - \frac{1}{N}D_2\left(\int_0^t \mu_2 C^N f(\tilde{x}_1^N, \tilde{x}_2^N)du\right) \,,$$

$t \geq 0$. For the same reasons as above, $\{x_1^N, x_2^N, \nu^N, \tilde{x}_1^N, \tilde{x}_2^N, \tilde{\nu}^N\}$ is relatively compact. Since

$$D_{\mathbb{R}_+^2}[0, +\infty) \ni y \mapsto \int_0^t f(y_1(u), y_2(u))du$$

is continuous, the continuous mapping theorem implies that the limit $(x_1, x_2, \nu, \tilde{x}_1, \tilde{x}_2, \tilde{\nu})$ of any convergent subsequence satisfies,

$$\tilde{x}_1(t) = x_1(0) + \lambda_1\tilde{\nu}([0,t] \times \mathcal{A}) - \int_0^t \mu_1\tilde{x}_1(u)du \,,$$
$$\tilde{x}_2(t) = x_2(0) + \lambda_2 t - \int_0^t \mu_2 C f(\tilde{x}_1, \tilde{x}_2)du \,,$$

for all $t \geq 0$. Now, if $\tilde{x}_1(u) + \tilde{x}_2(u) < C$ for some $t_2 > t_1 > 0$ for all $u \in [t_1, t_2]$, then $\tilde{\nu}([t_1, t_2] \times \mathcal{A}) = t_2 - t_1$. This follows from the definition of $\tilde{\nu}$ and the fact that $(\tilde{x}_1(\cdot), \tilde{x}_2(\cdot))$ is continuous on compact intervals. Hence, for a.e. $t \geq 0$ such that $\tilde{x}_1(t) + \beta\tilde{x}_2(t) < 2\epsilon$, we have $\dot{\tilde{x}}_1(t) = \lambda_1 - \mu_1\tilde{x}_1(t)$, for any $\beta > 0$. But then,

$$\dot{\tilde{x}}_1(t) + \dot{\tilde{x}}_2(t) \geq \lambda_1 - \mu_1 2\epsilon + \beta\lambda_2 - \beta \,,$$

which is positive for small enough $\beta > 0$. Thus, if $x_1(0) + \beta x_2(0) \geq 2\epsilon$, we must have $\tilde{x}_1(t) + \beta\tilde{x}_2(t) \geq 2\epsilon$, for all $t \geq 0$. Now $\tilde{x}_i^N \to \tilde{x}_i$, so for any fixed $t \geq 0$, $\liminf_N \inf_{u\in[0,t]} \tilde{x}_1^N(u) + \tilde{x}_2^N(u) > \epsilon$. But by the definition of $\tilde{x}^N$, if $\tilde{x}_1^N(u) + \tilde{x}_2^N(u) > \epsilon$ for all $u \in [0,t]$ then $(\tilde{x}_1^N(u), \tilde{x}_2^N(u)) = (x_1^N(u), x_2^N(u)), u \in [0,t]$. Therefore, if $x_1(0) + x_2(0) > \epsilon$ then $x_1(t) + x_2(t) \geq \epsilon$ for all $u \geq 0$. Since the only point of discontinuity of $(y_1, y_2) \mapsto y_2/(y_1 + y_2)$ is $(0,0)$, the continuous mapping theorem implies that $x_1(\cdot), x_2(\cdot)$ satisfy (4.4).

To show the second part, apply Lemma 4.3.1 to ν. A limiting argument similar to that in the first part, shows that $\pi_u(\mathcal{A})$ has the stated properties, for $x_1(u) + x_2(u) < C$ or $x_1(u) + x_2(u) > C$.

□

It turns out that the limit in Theorem 4.3.2 is unique.

Proposition 4.3.3. *For almost all $t \geq 0$ such that $x_1(t) + x_2(t) = C$,*

$$\dot{x}_1(t) = (\mu_1 x_1(t) + \mu_2 x_2(t) - \lambda_2)^+ \wedge \lambda_1 - \mu_1 x_1(t) \,, \tag{4.6}$$

$$\dot{x}_2(t) = \lambda_2 - \mu_2 x_2(t) \,. \tag{4.7}$$

Moreover, (4.4) *has a unique solution. Therefore* $(x_1^N(\cdot), x_2^N(\cdot))$ *converges weakly to* (4.4).

