Stabilization of Branching Queueing Networks*

Tomáš Brázdil[1] and Stefan Kiefer[2]

1 Faculty of Informatics, Masaryk University, Czech Republic
brazdil@fi.muni.cz

2 Department of Computer Science, University of Oxford, United Kingdom
stefan.kiefer@cs.ox.ac.uk

──── **Abstract** ────

Queueing networks are gaining attraction for the performance analysis of parallel computer systems. A Jackson network is a set of interconnected servers, where the completion of a job at server i may result in the creation of a new job for server j. We propose to extend Jackson networks by "branching" and by "control" features. Both extensions are new and substantially expand the modelling power of Jackson networks. On the other hand, the extensions raise computational questions, particularly concerning the stability of the networks, i.e, the ergodicity of the underlying Markov chain. We show for our extended model that it is decidable in polynomial time if there exists a controller that achieves stability. Moreover, if such a controller exists, one can efficiently compute a static randomized controller which stabilizes the network in a very strong sense; in particular, all moments of the queue sizes are finite.

1 Introduction

Queueing theory plays a central role in the performance analysis of computer systems. In particular, *queueing networks* are gaining attraction as models of parallel systems. A queueing network is a set of processing units (called *servers*), each of which performs tasks (called *jobs*) of a certain type. Each server has its own queue of jobs waiting to be processed. The successful completion of a job may trigger one (or more) new jobs (of possibly different type) that need to be processed as well. In addition to this "internal" job creation, so-called *open* queueing networks allow for new jobs to arrive "externally", i.e., from outside.

Queueing networks are a popular model for both hardware and software systems because of their simplicity and generality. On the hardware side, queueing networks can, e.g., be used for modeling multi-core processors, see e.g. [30] and the references in [9]. One advantage of queueing-based analyses is their scalability with growing parallelism; e.g., it is said in [21]: "Cycle-accurate full-system performance simulators do not scale well beyond a few tens of processor cores at best. As such, analytical models based on the theory of queueing systems, are a logical choice for developing a basic understanding of the fundamental tradeoffs in future, large-scale multi-core systems." On the software side, queueing networks are used for modeling message passing. It is said in [29]: "Two natural classes of systems can be modeled using such a framework: asynchronous programs on a multi-core computer and distributed programs communicating on a network." Of course, the realm of queueing networks stretches far beyond computer science, see [5, 8].

The simplest queueing networks are so-called *Jackson networks* [16]: Given two servers $i, j \in \{1, \dots, n\}$, there is a "rule" of the form $i \xrightarrow{p_{ij}} j$ which specifies the probability p_{ij} that the completion of an i-job results in the creation of a j-job. There are also rules $i \xrightarrow{p_{i0}} \varepsilon$ where $p_{i0} = 1 - \sum_j p_{ij}$ specifies the probability that no new job is created. Each server i has a *rate* μ_i with which an i-job is processed if there is one. In addition, there is a rate α_i with which i-jobs arrive from outside the network. The processing times and the external arrivals are exponentially distributed, so that a Jackson network describes a *continuous-time Markov chain (CTMC)*. It was shown in Jackson's paper [16] that if the rate λ_i of internal *and* external arrivals at server i is less than μ_i for all i, then the network is *stable*, i.e., the average queue length is finite and almost surely all queues are empty infinitely often. Moreover, Jackson networks allow a *product form*, i.e., the steady-state distribution of the queue lengths can be expressed as a product of functions $\pi_i(k)$, where $\pi_i(k)$ is the steady-state probability that queue i has length k.

▶ **Example 1** (network processor). In [1], Jackson networks are used to model network processors, i.e., chips that are specifically targeted at networking applications—think of a router. We describe the model from [1] (sections 4.1 and 4.2, slightly adapted). Before packets are processed in the "master processor" M, they pass through the "data plane" D, from which a fraction q of packets needs to be processed first in the "control plane" C:

$$D \xrightarrow{1-q} M \qquad D \xrightarrow{q} C \qquad C \xrightarrow{1} M \qquad M \xrightarrow{1} \varepsilon$$

An "arrival manager" A sends some packets (fraction d_0) directly to D, but others (fractions $d_1, \dots, d_n$ with $d_0 + d_1 + \cdots + d_n = 1$) are sent to "slave processors" $S_1, \dots, S_n$ to assist the master. Some packets (fraction b) still need the attention of the master after having been processed by a slave:

$$A \xrightarrow{d_0} D \qquad A \xrightarrow{d_i} S_i \qquad S_i \xrightarrow{b} D \qquad S_i \xrightarrow{1-b} \varepsilon \quad , \quad i \in \{1, \dots, n\}.$$

Jackson networks and their extensions have been thoroughly studied, but they are restricted in their modelling capabilities, as (i) the completion of a job may trigger at most one job, and (ii) there is no nondeterminism that would allow to control the output probabilities of a server. Considering (i), it seems unnatural to assume that a distributed program communicating on a network produces at most one message at the end of its computation. Considering (ii), the "arrival manager" A in Example 1 may want to flexibly pass incoming packets to the master or one of the slaves, possibly depending on the current load. These restrictions have not been fully addressed, not even in isolation. In this paper we introduce *controlled branching queueing networks*, which are Jackson-like networks but allow for both nondeterminism ("controlled") and the creation of more than one job ("branching").

Both extensions directly raise computational issues. We show in Example 2 on page 5 that even purely stochastic branching networks do not allow a product form, which illustrates the mathematical challenge[1] of this extension and poses the question for an effective criterion that allows to determine whether the network is stable, i.e., returns to the empty state infinitely often. Moreover, due to the nondeterminism, we now deal with *continuous-time Markov decision processes (CTMDPs)*. Our main theorem (Theorem 3) states that if there exists any scheduler resolving the nondeterminism in such a way that the controlled branching

---

[1] It is noted in [15] that "[. . .] virtually all of the models that have been successfully analyzed in classical queueing network theory are models having a so-called product form stationary distribution."

network is stable, then there exists a *randomized static* scheduler that achieves stability as well, where by "randomized static" we mean that the decisions may be randomized but not dependent on the current state (such as the load) of the system. Moreover, the existence of such a stabilizing scheduler and the scheduler itself can be determined in polynomial time, and, finally, the randomized static scheduler is stabilizing in a very strong sense, in particular, all moments of the queue sizes are finite.

Related work. We use nondeterminism to describe systems whose behaviour is not completely specified. A system *designer* can then resolve the nondeterminism to achieve certain goals, in our case stability. Although nondeterminism is a very well established modelling feature of probabilistic systems (see e.g. [20]), the literature on automatic design of stabilizing controllers for queueing networks is sparse. *Flow-controlled* networks [28, 22] allow to control only the external arrival stream or the service rates (see also [2] and the references therein). The authors of [19, 14] consider queueing networks with fewer servers than job types, so that the controller needs to assign servers to queues. As in [19, 14], we also use linear programming to design a controller, but our aim is different: we allow the controller to influence the production of the individual queues, and we study the complexity of designing stabilizing controllers and the nature of such controllers. There has been a substantial amount of work in the last years analyzing probabilistic systems with "branching features", most prominently on *recursive Markov chains* [13, 12] and *probabilistic pushdown systems* [11, 6]. While these models allow for a probabilistic splitting of tasks by pushing new procedures on a stack, the produced tasks are processed in a strictly sequential manner, whereas the queues in a queueing network process jobs in parallel and in continuous time. Recently, *probabilistic split-join systems* were introduced [17], which allow for branching but not for external arrivals, and assume unlimited parallelism. In [18, chapter 8] a queueing model with multiple classes of tasks and "feedback" is discussed, which is similar to our branching except that there is only one server, hence there is no parallelism. Algorithmic theory of queueing systems has also attracted some attention in the past. In particular, for *closed* (i.e., without external arrivals) queueing systems, [26] shows EXP-completeness of minimizing a weighted throughput of the queues.

2   Preliminaries

Numbers. We use $\mathbb{Z}, \mathbb{Q}, \mathbb{R}$ for the sets of integer, rational, real numbers, respectively, and $\mathbb{N}, \mathbb{Q}_{\geq 0}, \mathbb{R}_{\geq 0}$ for their respective subsets of nonnegative numbers.

Vectors and Matrices. Let $n \geq 1$. We use boldface letters for denoting vectors $\boldsymbol{x} = (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n) \in \mathbb{R}^n$. Vectors are row vectors per default, and we use superscript ^{T} for transpose, so that $\boldsymbol{x}^T$ denotes a column vector. If the dimension n is clear from the context, we write $\boldsymbol{0} := (0, \ldots, 0)$, $\boldsymbol{1} := (1, \ldots, 1)$, and $\boldsymbol{e}^{(i)} = (0, \ldots, 0, 1, 0, \ldots, 0)$ for the vector with the 1 at the ith component $(1 \leq i \leq n)$. It is convenient to define $\boldsymbol{e}^{(0)} := \boldsymbol{0}$. For two vectors $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$ we write $\boldsymbol{x} \sim \boldsymbol{y}$ with $\sim \in \{=, <, \leq, >, \geq\}$ if the respective relation holds componentwise. For a vector $\boldsymbol{x} \in \mathbb{R}^n$ we denote its *1-norm* by $\|\boldsymbol{x}\| := \sum_{i=1}^{n} |\boldsymbol{x}_i|$. When $\boldsymbol{x} \in \mathbb{N}^n$ is a vector of queue sizes, we refer to $\|\boldsymbol{x}\|$ as the *total queue size*. For a matrix $A \in \mathbb{R}^{n \times n}$, we write A_i for its ith row, i.e., $A_i = (A_{i1}, \ldots, A_{in})$.

