

Stability of Parallel Queueing Systems with Coupled Service Rates

Sem Borst · Matthieu Jonckheere · Lasse Leskelä

Received: 2 February 2007 / Accepted: 26 July 2007 /
Published online: 11 September 2007
© Springer Science + Business Media, LLC 2007

Abstract This paper considers a parallel system of queues fed by independent arrival streams, where the service rate of each queue depends on the number of customers in all of the queues. Necessary and sufficient conditions for the stability of the system are derived, based on stochastic monotonicity and marginal drift properties of multiclass birth and death processes. These conditions yield a sharp characterization of stability for systems where the service rate of each queue is decreasing in the number of customers in other queues, and has uniform limits as the queue lengths tend to infinity. The results are illustrated with applications where the stability region may be nonconvex.

Keywords Stability · Positive recurrence · Multiclass birth and death process · Foster–Lyapunov drift criterion · Stochastic comparison · Coupled processors

1 Introduction

We consider a parallel system of queues fed by independent arrival streams, where the service rate of each queue depends on the number of customers in all of the queues. This type of model is natural for manufacturing systems where a server is capable to process other queues when its own buffer is empty, or for cellular radio networks, where the available transmission rate for customers in a particular cell is

S. Borst · L. Leskelä
Department of Mathematics and Computer Science, Eindhoven University of Technology,
P.O. Box 513, 5600 MB Eindhoven, The Netherlands
e-mail: sem@win.tue.nl

M. Jonckheere (✉)
Kruislaan 413, Amsterdam 1098SJ, The Netherlands
e-mail: matthieu.jonckheere@cwi.nl

L. Leskelä
e-mail: lasse.leskela@iki.fi

decreasing in the number of customers in the neighboring cells (Bonald et al. 2004). Another important category of applications consists of processor sharing models, where several customer classes simultaneously use one or more servers, whose rate allocations and total processing rates may depend on the number of customers in each of the classes (Bonald et al. 2006). For example, in wireless data networks employing channel-aware scheduling, the total service rate available to all customers can be increasing in the total number of customers, due to multiuser diversity (Liu et al. 2003).

Stability is arguably the most fundamental property of a queueing system, and provides a crude yet useful first-order benchmark of the system performance. A general framework for analyzing stochastic stability consists of the Foster–Lyapunov criteria, which are based on finding a suitable test function having a positive or negative mean drift in almost all states of the state space (Meyn and Tweedie 1993; Fayolle et al. 1995). In the context of multiclass queueing systems with coupled servers, these techniques have been successfully applied to systems with utility-based service allocations (de Veciana et al. 2001). Fluid limit analysis is another powerful method for finding necessary and sufficient stability conditions for a wide class of multiclass queueing networks with work-conserving service disciplines (Dai 1995; Meyn 1995).

The stability analysis of multiclass queueing systems with general state-dependent service rates is difficult, because there is no systematic way of finding test functions satisfying the Foster–Lyapunov criteria, and the fluid limit techniques are often restricted to systems of work-conserving servers with fixed total rate. An alternative means for deriving stability conditions is to study whether the system of interest is stochastically comparable to a simpler system that is easier to analyze. This approach was first used in the multiclass queueing context by Rao and Ephremides (1988) and Szpankowski (1988), and later refined by Szpankowski (1994), to characterize the stability of buffered random access systems.

In this paper we provide an extension of the above ideas (tentatively discussed in Jonckheere and Borst 2006), by deriving marginal drift criteria for multiclass birth and death processes that allow us to analyze the stability of a broad class of parallel queueing systems. Moreover, we present conditions for partial stability, where only some of the queues are stable, and give a sharp stability characterization for systems where the service rate of each queue is decreasing in the number of customers in other queues, and has uniform limits as the queue lengths tend to infinity. For systems of at most three queues, where the service rates only depend on whether the queues are empty or not, our results yield as special cases stability characterizations that have earlier been found using transform methods (Fayolle and Iasnogorodski 1979; Cohen and Boxma 1983) and the ergodic theory of deflected random walks (Fayolle et al. 1995).