Proof. For almost all $t \geq 0$ such that $x_1(t) + x_2(t) = C$, $\dot{x}_1(t) + \dot{x}_2(t) = 0$, i.e.,

$$\pi_t(\mathcal{A}) = \frac{\mu_1 x_1(t) - \mu_2 x_2(t) - \lambda_2}{\lambda_1} \,, \tag{4.8}$$

from equation (4.4) and Theorem 4.3.2. But $0 \leq \pi_t(\mathcal{A}) \leq 1$, by Theorem 4.3.2.2, so the set of times for which $\lambda_2 \geq \mu_1 x_1(t) + \mu_2 x_2(t)$ or $\mu_1 x_1(t) + \mu_2 x_2(t) \geq \lambda_1 + \lambda_2$, has zero Lebesgue measure. Hence, (4.6),(4.7) are true for almost all t with $x_1(t) + x_2(t) = C$.

Now, the ODE defined by (4.4) and (4.6)-(4.7) has step discontinuities at the boundary. It is easy to see that for no values of the parameters, $\dot{x}_1(t) + \dot{x}_2(t)$ is both < 0 and > 0 in the interior and exterior, respectively, of a small neighborhood of $\{(y_1, y_2) \in \mathbb{R}_+^2 | y_1 + y_2 = C\}$. (See Figure 4.2.)

Therefore, by Filippov [20] this yields a unique solution.

□

The ODE (4.4), (4.6)-(4.7), implies that while on the boundary, the system can behave in three different ways, depending on the exact point $(x_1, x_2) = (x_1, C - x_1)$ (see Figure 4.3.1.):

1. $\mu_1 x_1 + \mu_2 x_2 > \lambda_1 + \lambda_2$, i.e. total departure rate is higher than total arrival rate. In this case the system cannot maintain itself near the boundary. Thus it will immediately move below the boundary, rendering any admission control unnecessary.

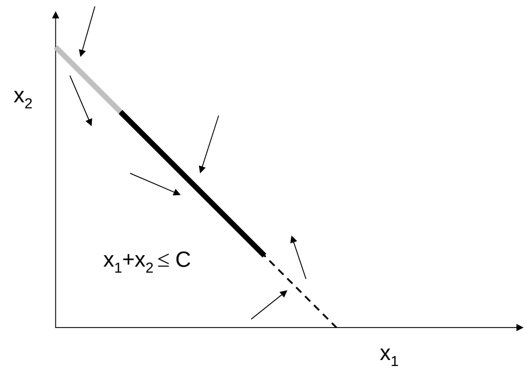

Figure 4.2. Possible behavior around the boundary $x_1 + x_2 = C$.

2. $\lambda_2 > \mu_1 x_1 + \mu_2 x_2$. In this case non-real-time flows arrive faster than flows finish, so the system will move inevitably in states above the boundary, irrespective of the fact that all subsequent real-time flows will be rejected. Thus in this case, admission control is innefective in providing enough bandwidth to the already connected real-time flows. Note that the time for the system to return in states below (or on) the boundary once it escapes, is $O(1)$.

3. $\lambda_2 \leq \mu_1 x_1 + \mu_2 x_2 \leq \lambda_1 + \lambda_2$. In this case admission control can compensate for the variations in the number of non-real-time flows.

The possibility of case 2 above, suggests that user-enforced admission control is not sufficient for providing enough bandwidth for the entire duration of real-time flows once they decide to enter the system. Any admission control that claims to avoid this unpleasant situation, should consider other information as well, such as the total number of real-time flows.

It is natural to ask whether there are any conditions that ensure that case 2 never happens. Intuitively this should be the case when real-time flows last for a small period of time relative to the duration of non-real-time flows, i.e. when $\mu_1^{-1} \ll \mu_2^{-1}$. In this case, departures of real-time flows could adapt fast enough in variations of the number of non-real-time flows. Interestingly enough, this is the case when $\mu_1^{-1} < \mu_2^{-1}$, a relation between the *average* real-time flow duration and the *average* file size. Hence, admission control is likely to be effective when the file sizes are large enough, or equivalently when the minimum bandwidth requirement of real-time flows is small enough.

4.3.2 Equilibrium

In this section we identify the fixed points of the limit process. Clearly, there are no fixed points if $\lambda_2 > \mu_2 C$, i.e., the system is unstable. In the stable case, there can be no fixed points lying above the boundary since then, the number of real-time flows would tend to decrease, i.e., the expected drift of the number of real-time flows would be negative. On the other hand, if there were no real-time flows in the system, then the stability condition forces the number of non-real-time flows to decrease on the average. Hence in the stable case, we expect any fixed points to lie below the boundary.