CTMDP. A *continuous-time Markov decision process (CTMDP)* consists of an at most countable set S of states, an initial state $s_1 \in S$, a set of actions Σ [2], and a transition rate

---

[2] Usually, each state has its own set of available actions. As this feature is not needed for queueing networks, we stick to the simpler version in which all actions are always available.

$q(s, \sigma, s') \geq 0$ for each pair of states $s, s' \in S$ and each action $\sigma \in \Sigma$ (here $q(s, \sigma, s') = 0$ means that the transition from s to s' never occurs). We define a *continuous-time Markov chain (CTMC)* to be a CTMDP whose set of actions Σ is a singleton (we usually do not write the only action explicitly, so the transition rates of a CTMC are denoted by $q(s, s')$, etc.).

Intuitively, a run of a CTMDP starts in s_1 and then evolves in so-called epochs. Assume that (after the previous epoch) the system is in a state s. The next epoch consists of the following phases: First, a scheduler chooses an action $\sigma \in \Sigma$ to be executed. Second, a waiting time for transition to each state $s' \in S$ is chosen according to the exponential distribution with the rate $q(s, \sigma, s')$ (here we assume that if $q(s, \sigma, s') = 0$, then the waiting time is ∞). The transitions compete in a way such that the one with the least waiting time is executed and the state of the CTMDP is chosen accordingly (the other transitions are subsequently discarded).

Formally, a *run* is an infinite sequence $s_1, \sigma_1, t_1, s_2, \sigma_2, t_2, \ldots \in (S \times \Sigma \times \mathbb{R}_{\geq 0})^\omega$. We denote by *Run* the set of all runs. A *scheduler* is a function Θ which assigns to every finite path $s_1, \sigma_1, t_1, s_2, \sigma_2, t_2, \ldots s_n \in (S \times \Sigma \times \mathbb{R}_{\geq 0})^* \times S$ a probability distribution $\Theta(w)$ on actions (i.e. $\Theta(w) : \Sigma \to [0, 1]$ satisfies $\sum_{\sigma \in \Sigma} \Theta(w)(\sigma) = 1$). For technical reasons, we have to restrict ourselves to measurable schedulers (for details see e.g. [24]).

We work with a measurable space of runs $(Run, \mathcal{F})$ where $\mathcal{F}$ is the smallest σ-algebra generated by basic cylinders (i.e. sets of runs with common finite prefix) in a standard way. Every scheduler Θ induces a unique probability measure Pr_Θ on $\mathcal{F}$ determined by the probabilities of the basic cylinders. For detailed definitions see [24]. Then each scheduler Θ *induces* a stochastic process $(x(t) \mid t \in \mathbb{R}_{\geq 0})$ on the probability space $(Run, \mathcal{F}, \mathrm{Pr}_\Theta)$ where $x(t)$ is the current state of the run in time t, i.e., each $x(t)$ is a random variable defined by

$$x(t)(s_1, \sigma_1, t_1, s_2, \sigma_2, t_2, \ldots) = s_i \quad , \quad \sum_{j=1}^{i-1} t_i \leq t \text{ and } \sum_{j=1}^{i} t_i \geq t.$$

A scheduler Θ is *memoryless* if for every path $w = s_1, \sigma_1, t_1, s_2, \sigma_2, t_2, \ldots s_{n+1} \in (S \times \Sigma \times \mathbb{R}_{\geq 0})^* \times S$ we have that $\Theta(w) = \Theta(s_{n+1})$.

Networks. Define $\mathsf{R}^{(n,K)} := \{\boldsymbol{r} \in \mathbb{N}^n \mid \boldsymbol{r}_1 + \cdots + \boldsymbol{r}_n \leq K\}$. A *production function* for (n, K) is a function $Prob : R \to \mathbb{Q} \cap (0, 1]$ with $R \subseteq \mathsf{R}^{(n,K)}$ such that $\sum_{\boldsymbol{r} \in R} Prob(\boldsymbol{r}) = 1$. A *controlled branching network with n queues and branching factor K* consists of an arrival rate $\mu_0 \in \mathbb{Q} \cap (0, \infty)$, queue rates $\mu_i \in \mathbb{Q} \cap (0, \infty)$ for $i \in \{1, \ldots, n\}$, an arrival production function $Prob_0 : R_0 \to \mathbb{Q} \cap (0, 1]$ for (n, K), and finite action sets $\Sigma_1, \ldots, \Sigma_n$ as follows. An action $\sigma_i \in \Sigma_i$ assigns to queue i a production function $Prob_i(\sigma_i) : R_i(\sigma_i) \to \mathbb{Q} \cap (0, 1]$ for (n, K). Define $\Sigma := \Sigma_1 \times \cdots \times \Sigma_n$. If $\sigma = (\sigma_1, \ldots, \sigma_n) \in \Sigma$, we write $R_i(\sigma)$, $Prob_i(\sigma)$ and $R_0(\sigma)$, $Prob_0(\sigma)$ to mean $R_i(\sigma_i)$, $Prob_i(\sigma_i)$ and R_0, $Prob_0$. Observe that the rates μ_i do not depend on actions. This simplification is without loss of generality.[3] We assume a nonzero arrival stream, i.e., there is $\boldsymbol{r} \in R_0$ with $\boldsymbol{r} \neq \boldsymbol{0}$. We define the *size* of a controlled network by $n + K + (\sum_{i=0}^{n} |\mu_i|) + |R_0| + |Prob_0| + \sum_{i=1}^{n} \sum_{\sigma_i \in \Sigma_i} |R_i(\sigma_i)| + |Prob_i(\sigma_i)|$, where $|\mu_i|$ etc. means the description size assuming the rationals are represented as fractions of integers in

---

[3] To show that this assumption is w.l.o.g. one can employ the standard "uniformization" trick. More precisely, assume that the actions of a queue i have different rates. Define μ_i to be the maximum of all rates of Σ_i and compensate by "adding self-loops", i.e., make the actions of Σ_i generate a new job for queue i with a suitable probability. This effectively substitutes a transition with longer delay by possibly several transitions with delay μ_i. As static schedulers can be easily translated between the original and the transformed system, our results remain valid.

binary. A controlled branching network induces a CTMDP with state space $\mathbb{N}^n$ (the queue sizes), initial state $\mathbf{0}$, action set Σ, and transition rates

$$q(\boldsymbol{x}, \sigma, \boldsymbol{y}) = \sum_{i \in \{0,1,\ldots,n\}:i=0 \vee \boldsymbol{x}_i \neq 0} \;\; \sum_{\boldsymbol{r} \in R_i(\sigma):\boldsymbol{y}=\boldsymbol{x}-\boldsymbol{e}^{(i)}+\boldsymbol{r}} \mu_i Prob_i(\sigma)(\boldsymbol{r}) \qquad \text{for } \boldsymbol{x}, \boldsymbol{y} \in \mathbb{N}^n, \sigma \in \Sigma.$$

Interpreting this definition, there is a "race" between external arrivals (rate μ_0) and the nonempty queues (rates μ_i); if the external arrivals win, new jobs are added according to $Prob_0(\sigma)$; if queue i wins, one i-job is removed and new jobs are added according to $Prob_i(\sigma)$.

An *purely stochastic branching network* is a controlled branching network with $\Sigma = \{\sigma\}$, i.e., with a unique action for each queue. Hence, the induced CTMDP is a CTMC. In the purely stochastic case we write only R_i for $R_i(\sigma)$ etc. If $Prob_i(\boldsymbol{r}) = p$ in the purely stochastic case, we use in examples the notation $i \xrightarrow{p} \boldsymbol{r}$, where we often write $\boldsymbol{r} \in \mathsf{R}^{(n,K)}$ as a multiset without curly brackets. For instance, if $n = 2$, we write $1 \xrightarrow{p} 1, 2$ and $1 \xrightarrow{p} 2, 2$ and $1 \xrightarrow{p} \varepsilon$ to mean $Prob_1(\boldsymbol{r}) = p$ with $\boldsymbol{r} = (1,1)$ and $\boldsymbol{r} = (0,2)$ and $\boldsymbol{r} = (0,0)$, respectively.

Fixing a controlled network $\mathcal{N}$ and a scheduler Θ for the CTMDP induced by $\mathcal{N}$, we obtain a stochastic process $\mathcal{N}_\Theta = (\boldsymbol{x}(t) \mid t \in \mathbb{R}_{\geq 0})$, where $\boldsymbol{x}(0) = \mathbf{0} \in \mathbb{N}^n$, which evolves according to the dynamics of $\mathcal{N}$ and the scheduler Θ. In the purely stochastic case we drop the subscript Θ, and so we identify a network $\mathcal{N}$ with its induced stochastic process.