The paper is organized as follows. Section 2 describes the model details and discusses a notion of stability that is convenient for the subsequent analysis. Section 3 presents a key coupling result and marginal drift criteria for the stability of multiclass birth and death processes, while the main results regarding the stability of queueing systems are given in Section 4, in decreasing level of generality. Section 5 illustrates the results with two applications, and Section 6 concludes the paper.

2 Model description

2.1 Parallel queueing system with coupled service rates

We consider a parallel system of N queues, where each queue i is fed by an independent Poisson arrival process of rate λ_i , and served at rate $\phi_i(X_1, \dots, X_N)$ that depends on the number of customers X_j in each of the queues $j = 1, \dots, N$. We assume that all customers require independent exponentially distributed amounts of service with unit mean, and that the system has unlimited buffer space to accommodate customers. The scheduling at each queue can be first-come first-served, processor sharing, random order of service, or any discipline that does not depend on the service requirements.

Under these assumptions, we can model $X = (X_1, \dots, X_N)$ as a continuous-time Markov process on \mathbb{Z}_+^N , with transitions $x \mapsto x + e_i$ occurring at rate λ_i and transitions $x \mapsto x - e_i \geq 0$ at rate $\phi_i(x)$, where e_i denotes the i -th unit vector in \mathbb{Z}_+^N . We assume that the allocation function $\phi = (\phi_1, \dots, \phi_N)$ is bounded, which guarantees that the process X is nonexplosive. Hence we may assume that X and all other stochastic processes treated in the sequel have paths in the space $D = D(\mathbb{R}_+, \mathbb{Z}_+^N)$ of right-continuous functions from \mathbb{R}_+ to \mathbb{Z}_+^N with finite left limits. Recall that a stochastic process with paths in D can be viewed as a random element on the measurable space (D, \mathcal{D}) , where \mathcal{D} denotes the Borel σ -algebra generated by the standard Skorokhod topology (Kallenberg 2002).

Observe that this model also covers scenarios where the service requirements of customers at queue i are exponentially distributed with parameter $\mu_i \neq 1$, via replacing $\phi_i(x)$ by $\mu_i \phi_i(x)$.

2.2 Stability notions

A stochastic process X taking values in a countable space S and having paths in $D(\mathbb{R}_+, S)$ is said to be *stable*, if for any $\epsilon > 0$ there exists a finite set K such that

$$P(X(t) \notin K) \leq \epsilon \quad \text{for all } t, \tag{2.1}$$

and otherwise X is said to be *unstable*. Further, the process X is called *transient*, if $X(t) \rightarrow \infty$ almost surely, that is, for any finite set K ,

$$P\left(\bigcup_{s \geq 0} \bigcap_{t \geq s} \{X(t) \notin K\}\right) = 1.$$

An alternative way to express Eq. 2.1 is to say that the family of distributions $\{P \circ X(t)^{-1}\}_{t \geq 0}$ is tight. Observe that an irreducible Markov process is stable if and only if it is positive recurrent (Kallenberg 2002, Theorem 12.25). The following proposition illustrates an intuitively clear relation between transience and instability.

Proposition 1 *Any transient stochastic process X having paths in $D(\mathbb{R}_+, S)$ is unstable.*

Proof If X is transient, then for any finite set K there exists an s such that $P(\cap_{t \geq s} \{X(t) \notin K\}) > 1/2$. Hence, $\sup_t P(X(t) \notin K) > 1/2$ for all finite K , so X cannot be stable. □

In most applications, it is natural to assume that the Markov process describing the system is irreducible, in which case stability is equivalent to the existence of a unique stationary distribution. In Section 4, where the service rates of the original system are modified in various ways, it may happen that some of the modified Markov processes are not irreducible. This is why we need the following result to guarantee the existence of a stationary distribution for a stable multiclass birth and death process under slightly weaker than usual assumptions on the reachability of states. We denote by $X[x]$ the version of a Markov process X started in state x .

Proposition 2 *Let X be a N -class birth and death process with strictly positive birth rates λ_i and bounded death rates $\phi_i(x)$. Then the following are equivalent:*

1. $X[x]$ is stable for some initial state x .
2. $X[x]$ is stable for all initial states x .
3. X has a unique stationary distribution π supported on a set C such that $P_x(X(t) \in \cdot) \rightarrow \pi$ in total variation for all initial states x , and $X[x]$ is irreducible and positive recurrent for any $x \in C$.