Considering equations (4.4),(4.6), and (4.7), the following result is immediate.

Proposition 4.3.4. *If $\lambda_2 < \mu_2 C$, the limit process $(x_1(\cdot), x_2(\cdot))$ has a unique fixed point (x_1^*, x_2^*) given by*

$$(x_1^*, x_2^*) = \begin{cases} \left(\rho_1, \rho_1 \dfrac{\rho_2}{C - \rho_2}\right) & \text{if } \rho_1 + \rho_2 < C, \\ (C - \rho_2, \rho_2) & \text{if } \rho_1 + \rho_2 \geq C, \end{cases} \tag{4.9}$$

where $\rho_i = \lambda_i / \mu_i$, $i = 1, 2$.

We call $\rho_1 + \rho_2 \geq C$, $\rho_1 + \rho_2 < C$, the *heavy-traffic* and *light-traffic* case respectively. In the heavy-traffic case the equilibrium is determined solely by the dynamics of non-real-time flows. Moreover, the equilibrium number of non-real-time flows is such *as if* these were real-time flows with average duration μ_2^{-1} and no admission control.

On the other hand, under light-traffic the equilibrium number of real-time flows is independent of the dynamics of non-real-time flows.

Bibliography

[1] M. Alanyali. Three resource management policies for a resource with partition constrains. *Adv. Appl. Prob.*, 31:795–818, 1999.

[2] K.E. Altman, E. Avrachenkov and R. Núñez Queija. Perturbation analysis for denumberable Markov chains with application to queueing models,. *Adv. Appl. Prob.*, 36:839–853, 2004.

[3] M. Andrews, K. Kumaran, K. Ramanan, A. Stolyar, R. Vijayakumar, and P. Whiting. Scheduling in a queueing system with asynchronously varying service rates. *Probab. Engrg. Inform. Sci.*, 18:191–217, 2004.

[4] F. Baccelli and P. Brémaud. Virtual customers in sensitivity and light traffic analysis via Campbell's formula for point processes. *Adv. Appl. Prob.*, 25:221–234, 1993.

[5] F. Baccelli and P. Brémaud. *Elements of Queueing Theory.* Springer-Verlag, 2nd edition, 2003.

[6] F. Baccelli, G. Cohen, and B. Gaujal. Recursive equation and basic properties of timed Petri nets. *Journal of Dynamical Discrete Event Systems*, 1:415–438, 1992.

[7] F. Baccelli and S. Foss. Ergodicity of Jackson-type queueing networks. *Queueing Systems*, 17:5–72, 1994.

[8] F. Baccelli and A. Makowski. Queues in random environment: Stability and bounds. *Stochastic models*, 2:281–291, 1986.

[9] N.G. Bean, R.J. Gibbens, and S. Zachary. Asymptotic analysis of single resource loss systems in heavy traffic, with applications to integrated networks. *Adv. Appl. Prob.*, 27:273–292, 1995.

[10] B. Blaszczyszyn. Factorial moment expansion for stochastic systems. *Stoch. Proc. Appl.*, 56:321–335, 1995.

[11] A. Brzezinski, G. Zussman, and E. Modiano. Enabling distributed throughput maximization in wireless mesh networks via local pooling. Technical Report 2696, MIT/LIDS, 2006.

[12] P. Coolen-Schrijner and E.A. van Doorn. The deviation matrix of a continuous-time Markov chain. *Probab. Engrg. Inform. Sci.*, 16:351–366, 2002.

[13] C. Courcoubetis, A. Dimakis, and Reiman M. Providing bandwidth guarantees over a best-effort network:call-admission and pricing. In *Proc. of IEEE INFOCOM'01*, pages 459–467, 2001.

[14] J. Dai. On positive Harris recurrence of multiclass queueing networks: a unified approach via fluid limit models. *Ann. Appl. Probab.*, 5:49–77, 1995.

[15] J. Dai and B. Prabhakar. The throughput of data switches with and without speedup. In *Proc. of IEEE INFOCOM 2000*, pages 556–564, Tel-Aviv, Israel, 2000.

[16] G. de Veciana, T.-J. Lee, and T. Konstantopoulos. Stability and performance of networks supporting rate control. could the internet be unstable? In *Proc. IEEE INFOCOM'99*, 1999.