▶ **Example 2** (no product form). Consider the purely stochastic branching network with $0 \xrightarrow{1} 1, 2$ and $1 \xrightarrow{1} \varepsilon$ and $2 \xrightarrow{1} \varepsilon$. If its stationary distribution π (for a definition of stationary distribution see before Theorem 3) had product from, the queues would be "independent in steady-state", i.e., $\pi(\boldsymbol{x}_2 \geq 1 \mid \boldsymbol{x}_1 \geq 1) = \pi(\boldsymbol{x}_2 \geq 1)$, where by $\boldsymbol{x}$ we mean $\boldsymbol{x}(t)$ in steady state. However, if μ_0 is much smaller than $\mu_1 = \mu_2$, then we have $\pi(\boldsymbol{x}_2 \geq 1 \mid \boldsymbol{x}_1 \geq 1) > \pi(\boldsymbol{x}_2 \geq 1)$, intuitively because $\boldsymbol{x}_1 \geq 1$ probably means that there was an arrival recently, so that $\boldsymbol{x}_2 \geq 1$ is more likely than usual. More concretely, let $\mu_0 = 1$ and $\mu_1 = \mu_2 = 3$ and consider the 2-state Markov chain obtained by assuming that each arrival leads to the state $(1,1)$ and each completion of any job leads to the state $(0,0)$. By computing the stationary distribution π' of this 2-state Markov chain in the standard way, we obtain $\pi'((0,0)) = 6/7$ and $\pi'((1,1)) = 1/7$. Since this 2-state Markov chain "underapproximates" the CTMC induced by the network, we have $\pi(\boldsymbol{x}_1 \geq 1 \wedge \boldsymbol{x}_2 \geq 1) \geq 1/7$. On the other hand, by considering the two queues separately, the standard formula for the M/M/1 queue gives $\pi(\boldsymbol{x}_1 \geq 1) = \pi(\boldsymbol{x}_2 \geq 1) = 1/3$. Product form would imply $\pi(\boldsymbol{x}_1 \geq 1 \wedge \boldsymbol{x}_2 \geq 1) = \pi(\boldsymbol{x}_1 \geq 1) \cdot \pi(\boldsymbol{x}_2 \geq 1) = 1/9$, contradicting the inequality above.

3 Results

We focus on the stability of purely stochastic and controlled branching networks. Our notion of stability requires that the network is completely empty infinitely many times. Given a stochastic process $(\boldsymbol{x}(t) \mid t \in \mathbb{R}_{\geq 0})$, we say that the process is *ergodic* if the expected return time to $\mathbf{0}$ is finite. More formally, define a random variable R by

$$R \quad := \quad \inf_{t>0} \{t \mid \boldsymbol{x}(t) = \mathbf{0}, \exists t' < t : \boldsymbol{x}(t') \neq \mathbf{0}\} \,.$$

Then the process is ergodic iff $\mathbb{E}[R] < \infty$. In the controlled case, we say that a scheduler Θ for $\mathcal{N}$ is *ergodic for* $\mathcal{N}$ if $\mathcal{N}_\Theta$ is ergodic. In the following we use stability and ergodicity interchangeably. A scheduler Θ is *static* if it always chooses the same fixed distribution on actions. Note that static schedulers are memoryless. If in a stochastic process $(\boldsymbol{x}(t) \mid t \in \mathbb{R}_{\geq 0})$

the limit $\pi(\boldsymbol{x}) := \lim_{t \to \infty} \Pr(\boldsymbol{x}(t) = \boldsymbol{x})$ exists for all $\boldsymbol{x} \in \mathbb{N}^n$ and $\sum_{\boldsymbol{x} \in \mathbb{N}^n} \pi(\boldsymbol{x}) = 1$, then $\pi : \mathbb{R}_{\geq 0} \to [0, 1]$ is called the *stationary distribution*.

▶ **Theorem 3.** *Let $\mathcal{N}$ be a controlled branching network. It is decidable in polynomial time whether there exists an (arbitrary) ergodic scheduler for $\mathcal{N}$. If it exists, one can compute, in polynomial time, a* static randomized *ergodic scheduler for $\mathcal{N}$ with stationary distribution π such that there exists an exponential moment of the total number of waiting jobs, i.e., there is $\delta > 0$ such that $\sum_{\boldsymbol{x} \in \mathbb{N}^n} \exp(\delta \|\boldsymbol{x}\|) \pi(\boldsymbol{x})$ exists.*

To prove Theorem 3 we generalize the concept of *traffic equations* (see e.g. [7]) from the theory of Jackson networks. Intuitively, the traffic equations express the fact that the inflow of jobs to a given queue must be equal to the outflow. Remarkably, the traffic equations characterize the stability of the Jackson network. More precisely, a Jackson network is stable if and only if there is a solution of the traffic equations whose components are strictly smaller than the rates of the corresponding queues (we call such a solution *deficient*).

We show how to extend the traffic equations so that they characterize the stability of controlled branching networks. For a smooth presentation, we start with *purely stochastic* branching networks and add control later on. Hence, the overall plan of the proof of Theorem 3 is as follows: Set up traffic equations for purely stochastic branching networks and show that if there is a deficient solution of these equations, then the network is stable. This result, presented in Section 3.1 (Proposition 4), is of independent interest and requires the construction of a suitable *Lyapunov function*. Then, in Section 3.2, we generalize the traffic equations to controlled branching networks and show that any ergodic scheduler determines a deficient solution (Proposition 10). This solution naturally induces a static scheduler, which, when fixed, determines an purely stochastic network with deficiently solvable traffic equations. Propositions 4 and 10 imply Theorem 3 and provide some additional results.

3.1   Purely stochastic branching networks

Assume that $\mathcal{N}$ is purely stochastic, i.e., there is a single action for each queue. In such a case the CTMDP induced by the network is in fact a CTMC. We associate the following quantities to a network, which will turn out to be crucial for its performance. Let $\boldsymbol{\mu} := (\mu_1, \ldots, \mu_n)$. Let $\boldsymbol{\alpha} \in \mathbb{R}_{\geq 0}^n$ be the vector with $\boldsymbol{\alpha}_i := \mu_0 \sum_{\boldsymbol{r} \in R_0} Prob_0(\boldsymbol{r}) \boldsymbol{r}_i$; i.e., $\boldsymbol{\alpha}_i$ indicates the expected number of external arrivals at queue i per time unit. Note that $\boldsymbol{\alpha} \neq \mathbf{0}$, as we assume a nonzero arrival stream. Let $A \in \mathbb{R}_{\geq 0}^{n \times n}$ be the matrix with $A_{ij} := \sum_{\boldsymbol{r} \in R_i} Prob_i(\boldsymbol{r}) \boldsymbol{r}_j$; i.e., A_{ij} indicates the expected production of j-jobs when queue i fires. W.l.o.g. we assume that all queues are "reachable", i.e., for all queues i there is $j \in \mathbb{N}$ with $(\boldsymbol{\alpha} A^j)_i \neq 0$. We define a set of *traffic equations*

$$\lambda_j \quad = \quad \alpha_j + \sum_{i=1}^{n} \lambda_i \cdot A_{ij} \quad , \quad j \in \{1, \ldots, n\}, \tag{1}$$

in matrix form:

$$\boldsymbol{\lambda} = \boldsymbol{\alpha} + \boldsymbol{\lambda} A. \tag{2}$$

We prove the following proposition.

▶ **Proposition 4.** *Assume that $\boldsymbol{\lambda} \in \mathbb{R}_{\geq 0}^n$ solves the traffic equations* (2) *and satisfies $\boldsymbol{\lambda} < \boldsymbol{\mu}$. Then the following conclusions hold:*
1. *The process $\mathcal{N}$ is ergodic, i.e., the expected return time to $\mathbf{0}$ is finite.*

2. *There exists a stationary distribution π such that there exists an exponential moment of the total queue size, i.e., there is $\delta > 0$ such that $\sum_{\boldsymbol{x} \in \mathbb{N}^n} \exp(\delta \|\boldsymbol{x}\|)\pi(\boldsymbol{x})$ exists.*

The key step to the proof of Proposition 4 is to construct a so-called *Lyapunov function* with respect to which the process $\mathcal{N}$ exhibits a "negative drift". This is in fact a classical technique for showing the stability of queueing systems [23]; the difficulty lies in finding a suitable Lyapunov function. The "drift" of $\mathcal{N}$ is given by the *mean velocity vector* $\boldsymbol{\Delta}(\boldsymbol{x}) \in \mathbb{R}^n_{\geq 0}$ of $\mathcal{N}$, defined by $\boldsymbol{\Delta}(\boldsymbol{x}) := \lim_{h \to 0^+} \mathbb{E}\left[\boldsymbol{x}(t+h) - \boldsymbol{x}(t) \mid \boldsymbol{x}(t) = \boldsymbol{x}\right]/h$. The limit exists, is independent of t, and we have

$$\boldsymbol{\Delta}(\boldsymbol{x}) = \boldsymbol{\alpha} + \sum_{i:\boldsymbol{x}_i \neq 0} \mu_i(-\boldsymbol{e}^{(i)} + A_i)\,. \tag{3}$$

The following lemma is implicitly proved in [10, theorem 1].

▶ **Lemma 5.** *Suppose that a function $\widetilde{V} : \mathbb{R}^n_{\geq 0} \to \mathbb{R}_{\geq 0}$ is two times continuously differentiable, $\widetilde{V}(\boldsymbol{x}) = 0$ implies $\boldsymbol{x} = \boldsymbol{0}$, and that there is $\gamma > 0$ such that we have*

$$\boldsymbol{\Delta}(\boldsymbol{x})\big(\widetilde{V}'(\boldsymbol{x})\big)^T \leq -\gamma \qquad \text{for all } \boldsymbol{x} \neq \boldsymbol{0},$$

where $\widetilde{V}'(\boldsymbol{x})$ denotes the gradient of $\widetilde{V}$ at $\boldsymbol{x}$. Then the conclusions of Proposition 4 holds.

Following [10], we construct the Lyapunov function $\widetilde{V}$ in two stages: we first define a suitable *piecewise linear* function $V : \mathbb{N}^n \to \mathbb{R}_{\geq 0}$; then V is *smoothed* to obtain $\widetilde{V}$. For the definition of V we need the following lemma.