Moreover, X is unstable if and only if $X(t) \rightarrow \infty$ in probability, regardless of the initial state.

Proof See [Appendix](#). □

The following stability characterization of vector-valued stochastic processes is well-known. Because the proof is short, we give it here for completeness.

Proposition 3 *A stochastic process $X = (X_1, \dots, X_N)$ taking values in a countable space $S_1 \times \dots \times S_N$ is stable if and only if X_i is stable for each i .*

Proof First assume that X is stable. Given $\epsilon > 0$, let us fix a finite set K such that $\sup_t P(X(t) \notin K) \leq \epsilon$, and choose a finite rectangle $K_1 \times \dots \times K_N$ that contains K . Then for any i and all t , $P(X_i(t) \notin K_i) \leq P(X(t) \notin K)$, so it follows that X_i is stable.

For the reverse direction, it suffices to observe that for an arbitrary finite set $K = K_1 \times \dots \times K_N$, $P(X(t) \notin K) \leq \sum_i P(X_i \notin K_i)$. □

3 Multiclass birth and death processes

3.1 Stochastic comparison

When X and Y are random elements taking values in a partially ordered measurable space, we denote $X \leq_{st} Y$ and say that X is *stochastically less* than Y , if $E f(X) \leq E f(Y)$ for all positive increasing measurable functions. We use the terms increasing

and positive in the weak sense, so that a function f is increasing, if $f(x) \leq f(y)$ for all $x \leq y$, and positive if $f(x) \geq 0$ for all x . Moreover, we denote $X =_{\text{st}} Y$, if the distributions of X and Y are equal.

Let us endow the spaces \mathbb{Z}_+^N and $D(\mathbb{R}_+, \mathbb{Z}_+^N)$ with the usual coordinate-wise partial orders, so that $x \leq y$ in \mathbb{Z}_+^N if and only if $x_i \leq y_i$ for all i ; and $x \leq y$ in $D(\mathbb{R}_+, \mathbb{Z}_+^N)$ if and only if $x_i(t) \leq y_i(t)$ for all i and t . Recall that by Strassen’s theorem? (Kamae et al. 1977, Theorem 1), the stochastic processes X and Y having paths in $D(\mathbb{R}_+, \mathbb{Z}_+^N)$ satisfy $X \leq_{\text{st}} Y$ if and only if there exist processes \tilde{X} and \tilde{Y} defined on a common probability space such that $\tilde{X} =_{\text{st}} X$, $\tilde{Y} =_{\text{st}} Y$, and $\tilde{X}_i(t) \leq \tilde{Y}_i(t)$ for all i and t almost surely. The following result indicates a fundamental relation between stochastic ordering and stability.

Proposition 4 *Let X and Y be stochastic processes with paths in $D(\mathbb{R}_+, \mathbb{Z}_+^N)$ and assume that $X \leq_{\text{st}} Y$.*

1. *If X is transient, then so is Y .*
2. *If Y is stable, then so is X .*

Proof The first claim is a direct consequence of Strassen’s theorem, while the second follows directly from the definition of stability, because $P(|X(t)| > r) \leq P(|Y(t)| > r)$ for all r and t . □

A Markov process having paths in $D(\mathbb{R}_+, \mathbb{Z}_+^N)$ is called a *multiclass birth and death process*, if its transitions are given by

$$\begin{aligned} x &\mapsto x + e_i && \text{at rate } \lambda_i(x), \\ x &\mapsto x - e_i && \text{at rate } \phi_i(x)1(x_i > 0), \end{aligned}$$

where $\lambda_i(x)$ and $\phi_i(x)$ are some positive functions on \mathbb{Z}_+^N , called the class- i birth rates and death rates, respectively. The following lemma, which is proved using a direct coupling construction, gives a sufficient condition for the comparability of two multiclass birth and death processes.