[17] R. Durrett. *Probability: Theory and Examples.* Duxbury Press, 3rd edition, 1995.

[18] S. N. Ethier and T. G. Kurtz. *Markov Processes: Characterization and Convergence.* Wiley, New York, 1986.

[19] G. Fayolle, V.A. Malyshev, M.V. Menshikov, and A.F. Sidorenko. Lyapunov functions for Jackson networks. *Mathematics of Operations Research*, 18:916–927, 1993.

[20] A.F. Filippov. Differential equations with discontinuous right hand sides. *Translations of the American Mathematical Society*, 62, 1960.

[21] E. Gelenbe and C. Rosenberg. Queues with slowly varying arrival and service processes. *Management Science*, 36:928–937, 1990.

[22] P. Gupta, Y. Sankarasubramaniam, and A. Stolyar. Random-access scheduling with service differentiation in wireless networks. In *Proc. of IEEE INFOCOM'05*, pages 1815–1825, 2005.

[23] R. Gupta and J. Walrand. Achieving fairness in a distributed ad-hoc MAC. In *Advances in Control, Communication Networks, and Transportation Systems*, In Honor of Pravin Varaiya. Birkhäuser, 2005.

[24] P.J. Hunt and T.G. Kurtz. Large loss networks. *Stochastic Processes and their Applications*, 53:363–378, 1994.

[25] F.P. Kelly. Blocking probabilities in large circuit-switched networks. *Adv. Appl. Prob.*, 18:473–505, 1986.

[26] F.P. Kelly, P.B. Key, and S. Zachary. Distributed admission control. *IEEE Journal on Selected Areas in Communications*, 18:2617–2628, 2000.

[27] G. Koole. The deviation matrix of the M/M/1 and M/M/1/N queue, with applications to controlled queueing models. In *Proc. of the 37th IEEE CDC*, pages 56–59, Tampa, 1998.

[28] S. Kumar, P. Giaccone, and E. Leonardi. Rate stability of stable marriage scheduling algorithms in input-queued switches. In *Proc. of 40th Annual Allerton Conference*, 2002.

[29] T. G. Kurtz. Averaging for martingale problems and stochastic approximation. In *Proc. US-French Workshop*, volume 177 of *Lecture Notes in Control and Information Sciences*, pages 186–209. Springer, Berlin, 1992.

[30] V.A. Malyshev. Networks and dynamical systems. *Adv. Appl. Prob.*, 25:140–175, 1993.

[31] L. Massoulié, P. Key, A. Bain, and F. Kelly. A network flow model for mixtures of file transfers and streaming traffic. Technical Report 2003-37, Microsoft Reseach Technical Report, 2003.

[32] L. Massoulié and J. Roberts. Bandwidth sharing: objectives and algorithms. *IEEE/ACM Transactions on Networking*, 10:320–328, 2002.

[33] N. McKeown. *Scheduling Algorithms for Input-Queued Cell Switches*. PhD thesis, Univ. of Cal. at Berkeley, 1995.

[34] S.P. Meyn and R.L. Tweedie. *Markov Chains and Stochastic Stability*. Springer, London, 1993.

[35] M. Neuts. The M/M/1 queue with randomly varying arrival and service rates. *Opsearch*, 15:139–157, 1978.

[36] M. Reiman and B. Simon. Open queueing systems in light traffic. *Mathematics of Operations Research*, 14:26–59, 1989.

[37] S. Shakkottai and A. Stolyar. Scheduling for multiple flows sharing a time-varying channel: The exponential rule. In *Analytic Methods in Applied Probability*, pages 185–202. AMS, Providence, RI, 2002.

[38] A. Stolyar. On the stability of multiclass queueing networks: a relaxed sufficient condition via limiting fluid processes. *Markov Process and Related Fields*, 1:491–512, 1995.

[39] A. Stolyar. Maxweight scheduling in a generalized switch: State space collapse and workload minimization in heavy traffic. *Ann. Appl. Probab.*, 14:1–53, 2004.

[40] L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Trans. Automat. Control*, 47:1936–1948, 1992.

[41] G.G. Yin and Q. Zhang. *Continuous-time Markov Chains and Applications*. Springer-Verlag, 1998.

[42] S. Zachary and I. Ziedins. A refinement of the Hunt-Kurtz theory of large loss networks, with an application to virtual partitioning. *Ann. Appl. Probab.*, 12, 2002.