▶ **Lemma 6.** *The matrix series $A^* := \sum_{i=0}^{\infty} A^i$ converges ("exists") in $\mathbb{R}^{n \times n}_{\geq 0}$ and is equal to $(I - A)^{-1}$.*

Define vectors $\boldsymbol{q}^{(1)}, \ldots, \boldsymbol{q}^{(n)} \in \mathbb{R}^n_{\geq 0}$ by setting $\boldsymbol{q}^{(i)T} := \boldsymbol{a}^{(i)T}/\|\boldsymbol{a}^{(i)}\|$, where $\boldsymbol{a}^{(i)T}$ is the ith column of A^*. Observe that we have $\boldsymbol{1}\boldsymbol{q}^{(i)T} = 1$ for all i. Define the function $V : \mathbb{R}^n_{\geq 0} \to \mathbb{R}_{\geq 0}$ by $V(\boldsymbol{x}) := \max_i\{\boldsymbol{x}\boldsymbol{q}^{(i)T}\}$. We will use the following property of V:

▶ **Lemma 7.** *If $\boldsymbol{0} \neq \boldsymbol{x} \in \mathbb{R}^n_{\geq 0}$ and $\boldsymbol{x}_i = 0$, then $\boldsymbol{x}\boldsymbol{q}^{(i)T} < V(\boldsymbol{x})$.*

Lemma 7 is not obvious; in the appendix we use Farkas' lemma for the proof. The following lemma describes the crucial "negative drift" property of V:

▶ **Lemma 8.** *There is $\gamma > 0$ such that we have*

$$\boldsymbol{\Delta}(\boldsymbol{x})\big(V'(\boldsymbol{x})\big)^T \leq -\gamma \qquad \text{for all } \boldsymbol{x} \neq \boldsymbol{0}$$

and all subgradient vectors $V'(\boldsymbol{x})$ of V at $\boldsymbol{x}$. More precisely, one can choose

$$\gamma := \min_i(\mu_i - \boldsymbol{\lambda}_i)/\left\|\boldsymbol{a}^{(i)}\right\|,$$

where $\boldsymbol{a}^{(i)T}$ is the ith column of A^*.

▶ **Example 9.** Consider the network with

$$0 \xrightarrow{1} 1 \qquad \begin{array}{l} 1 \xrightarrow{1/5} 2,2 \\ 1 \xrightarrow{4/5} \varepsilon \end{array} \qquad \begin{array}{l} 2 \xrightarrow{1/6} 1,2 \\ 2 \xrightarrow{5/6} \varepsilon \end{array}\,,$$

arrival rate $\mu_0 = 7/30$, and $\boldsymbol{\mu} = (5/12, 7/20)$, so that $\mu_0 + \mu_1 + \mu_2 = 1$. Let us write $[0] := \boldsymbol{\alpha} = \mu_0(1, 0)$, $[1] := \mu_1(-\boldsymbol{e}^{(1)} + A_1)$, $[2] := \mu_2(-\boldsymbol{e}^{(2)} + A_2)$, $[01] := [0] + [1]$,
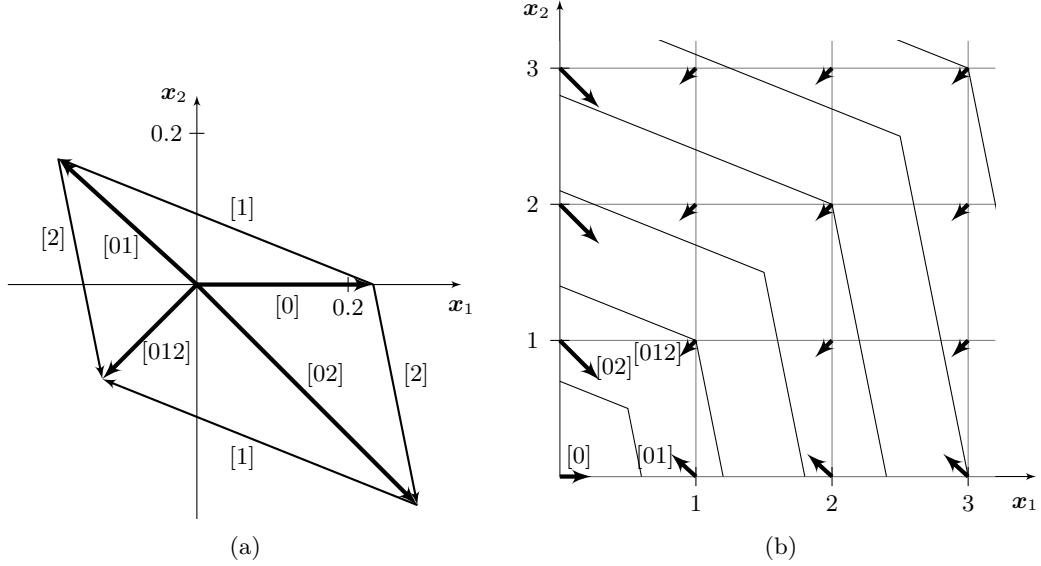
■ **Figure 1** Illustration of negative drift.

$[02] := [0] + [2]$, $[012] := [0] + [1] + [2]$. These vectors are shown in Figure 1 (a). The mean velocity vector $\boldsymbol{\Delta}(\boldsymbol{x})$ is one of the vectors $[0], [01], [02], [012]$, depending on which components of $\boldsymbol{x}$ are nonzero. The vector field in Figure 1 (b) shows the corresponding vectors for several $\boldsymbol{x} \in \mathbb{N}^2$. The connected line segments indicate points $\boldsymbol{x}$ with the same value of $V(\boldsymbol{x}) = \max\{\boldsymbol{x}\boldsymbol{q}^{(1)T}, \boldsymbol{x}\boldsymbol{q}^{(2)T}\} = \max\{\frac{5}{6}\boldsymbol{x}_1 + \frac{1}{6}\boldsymbol{x}_2, \frac{2}{7}\boldsymbol{x}_1 + \frac{5}{7}\boldsymbol{x}_2\}$ (values $0.5, 1, 1.5, \ldots$). It can be seen from the figure that the drift is negative with respect to the gradient of V, if $\boldsymbol{x} \neq \boldsymbol{0}$.

Proof of Lemma 8. Let $\boldsymbol{x} \neq \boldsymbol{0}$. We need to show $\boldsymbol{\Delta}(\boldsymbol{x})\boldsymbol{q}^{(i)T} \leq -\gamma$ for all i with $\boldsymbol{x}\boldsymbol{q}^{(i)T} = V(\boldsymbol{x})$. W.l.o.g. we assume that $\boldsymbol{x}\boldsymbol{q}^{(1)T} = V(\boldsymbol{x})$ and show only $\boldsymbol{\Delta}(\boldsymbol{x})\boldsymbol{q}^{(1)T} \leq -\gamma$. By Lemma 7 we have $\boldsymbol{x}_1 \neq 0$. It follows from the property $(I - A)A^* = I$ and the definition of $\boldsymbol{q}^{(1)}$ that we have

$$(-\boldsymbol{e}^{(1)} + A_1)\boldsymbol{q}^{(1)T} = -1/\left\|\boldsymbol{a}^{(1)}\right\| \qquad \text{and} \qquad (-\boldsymbol{e}^{(i)} + A_i)\boldsymbol{q}^{(1)T} = 0 \quad \text{for } 2 \leq i \leq n. \quad (4)$$

Hence we have:

$$\begin{aligned}
\boldsymbol{\Delta}(\boldsymbol{x})\boldsymbol{q}^{(1)T} &= \left(\boldsymbol{\alpha} + \sum_{i:\boldsymbol{x}_i \neq 0} \mu_i(-\boldsymbol{e}^{(i)} + A_i)\right)\boldsymbol{q}^{(1)T} && \text{by (3)} \\
&= \boldsymbol{\alpha}\boldsymbol{q}^{(1)T} - \mu_1/\left\|\boldsymbol{a}^{(1)}\right\| && \text{by (4) and } \boldsymbol{x}_1 \neq 0 \\
&\leq -\gamma + \boldsymbol{\alpha}\boldsymbol{q}^{(1)T} - \boldsymbol{\lambda}_1/\left\|\boldsymbol{a}^{(1)}\right\| && \text{by the definition of } \gamma \\
&= -\gamma + \left(\boldsymbol{\alpha} + \sum_{i=1}^{n} \boldsymbol{\lambda}_i(-\boldsymbol{e}^{(i)} + A_i)\right)\boldsymbol{q}^{(1)T} && \text{by (4)} \\
&= -\gamma + (\boldsymbol{\alpha} + \boldsymbol{\lambda}(-I + A))\boldsymbol{q}^{(1)T} \\
&= -\gamma + \boldsymbol{0}\boldsymbol{q}^{(1)T} = -\gamma && \text{by the traffic equation.}
\end{aligned}$$

◀

Using integration, one can obtain a two times continuously differentiable function $\widetilde{V}$ satisfying the conditions in Lemma 5: the function V is smoothed by defining $\widetilde{V}(\boldsymbol{x})$, for all $\boldsymbol{x}$, as an "average" of the values $V(\boldsymbol{y})$ where $\boldsymbol{y}$ belongs to a small ball around $\boldsymbol{x}$; see the appendix of [10] for the formal details. This concludes the proof of Proposition 4.