Lemma 1 *Let $X = (X_1, \dots, X_I)$ and $Y = (Y_1, \dots, Y_J)$ be multiclass birth and death processes such that X has birth rates $\lambda_i(x)$ and death rates $\phi_i(x)$, and Y has birth rates $\eta_j(y)$ and death rates $\psi_j(y)$. Assume that for all $i = 1, \dots, I \wedge J$, and all $x \in \mathbb{Z}_+^I$ and $y \in \mathbb{Z}_+^J$ such that $x_i = y_i$ and $(x_1, \dots, x_{I \wedge J}) \leq (y_1, \dots, y_{I \wedge J})$,*

$$\lambda_i(x) \leq \eta_i(y), \tag{3.1}$$

$$\phi_i(x) \geq \psi_i(y). \tag{3.2}$$

Then for all $x \in \mathbb{Z}_+^I$ and $y \in \mathbb{Z}_+^J$ such that $(x_1, \dots, x_{I \wedge J}) \leq (y_1, \dots, y_{I \wedge J})$,

$$(X_1[x], \dots, X_{I \wedge J}[x]) \leq_{\text{st}} (Y_1[y], \dots, Y_{I \wedge J}[y]),$$

where $X[x]$ and $Y[y]$ are versions of X and Y started in x and y , respectively.

Proof Let (\tilde{X}, \tilde{Y}) be the Markov process with paths in $D(\mathbb{R}_+, U)$, where $U = \{(x, y) \in \mathbb{Z}_+^I \times \mathbb{Z}_+^J : (x_1, \dots, x_{I \wedge J}) \leq (y_1, \dots, y_{I \wedge J})\}$, having the upward transitions

$$\begin{array}{lll}
 (x, y) \mapsto (x + e_i, y) & \text{at rate } \lambda_i(x), & i \leq I \wedge J, x_i < y_i, \\
 (x, y) \mapsto (x, y + e_i) & \text{at rate } \eta_i(y), & i \leq I \wedge J, x_i < y_i, \\
 (x, y) \mapsto (x + e_i, y + e_i) & \text{at rate } \lambda_i(x), & i \leq I \wedge J, x_i = y_i, \\
 (x, y) \mapsto (x, y + e_i) & \text{at rate } \eta_i(y) - \lambda_i(x), & i \leq I \wedge J, x_i = y_i, \\
 (x, y) \mapsto (x + e_i, y) & \text{at rate } \lambda_i(x), & i > I \wedge J, \\
 (x, y) \mapsto (x, y + e_i) & \text{at rate } \eta_i(y), & i > I \wedge J,
 \end{array}$$

and the downward transitions

$$\begin{array}{lll}
 (x, y) \mapsto (x - e_i, y) & \text{at rate } \phi_i(x), & i \leq I \wedge J, 0 < x_i < y_i, \\
 (x, y) \mapsto (x, y - e_i) & \text{at rate } \psi_i(y), & i \leq I \wedge J, 0 < x_i < y_i, \\
 (x, y) \mapsto (x - e_i, y - e_i) & \text{at rate } \psi_i(y), & i \leq I \wedge J, 0 < x_i = y_i, \\
 (x, y) \mapsto (x - e_i, y) & \text{at rate } \phi_i(x) - \psi_i(y), & i \leq I \wedge J, 0 < x_i = y_i, \\
 (x, y) \mapsto (x - e_i, y) & \text{at rate } \phi_i(x), & i > I \wedge J, \\
 (x, y) \mapsto (x, y - e_i) & \text{at rate } \psi_i(y), & i > I \wedge J.
 \end{array}$$

In light of Eqs. 3.1 and 3.2, we see that all transition rates described above are positive. Moreover, because each of the transitions are mappings from U into U , we can be assured that the process (\tilde{X}, \tilde{Y}) exists.

By studying the marginals of the transition rates, we see that both \tilde{X} and \tilde{Y} are Markov, and that their intensity matrices coincide with those of X and Y , respectively. Hence, for all x and y such that $(x, y) \in U$, we have constructed versions of $X[x]$ and $Y[y]$ on a common probability space such that $(X_1[x](t), \dots, X_{I \wedge J}[x](t)) \leq (Y_1[y](t), \dots, Y_{I \wedge J}[y](t))$ for all t almost surely. □

3.2 Marginal drift conditions

In this section, we develop necessary and sufficient conditions for the stability of a multiclass birth and death process, given that each coordinate process except one is known to be stable. The following proposition extends the classical Neuts’ mean drift condition (Neuts 1978), see also Tweedie (1982). The proof follows closely the principles in Section 19 of Meyn and Tweedie (1993).