3.2 Controlled branching networks

In this subsection we generalize the traffic equations (2) to deal with an arbitrary controlled branching network $\mathcal{N}$. To obtain a distribution on actions for a static randomized ergodic scheduler, we assign variables to actions instead of queues, i.e., for every action ξ of the network we introduce a variable λ_ξ capturing the *rate of firing the action ξ*. Denote by $\bar{\Sigma}$ the set $\bigcup_{i=1}^n \Sigma_i$. Given $\zeta \in \bar{\Sigma}$ and $j \in \{1, \ldots, n\}$, we denote by $A_{\zeta j}$ the average number of jobs added to the queue j when the action ζ fires, i.e., for $\zeta \in \Sigma_i$ we set

$$A_{\zeta j} := \sum_{\boldsymbol{r} \in R_i(\zeta)} Prob_i(\zeta)(\boldsymbol{r}) \cdot \boldsymbol{r}_j \,.$$

We generalize (2) to the *traffic LP* presented in Figure 2, where the variable δ is intended to bound, for all j, the probability that queue j is busy.

min δ subject to

$$\sum_{\xi \in \Sigma_j} \lambda_\xi \;\; = \;\; \boldsymbol{\alpha}_j + \sum_{i=1}^n \sum_{\zeta \in \Sigma_i} \lambda_\zeta \cdot A_{\zeta j} \qquad\qquad j \in \{1, \ldots, n\}$$

$$\delta \;\; \geq \;\; \frac{\sum_{\xi \in \Sigma_j} \lambda_\xi}{\mu_j} \qquad\qquad\qquad\qquad j \in \{1, \ldots, n\}$$

$$\lambda_\xi \;\; \geq \;\; 0 \qquad\qquad\qquad\qquad\qquad\qquad \xi \in \bar{\Sigma}$$

■ **Figure 2** The traffic LP.

We prove the following

▶ **Proposition 10.**
1. *If there exists an arbitrary ergodic scheduler for $\mathcal{N}$, then the traffic LP can be solved with* min $\delta < 1$.
2. *If the traffic LP is solved with* min $\delta < 1$, *one can compute in polynomial time a static randomized ergodic scheduler Θ_s for $\mathcal{N}$. Moreover, denoting by ρ_i the* utilization $\lim_{t \to \infty} \Pr(\boldsymbol{x}_i(t) \neq 0)$ *of the queue i, the scheduler Θ_s minimizes* $\max_i \rho_i$ *among all memoryless ergodic schedulers.*
Hence one can decide in polynomial time whether an arbitrary ergodic scheduler exists; if yes, one can compute in polynomial time a static randomized ergodic scheduler.

Let us first concentrate on part 1. Let Θ be an ergodic scheduler. Roughly speaking, we prove that a feasible solution of the traffic LP can be constructed using (limit) frequencies of firing individual actions in $\mathcal{N}_\Theta$. Formally, given a run ω of $\mathcal{N}_\Theta$, $t \in \mathbb{R}_{\geq 0}$ and $\xi \in \bar{\Sigma}$, we denote by $O_\xi^{\leq t}(\omega)$ the number of times the action ξ is fired up to time t on ω. For memoryless Θ we have the following result.

▶ **Lemma 11.** *Assume that Θ is a* memoryless *ergodic scheduler. For every $\xi \in \bar{\Sigma}$ there is a constant O_ξ such that for almost all runs ω of $\mathcal{N}_\Theta$ the limit*

$$\lim_{t \to \infty} \frac{O_\xi^{\leq t}(\omega)}{t}$$

exists and is equal to O_ξ. There is $\bar{\delta} < 1$ such that $(\bar{\delta}, O_\xi \mid \xi \in \bar{\Sigma})$ solves the traffic LP. Moreover, for every $i \in \{1, \ldots, n\}$ the utilization ρ_i of the queue i in $\mathcal{N}_\Theta$ is equal to $\frac{\sum_{\xi \in \Sigma_i} O_\xi}{\mu_i}$.

We prove Lemma 11 in Appendix A.3. If there exists an arbitrary (i.e. possibly history-dependent) ergodic scheduler, then by Theorem 7.3.8 of [27] there exists also a memoryless (and deterministic) ergodic scheduler.[4] This fact, combined with Lemma 11, implies part 1. of Proposition 10.

Now let us concentrate on part 2. of Proposition 10.

▶ **Lemma 12.** *Any feasible solution $(\bar{\delta}, \bar{\lambda}_\xi \mid \xi \in \bar{\Sigma})$ of the traffic LP with $\bar{\delta} < 1$ induces a static randomized ergodic scheduler whose utilization of any queue i is equal to $\frac{\sum_{\xi \in \Sigma_i} \bar{\lambda}_\xi}{\mu_i}$.*

Proof. We construct a static randomized scheduler Θ which chooses an action $\xi \in \Sigma_i$ for the queue i with probability

$$P_\xi = \frac{\bar{\lambda}_\xi}{\sum_{\zeta \in \Sigma_i} \bar{\lambda}_\zeta} \quad , \quad \sum_{\zeta \in \Sigma_i} \bar{\lambda}_\zeta > 0 \,. \tag{5}$$

Otherwise, if $\sum_{\zeta \in \Sigma_i} \bar{\lambda}_\zeta = 0$, we may control the queue i arbitrarily because no jobs ever come to the queue. We further assume (w.l.o.g.) that such queues have been removed from the network, i.e., that P_ξ is defined using (5) for all $\xi \in \bar{\Sigma}$. Note that $\sum_{\xi \in \Sigma_i} P_\xi = 1$ for every $i \in \{1, \ldots, n\}$.

Fixing the scheduler Θ we obtain a purely stochastic branching network whose traffic equations are deficiently solvable. Formally, we define a new purely stochastic branching network $\mathcal{N}'$ with n queues with the same arrival rate, the same arrival production function and the same queue rates as $\mathcal{N}$. Further, $\mathcal{N}'$ has $R_i' = \bigcup_{\xi \in \Sigma_i} R_i(\xi)$ and the following production functions $Prob_i'$ associated to queues:

$$Prob_i'(\boldsymbol{r}) = \sum_{\xi \in \Sigma_i} P_\xi \cdot Prob_i(\xi)(\boldsymbol{r}) \quad , \quad \boldsymbol{r} \in R_i'$$

(Here we formally assume $Prob_i(\xi)(\boldsymbol{r}) = 0$ for $\boldsymbol{r} \notin R_\xi$.) The traffic equations (1) for $\mathcal{N}'$ have the following form:

$$\lambda_j \quad = \quad \alpha_j + \sum_{i=1}^{n} \lambda_i \cdot A_{ij}' \quad , \quad j \in \{1, \ldots, n\} \tag{6}$$

with

$$A_{ij}' := \sum_{\boldsymbol{r} \in R_i'} Prob_i'(\boldsymbol{r}) \cdot \boldsymbol{r}_j = \sum_{\xi \in \Sigma_i} P_\xi \sum_{\boldsymbol{r} \in R_i'} Prob_i(\xi)(\boldsymbol{r}) \cdot \boldsymbol{r}_j =$$

$$= \sum_{\xi \in \Sigma_i} \frac{\bar{\lambda}_\xi}{\sum_{\zeta \in \Sigma_i} \bar{\lambda}_\zeta} \sum_{\boldsymbol{r} \in R_i'} Prob_i(\xi)(\boldsymbol{r}) \cdot \boldsymbol{r}_j = \sum_{\xi \in \Sigma_i} \frac{\bar{\lambda}_\xi}{\sum_{\zeta \in \Sigma_i} \bar{\lambda}_\zeta} \cdot A_{\xi j} \,.$$

---

[4] To be formally correct, we apply Theorem 7.3.8 of [27] to the *embedded* discrete time MDP and obtain a scheduler which returns to the state $\boldsymbol{0}$ in finitely many steps (on average). As there are only finitely many rates in our system, this means that also the expected return time to $\boldsymbol{0}$ is finite.

Setting $\lambda_i := \sum_{\xi \in \Sigma_i} \bar{\lambda}_\xi$ for every $i \in \{1, \ldots, n\}$, we obtain $\lambda_i A'_{ij} = \sum_{\xi \in \Sigma_i} \bar{\lambda}_\xi A_{\xi j}$. If we put this equality into the first equation of the traffic LP, we see that $(\lambda_1, \ldots, \lambda_n)$ solves (6). Also, $\lambda_j < \mu_j$ for all $j \in \{1, \ldots, n\}$. Proposition 4 then implies that the scheduler Θ is ergodic.

Finally, let us concentrate on the utilization. Note that the utilization of any queue i is the same in $\mathcal{N}'$ as in $\mathcal{N}_\Theta$, so it suffices to concentrate on $\mathcal{N}'$. Observe that the matrix $I - A'$ is invertible by Lemma 6. This means that $(\lambda_1, \ldots, \lambda_n)$ is, in fact, the *unique* solution of (6). Then however, by Lemma 11, the utilization ρ_i of queue i in $\mathcal{N}'$ (and thus also in $\mathcal{N}_\Theta$) is equal to $\frac{\lambda_i}{\mu_i} = \frac{\sum_{\xi \in \Sigma_i} \bar{\lambda}_\xi}{\mu_i}$. ◄

To complete the proof of Proposition 10, we consider the problem of minimizing the maximal utilization $\max_i \rho_i$. Let Θ_s be a static randomized ergodic scheduler induced by a solution of the traffic LP in the sense of Lemma 12 (here we consider a solution which minimizes δ). Observe that the scheduler Θ_s minimizes $\max_i \rho_i$ among all schedulers induced by solutions of the traffic LP. However, by Lemmas 11 and 12, for every memoryless scheduler Θ there exists a static randomized scheduler induced by a solution of the traffic LP which has the same utilization of each queue as Θ. Thus Θ_s minimizes $\max_i \rho_i$ among all memoryless ergodic schedulers.

4 Conclusions

We have suggested and studied controlled branching networks, a queueing model which extends Jackson networks by nondeterministic and branching features as required to model parallel systems. Although much of the classical theory (such as product-form stationary distributions) no longer holds for controlled branching networks, we have shown that the traffic equations can be generalized. This enabled us to construct a suitable Lyapunov function which we have used to establish strong stability properties. We have shown for the controlled model that static randomized schedulers are sufficient to achieve those strong stability properties. Linear programming can be used to efficiently compute such a scheduler, which at the same time minimizes the maximal queue utilization.

Future work should include the investigation of more performance measures, e.g., the long-time average queue size. Can non-static schedulers help to minimize it?

References

1   M. Ahmadi and S. Wong. A performance model for network processor architectures in packet processing system. In *Proceedings of the 19th IASTED International Conference on Parallel and Distributed Computing and Systems*, pages 176–181. ACTA Press, 2007.

2   A. Azaron and S.M.T. Fatemi Ghomi. Optimal control of service rates and arrivals in Jackson networks. *European Journal of Operational Research*, 147(1):17–31, 2003.

3   A. Berman and R.J. Plemmons. *Nonnegative matrices in the mathematical sciences*. Academic Press, 1979.

4   P. Billingsley. *Probability and Measure*. Wiley, 1995.

5   G. Bolch, S. Greiner, H. de Meer, and K.S. Trivedi. *Queueing Networks and Markov Chains*. John Wiley and Sons, 2006.

6   T. Brázdil, S. Kiefer, A. Kučera, and I. Hutařová Vařeková. Runtime analysis of probabilistic programs with unbounded recursion. In *Proceedings of ICALP*, volume 6756 of *LNCS*, pages 319–331, 2011.

7   H. Chen and D. Yao. *Fundamentals of Queueing Networks*. Springer, 2001.

8 J. Daigle. *Queueing Theory with Applications to Packet Telecommunication.* Springer, 2010.

9 J.D. Deng and M.K. Purvis. Multi-core application performance optimization using a constrained tandem queueing model. *Journal of Network and Computer Applications*, In Press, Corrected Proof, 2011. DOI: 10.1016/j.jnca.2011.07.004.

10 D. Down and S.P. Meyn. Piecewise linear test functions for stability and instability of queueing networks. *Queueing Systems*, 27:205–226, 1997.

11 J. Esparza, A. Kučera, and R. Mayr. Model checking probabilistic pushdown automata. In *LICS 2004*, pages 12–21. IEEE, 2004.

12 K. Etessami, D. Wojtczak, and M. Yannakakis. Recursive stochastic games with positive rewards. In *Proceedings of ICALP 2008*, pages 711–723, 2008.

13 K. Etessami and M. Yannakakis. Recursive Markov chains, stochastic grammars, and monotone systems of nonlinear equations. *Journal of the ACM*, 56(1):1–66, 2009.

14 K.D. Glazebrook and J. Niño-Mora. A linear programming approach to stability, optimisation and performance analysis for Markovian multiclass queueing networks. *Annals of Operations Research*, 92:1–18, 1999.

15 J.M. Harrison and R.J. Williams. Brownian models of feedforward queueing networks: Quasireversibility and product form solutions. *Annals of Applied Probability*, 2(2):263–293, 1992.

16 J.R. Jackson. Networks of waiting lines. *Operations Research*, 5(4):518–521, 1957.

17 S. Kiefer and D. Wojtczak. On probabilistic parallel programs with process creation and synchronisation. In *Proceedings of TACAS*, volume 6605 of *LNCS*, pages 296–310. Springer, 2011.

18 M.Y. Kitaev and V.V.Rykov. *Controlled queueing systems.* CRC Press, 1995.

19 P.R. Kumar and S.P. Meyn. Duality and linear programs for stability and performance analysis of queueing networks and scheduling policies. *IEEE Transactions on Automatic Control*, 42(1):4–17, 1996.

20 M. Kwiatkowska, G. Norman, and D. Parker. Prism 4.0: Verification of probabilistic real-time systems. In *Proceedings of CAV*, volume 6806 of *LNCS*, pages 585–591, 2011.

21 N. Madan, A. Buyuktosunoglu, P. Bose, and M. Annavaram. A case for guarded power gating for multi-core processors. In *High Performance Computer Architecture (HPCA)*, pages 291–300, 2011.

22 W.A. Massey and R. Srinivasan. A heavy traffic analysis for semi-open networks. *Performance Evaluation*, 13(1):59–66, 1991.

23 S.P. Meyn and R.L. Tweedie. *Markov Chains and Stochastic Stability.* Springer, 1993.

24 M. Neuhäußer, M. Stoelinga, and J.-P. Katoen. Delayed nondeterminism in continuous-time Markov decision processes. In *Proceedings of FoSSaCS 2009*, volume 5504 of *LNCS*, pages 364–379. Springer, 2009.

25 J.R. Norris. *Markov Chains.* Oxford University Press, first edition, 1998.

26 C.H. Papadimitriou and J.N. Tsitsiklis. The complexity of optimal queueing network control. *Mathematics of Operations Research*, 24:293–305, 1994.

27 M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming.* John Wiley and Sons, 2008.

28 S. Stidham, Jr. Optimal control of admission to a queueing system. *IEEE Transactions on Automatic Control*, 30(8):705–713, 1985.

29 S. La Torre, P. Madhusudan, and G. Parlato. Context-bounded analysis of concurrent queue systems. In *Proceedings of TACAS*, volume 4963 of *LNCS*, pages 299–314, 2008.

30 H. Zisgen, I. Meents, B.R. Wheeler, and T. Hanschke. A queueing network based system to model capacity and cycle time for semiconductor fabrication. In *Proceedings of the 40th Conference on Winter Simulation*, pages 2067–2074, 2008.

A   Proofs

A.1   Proof of Lemma 6

Lemma 6.   *The matrix series $A^* := \sum_{i=0}^{\infty} A^i$ converges ("exists") in $\mathbb{R}_{\geq 0}^{n \times n}$ and is equal to $(I - A)^{-1}$.*

Proof. By the traffic equation $\boldsymbol{\lambda} = \boldsymbol{\alpha} + \boldsymbol{\lambda} A$ we have $\boldsymbol{\lambda} A \leq \boldsymbol{\lambda}$, with $\boldsymbol{\lambda} > \boldsymbol{0}$. By [3, Theorem 2.1.11] this implies that the spectral radius of A is at most 1 (where the spectral radius is the largest absolute value of the eigenvalues of A).

To show the statement of the lemma, it now suffices to show that $I - A$ is invertible. Consider the (monotone) function $\boldsymbol{f} : \mathbb{R}_{\geq 0}^n \to \mathbb{R}_{\geq 0}^n$ with $\boldsymbol{f}(\boldsymbol{x}) := \boldsymbol{\alpha} + \boldsymbol{x} A$. By the traffic equation, $\boldsymbol{\lambda}$ is a fixed point of $\boldsymbol{f}$. Assume for a contradiction that $I - A$ is singular, i.e., there is $\boldsymbol{y} \in \mathbb{R}^n$ with $\boldsymbol{y} \neq \boldsymbol{0}$ and $\boldsymbol{y} A = \boldsymbol{y}$. Then $\boldsymbol{\lambda} + r\boldsymbol{y}$ is a fixed point for all $r \in \mathbb{R}$. Choose r so that $\boldsymbol{z} := \boldsymbol{\lambda} + r\boldsymbol{y} \in \mathbb{R}_{\geq 0}^n$, but $\boldsymbol{z}_i = 0$ for some $i \in \{1, \ldots, n\}$. By the monotonicity of $\boldsymbol{f}$, all points $\boldsymbol{u} \in \{\boldsymbol{0}, \boldsymbol{f}(\boldsymbol{0}), \boldsymbol{f}(\boldsymbol{f}(\boldsymbol{0})), \ldots\}$ satisfy $\boldsymbol{0} \leq \boldsymbol{u} \leq \boldsymbol{z} = \boldsymbol{f}(\boldsymbol{z})$. It follows that $\boldsymbol{u}_i = 0$ for all $\boldsymbol{u}$. This contradicts the fact that queue i is "reachable" (recall that "queue i is reachable" means there is $j \in \mathbb{N}$ with $(\boldsymbol{\alpha} A^j)_i \neq 0$). ◄

A.2   Proof of Proposition 4

In this section we complete the proof of

Proposition 4.   *Assume that $\boldsymbol{\lambda} \in \mathbb{R}_{\geq 0}^n$ solves the traffic equations (2) and satisfies $\boldsymbol{\lambda} < \boldsymbol{\mu}$. Then the following conclusions hold:*

1. *The process $\mathcal{N}$ is ergodic, i.e., the expected return time to $\boldsymbol{0}$ is finite.*
2. *There exists a stationary distribution π such that there exists an exponential moment of the total queue size, i.e., there is $\delta > 0$ such that $\sum_{\boldsymbol{x} \in \mathbb{N}^n} \exp(\delta \|\boldsymbol{x}\|) \pi(\boldsymbol{x})$ exists.*

Recall from Lemma 6 that the matrix series $A^* := \sum_{i=0}^{\infty} A^i$ converges ("exists") in $\mathbb{R}_{\geq 0}^{n \times n}$ and equals $(I - A)^{-1}$. Observe that for any vector $\boldsymbol{v}^T = A^* \boldsymbol{w}^T$ with $\boldsymbol{w} \in \mathbb{R}^n$, we have $\boldsymbol{v}^T = A\boldsymbol{v}^T + \boldsymbol{w}^T$. We prove Lemma 7:

Lemma 7.   *If $\boldsymbol{0} \neq \boldsymbol{x} \in \mathbb{R}_{\geq 0}^n$ and $\boldsymbol{x}_i = 0$, then $\boldsymbol{x} \boldsymbol{q}^{(i)T} < V(\boldsymbol{x})$.*

Proof. For the proof we use the following notation: For a matrix M we denote by $M_{i..j}$ the matrix obtained by restricting M to its rows indexed by $i, \ldots, j$. Similarly, we denote by $M_{i..j,k..\ell}$ the matrix M restricted to the corresponding rows and columns.

Let $Q \in \mathbb{R}_{\geq 0}^{n \times n}$ denote the matrix whose columns are $\boldsymbol{q}^{(1)T}, \ldots, \boldsymbol{q}^{(n)T}$. Note that $\boldsymbol{1}Q = \boldsymbol{1}$. We have $Q = A^* D$ where $D \in \mathbb{R}_{\geq 0}^{n \times n}$ is the diagonal matrix with $1/D_{ii} = (\boldsymbol{1}A^*)_i$. Further, we define the matrix $P \in \{-1, 0, 1\}^{n \times (n-1)}$ with

$$
P := \begin{pmatrix} -1 & -1 & -1 & \cdots & -1 & -1 \\ 1 & 0 & 0 & \cdots & 0 & 0 \\ 0 & 1 & 0 & \cdots & 0 & 0 \\ & & & \ddots & & \\ 0 & 0 & 0 & \cdots & 0 & 1 \end{pmatrix}.
$$

W.l.o.g. let $\boldsymbol{x}_1 = 0$. With the notation above and writing $\bar{\boldsymbol{x}} := (x_2, \ldots, x_n) \in \mathbb{R}_{\geq 0}^{n-1}$, the lemma states that

there is no $\bar{\boldsymbol{x}} \in \mathbb{R}_{\geq 0}^{n-1}$ with $\bar{\boldsymbol{x}} \neq \boldsymbol{0}$ and $\bar{\boldsymbol{x}}Q_{2..n}P \leq \boldsymbol{0}$. $\hspace{2cm}$ (7)

We use Farkas' lemma for the proof:

▶ **Lemma 13** (Farkas' lemma). *Let $M \in \mathbb{R}^{m \times n}$ and $\boldsymbol{b} \in \mathbb{R}^m$. Then exactly one of the following is true:*
1. *There exists $\boldsymbol{x} \in \mathbb{R}^m$ with $\boldsymbol{x}M \leq \boldsymbol{0}$ and $\boldsymbol{x}\boldsymbol{b}^T > 0$.*
2. *There exists $\boldsymbol{y} \in \mathbb{R}^n$ with $M\boldsymbol{y}^T = \boldsymbol{b}^T$ and $\boldsymbol{y} \geq \boldsymbol{0}$.*

Using Farkas' lemma we prove (7) by exhibiting $\boldsymbol{b} > \boldsymbol{0}$ and $\boldsymbol{y} \geq \boldsymbol{0}$ such that $Q_{2..n}P\boldsymbol{y}^T = \boldsymbol{b}^T$. We choose

$$\boldsymbol{b}^T := (A_{2..n,2..n})^* \boldsymbol{1}^T \geq \boldsymbol{1}^T > \boldsymbol{0}^T .$$

We have $\boldsymbol{1}Q = \boldsymbol{1}$, which implies $\boldsymbol{1}Q^{-1} = \boldsymbol{1}QQ^{-1} = \boldsymbol{1}$, hence $P(Q^{-1})_{2..n} = Q^{-1} - E$ where $E \in \{0,1\}^{n \times n}$ denotes the matrix with $\boldsymbol{1}$ in the first row and $\boldsymbol{0}$ in the other rows. It follows

$$P(Q^{-1})_{2..n}P = Q^{-1}P . \hspace{2cm} (8)$$

Setting

$$\boldsymbol{y}^T := (Q^{-1})_{2..n}P\boldsymbol{b}^T$$

we therefore have

$$Q_{2..n}P\boldsymbol{y}^T = Q_{2..n}P(Q^{-1})_{2..n}P\boldsymbol{b}^T \overset{(8)}{=} Q_{2..n}Q^{-1}P\boldsymbol{b}^T = P_{2..n}\boldsymbol{b}^T = \boldsymbol{b}^T ,$$

as desired. It remains to prove $\boldsymbol{y} \geq \boldsymbol{0}$. As $Q = A^*D$, we have $Q^{-1} = D^{-1}(I - A)$ with D^{-1} nonnegative, so it suffices to prove that $(I - A)_{2..n}P\boldsymbol{b}^T \geq \boldsymbol{0}^T$. Indeed we have

$$(I - A)_{2..n}P\boldsymbol{b}^T = (I - A_{2..n}P)\boldsymbol{b}^T = \boldsymbol{b}^T - A_{2..n}P\boldsymbol{b}^T \geq \boldsymbol{b}^T - A_{2..n,2..n}\boldsymbol{b}^T = \boldsymbol{1}^T ,$$

where the inequality holds as $\boldsymbol{b}^T$ and the first column of $A_{2..n}$ are nonnegative and the first row of P is negative.

$\hspace{13cm}$ ◀

Recall the following lemma:

Lemma 8. *There is $\gamma > 0$ such that we have*

$$\boldsymbol{\Delta}(\boldsymbol{x})\bigl(V'(\boldsymbol{x})\bigr)^T \leq -\gamma \hspace{1cm} \text{for all } \boldsymbol{x} \neq \boldsymbol{0}$$

and all subgradient vectors $V'(\boldsymbol{x})$ of V at $\boldsymbol{x}$. More precisely, one can choose

$$\gamma := \min_i (\mu_i - \boldsymbol{\lambda}_i) / \left\| \boldsymbol{a}^{(i)} \right\| ,$$

where $\boldsymbol{a}^{(i)T}$ is the ith column of A^*.

Proposition 4 follows from Lemma 8 using exactly the reasoning from theorem 1 and lemma 5 of [10].

A.3 Proof of Lemma 11

Lemma 11. *Assume that Θ is a memoryless ergodic scheduler. For every $\xi \in \bar{\Sigma}$ there is a constant O_ξ such that for almost all runs ω of $\mathcal{N}_\Theta$ the limit*

$$\lim_{t \to \infty} \frac{O_\xi^{\leq t}(\omega)}{t}$$

exists and is equal to O_ξ. There is $\bar{\delta} < 1$ such that $\left(\bar{\delta}, O_\xi \mid \xi \in \bar{\Sigma} \right)$ solves the traffic LP. Moreover, for every $i \in \{1, \ldots, n\}$ the utilization ρ_i of the queue i in $\mathcal{N}_\Theta$ is equal to $\frac{\sum_{\xi \in \Sigma_i} O_\xi}{\mu_i}$.

We divide our proof of Lemma 11 into four claims.

▶ Claim 1. For every $\xi \in \bar{\Sigma}$ there is a constant O_ξ such that for almost all runs ω of $\mathcal{N}_\Theta$ the limit

$$\lim_{t \to \infty} \frac{O_\xi^{\leq t}(\omega)}{t}$$

exists and is equal to O_ξ.

Proof. Given an action $\xi \in \bar{\Sigma}$ and $i \geq 1$, denote by M_i^ξ a random variable giving the number of times the action ξ is fired between the i-th and $i+1$-st visit to $\mathbf{0}$ (note that the network starts with all empty queues so the first visit happens at time 0). Observe that all M_i^ξ are identically distributed. Also, as the expected return time to $\mathbf{0}$ is finite and there are only finitely many distinct rates (i.e. the expected time between two state transitions is bounded from below by a positive constant), the expectation $\mathbb{E}\left[M_1^\xi \right]$ is finite. By the strong law of large numbers (see e.g. [4]), there is a constant M^ξ such that for almost all runs ω the limit $\lim_{n \to \infty} \frac{\sum_{i=1}^n M_i^\xi(\omega)}{n}$ exists and is equal to M^ξ. For $i \geq 1$ we define R_i to be a random variable giving the total time between the i-th and $i+1$-st visits to $\mathbf{0}$. Note that R_1 is precisely the variable R giving the return time to $\mathbf{0}$. By the strong law of large numbers, for almost all runs ω the limit $\lim_{n \to \infty} \frac{\sum_{i=1}^n R_i(\omega)}{n}$ exists and is equal to $\mathbb{E}\left[R_1 \right] = \mathbb{E}\left[R \right] < \infty$.

Now note that for $t \geq 0$ and $n \geq 1$ satisfying $\sum_{i=1}^{n-1} R_i(\omega) \leq t \leq \sum_{i=1}^n R_i(\omega)$ (here for $n = 1$ we assume $\sum_{i=1}^{n-1} R_i(\omega) = 0$) we have

$$\frac{\sum_{i=1}^{n-1} M_i^\xi(\omega)}{\sum_{i=1}^n R_i(\omega)} \quad \leq \quad \frac{O_\xi^{\leq t}(\omega)}{t} \quad \leq \quad \frac{\sum_{i=1}^n M_i^\xi(\omega)}{\sum_{i=1}^n R_i(\omega)} \tag{9}$$

The limit of the rightmost term of (9) exists:

$$\frac{M^\xi}{\mathbb{E}\left[R \right]} \quad = \quad \lim_{n \to \infty} \frac{\sum_{i=1}^n M_i^\xi(\omega)}{n} \cdot \lim_{n \to \infty} \frac{n}{\sum_{i=1}^n R_i(\omega)} \quad = \quad \lim_{n \to \infty} \frac{\sum_{i=1}^n M_i^\xi(\omega)}{\sum_{i=1}^n R_i(\omega)}$$

Similarly, the limit of the left most term of (9) exists:

$$\frac{M^\xi}{\mathbb{E}\left[R \right]} \quad = \quad \left(\lim_{n \to \infty} \frac{\sum_{i=1}^{n-1} M_i^\xi(\omega)}{n-1} \cdot \lim_{n \to \infty} \frac{n-1}{n} \right) \cdot \lim_{n \to \infty} \frac{n}{\sum_{i=1}^n R_i(\omega)} \quad = \quad \lim_{n \to \infty} \frac{\sum_{i=1}^n M_i^\xi(\omega)}{\sum_{i=1}^n R_i(\omega)}$$

Thus by basic properties of limits, the limit $\lim_{t \to \infty} \frac{O_\xi^{\leq t}(\omega)}{t}$ exists and is equal to $\frac{M^\xi}{\mathbb{E}[R]}$. ◀

The following two claims show that the frequencies indeed solve the traffic LP with $\delta < 1$.

▶ Claim 2. For every $j \in \{1, \ldots, n\}$ we have

$$\sum_{\xi \in \Sigma_j} O_\xi = \sum_{i=0}^n \sum_{\zeta \in \Sigma_i} O_\zeta \cdot A_{\zeta j} \tag{10}$$

Here we assume (in order to simplify notation) that the external source is a (symbolic) queue 0, with action set $\Sigma = \{\iota\}$ where ι has rate μ_0 and a production function $Prob_0$ (thus $O_\iota \cdot A_{\iota j} = \boldsymbol{\alpha}_j$).

Proof. Denote by $O_{\zeta \to j}^{\le t}(\omega)$ the number of jobs produced for queue j by firing the action ζ. The crucial observation is that for almost all runs ω we have

$$\lim_{t \to \infty} \frac{\sum_{\xi \in \Sigma_j} O_\xi^{\le t}(\omega)}{t} \quad = \quad \lim_{t \to \infty} \frac{\sum_{i=0}^n \sum_{\zeta \in \Sigma_i} O_{\zeta \to j}^{\le t}(\omega)}{t}$$

Indeed, $\sum_{\xi \in \Sigma_j} O_\xi^{\le t}(\omega) \le \sum_{i=0}^n \sum_{\zeta \in \Sigma_i} O_{\zeta \to j}^{\le t}(\omega)$ for all t and ω because $\sum_{i=0}^n \sum_{\zeta \in \Sigma_i} O_{\zeta \to j}^{\le t}(\omega)$ is precisely the number of jobs that enter the queue j up to time t. On the other hand, $\sum_{i=0}^n \sum_{\zeta \in \Sigma_i} O_{\zeta \to j}^{\le t}(\omega) \le \sum_{\xi \in \Sigma_j} O_\xi^{\le t}(\omega)$ for infinitely many t and almost all ω because (almost surely) queue j becomes empty infinitely many times.

We obtain that

$$\begin{aligned}
\sum_{\xi \in \Sigma_j} O_\xi &= \sum_{\xi \in \Sigma_j} \lim_{t \to \infty} \frac{O_\xi^{\le t}(\omega)}{t} \\
&= \lim_{t \to \infty} \frac{\sum_{\xi \in \Sigma_j} O_\xi^{\le t}(\omega)}{t} \\
&= \lim_{t \to \infty} \frac{\sum_{i=0}^n \sum_{\zeta \in \Sigma_i} O_{\zeta \to j}^{\le t}(\omega)}{t} \\
&= \sum_{i=0}^n \sum_{\zeta \in \Sigma_i} \lim_{t \to \infty} \frac{O_{\zeta \to j}^{\le t}(\omega)}{t} \\
&= \sum_{i=0}^n \sum_{\zeta \in \Sigma_i} \lim_{t \to \infty} \frac{O_\zeta^{\le t}(\omega)}{t} \cdot \lim_{t \to \infty} \frac{O_{\zeta \to j}^{\le t}(\omega)}{O_\zeta^{\le t}(\omega)} \\
&= \sum_{i=0}^n \sum_{\zeta \in \Sigma_i} O_\zeta \cdot A_{\zeta j}
\end{aligned}$$

Here the last equality follows from Claim 1 and the fact that $\lim_{t \to \infty} \frac{O_{\zeta \to j}^{\le t}(\omega)}{O_\zeta^{\le t}(\omega)}$ is the average number of jobs for queue j produced by firing ζ, i.e., $A_{\zeta j}$.   ◀

▶ Claim 3. For every $i \in \{1, \ldots, n\}$ we have

$$\frac{\sum_{\xi \in \Sigma_i} O_\xi}{\mu_i} < 1 \tag{11}$$

Proof. Intuitively, we show that $\frac{\sum_{\xi \in \Sigma_i} O_\xi}{\mu_i} + \pi(\mathbf{0})$ is equal to the proportion of time in which either the queue i operates a job or all queues are empty. As this proportion is at most 1 and $\pi(\mathbf{0}) > 0$, we obtain the desired result.

Formally, denote by T_n the total amount of time in which the first n jobs are processed by queue i (by this we mean the total time from the beginning of the computation needed to

finish n jobs using the "machine" operating queue i). Denote by N_j the total time it takes to process the j-th job from queue i (by this we mean the time the "machine" actively spends by processing the job).

Observe that for almost all runs ω we have that $\lim_{n \to \infty} \frac{\sum_{j=1}^{n} N_j(\omega)}{n}$ is the average processing time of the queue i which is $\frac{1}{\mu_i}$, and that $\frac{n}{T_n(\omega)}$ is equal to

$$\frac{\sum_{\xi \in \Sigma_i} O_\xi^{\leq T_n(\omega)}}{T_n(\omega)}$$

which means that by Claim 1,

$$\lim_{n \to \infty} \frac{n}{T_n(\omega)} \quad = \quad \lim_{n \to \infty} \frac{\sum_{\xi \in \Sigma_i} O_\xi^{\leq T_n(\omega)}}{T_n(\omega)} \quad = \quad \sum_{\xi \in \Sigma_i} O_\xi$$

Thus

$$\lim_{n \to \infty} \frac{\sum_{j=1}^{n} N_j(\omega)}{T_n(\omega)} \quad = \quad \lim_{n \to \infty} \frac{\sum_{j=1}^{n} N_j(\omega)}{n} \cdot \lim_{n \to \infty} \frac{n}{T_n(\omega)} \quad = \quad \frac{\sum_{\xi \in \Sigma_i} O_\xi}{\mu_i} \tag{12}$$

Denote by E_n the total amount of time the network spends in state $\mathbf{0}$ before queue i finishes n jobs (by this we mean the sum of all time intervals up to time T_n in which all queues are simultaneously idle before n jobs are finished by the "machine" operating queue i). Intuitively, $\lim_{n \to \infty} \frac{E_n(\omega)}{T_n(\omega)}$ is equal to the proportion of time spent in $\mathbf{0}$ which is equal to $\pi(\mathbf{0}) > 0$. More precisely, by the Ergodic Theorem, see Theorem 3.8.1 of [25], for almost all runs ω we have

$$0 \quad < \quad \pi(\mathbf{0}) \quad = \quad \lim_{t \to \infty} \frac{1}{t} \int_0^t I[\mathbf{x}(\omega) = \mathbf{0}](s) ds \quad = \quad \lim_{n \to \infty} \frac{E_n(\omega)}{T_n(\omega)} \tag{13}$$

Finally, by (12) and (13),

$$\frac{\sum_{\xi \in \Sigma_i} O_\xi}{\mu_i} \quad < \quad \frac{\sum_{\xi \in \Sigma_i} O_\xi}{\mu_i} + \pi(\mathbf{0}) \quad =$$

$$\frac{\sum_{j=1}^{n} N_j(\omega)}{n} \cdot \lim_{n \to \infty} \frac{n}{T_n(\omega)} + \lim_{n \to \infty} \frac{E_n(\omega)}{T_n(\omega)} \quad = \quad \lim_{n \to \infty} \frac{\sum_{j=1}^{n} N_j(\omega) + E_n(\omega)}{T_n(\omega)} \quad \leq \quad 1$$

◄

▶ **Claim 4.** For every $i \in \{1, \ldots, n\}$, $\rho_i = \frac{\sum_{\xi \in \Sigma_i} O_\xi}{\mu_i}$.

Proof. By the Ergodic Theorem, see Theorem 3.8.1 of [25], and by (12), for almost all runs ω we have

$$\rho_i \quad = \quad \lim_{t \to \infty} \Pr(\mathbf{x}_i(t) \neq 0) \quad = \quad \sum_{\mathbf{x} \in \mathbb{N}^n : \mathbf{x}_i \neq 0} \pi(\mathbf{x})$$

$$= \quad \lim_{t \to \infty} \frac{1}{t} \int_0^t I[\mathbf{x}_i(\omega) \neq 0](s) ds \quad = \quad \lim_{n \to \infty} \frac{\sum_{j=1}^{n} N_j(\omega)}{T_n(\omega)} \quad = \quad \frac{\sum_{\xi \in \Sigma_i} O_\xi}{\mu_i}$$

◄

This finishes a proof of Lemma 11.