Proposition 5 *Let $X = (X_1, \dots, X_{n+1})$ be an $(n + 1)$ -class birth and death process with strictly positive birth rates λ_i and bounded death rates $\phi_i(x)$ such that $\phi_i(x) = \phi_i(x_1, \dots, x_n)$ only depends on the first n input arguments for all $i = 1, \dots, n$. Assume that:*

1. *The Markov process $X^n = (X_1, \dots, X_n)$ is stable and has the stationary distribution π .*

2. The birth rate of X_{n+1} satisfies the condition

$$\lambda_{n+1} < \sum_{x^n \in \mathbb{Z}_+^n} \left\{ \liminf_{x_{n+1} \rightarrow \infty} \phi_{n+1}(x^n, x_{n+1}) \right\} \pi(x^n).$$

Then the process $X = (X^n, X_{n+1})$ is stable.

In particular, if $\phi_{n+1}(x) = \phi_{n+1}(x_{n+1})$ only depends on x_{n+1} , then

$$\lambda_{n+1} < \liminf_{x_{n+1} \rightarrow \infty} \phi_{n+1}(x_{n+1})$$

is sufficient for the stability of X^{n+1} , regardless of the initial state.

To prove the above proposition, we use the following lemma:

Lemma 2 Let (X, Y) be a stochastic process with values in $S \times \mathbb{Z}_+$, where S is countable. Assume that

1. $X_t \rightarrow \pi$ weakly for some probability distribution π on S ,
2. $Y_t \rightarrow \infty$ in probability.

Then for all bounded f ,

$$\limsup_{t \rightarrow \infty} E f(X_t, Y_t) \leq \sum_x \left\{ \limsup_{y \rightarrow \infty} f(x, y) \right\} \pi(x). \tag{3.3}$$

Proof Assume first that S is finite, and let $\bar{f}(x) = \limsup_{y \rightarrow \infty} f(x, y)$. Then for any $\epsilon > 0$ there exists an r such that $f(x, y) \leq \bar{f}(x) + \epsilon$ for all x and all $y > r$. It follows that

$$\begin{aligned} E f(X_t, Y_t) &\leq E(\bar{f}(X_t) + \epsilon)1(Y_t > r) + E f(X_t, Y_t)1(Y_t \leq r) \\ &= E \bar{f}(X_t) + \epsilon + E[f(X_t, Y_t) - \bar{f}(X_t) - \epsilon]1(Y_t \leq r) \\ &\leq E \bar{f}(X_t) + \epsilon + 2\|f\|P(Y_t \leq r), \end{aligned}$$

where $\|f\| = \sup_{x,y} |f(x, y)|$. By letting $t \rightarrow \infty$ and recalling that ϵ is arbitrary, we see that the claim holds for a finite set S .

If S is countably infinite, then for any finite set K , the claim holds for the function $f_K(x, y) = f(x, y)1(x \in K)$. Hence, because $\bar{f}_K(x) = \bar{f}(x)1(x \in K)$, we get

$$\limsup_{t \rightarrow \infty} E f(X_t, Y_t) \leq \sum_x \bar{f}(x) \pi(x) + 2\|f\|\pi(K^c).$$

Thus the claim follows, because we can make $\pi(K^c)$ arbitrarily small by choosing K large enough. □

Proof of Proposition 5 Let us define $V(x) = x_{n+1}$ and denote the mean drift of V with respect to X by

$$\Delta V(x) = \lambda_{n+1} - \phi_{n+1}(x)1(x_{n+1} > 0).$$

Let us also define $V_M(x) = V(x)1(x_{n+1} \leq M)$ for $M > 0$. Then by Kolmogorov’s forward equation (Kallenberg 2002, Theorem 19.6) we have

$$\int_0^t E_x \Delta V_M(X(s)) ds = E_x V_M(X(t)) - V_M(x).$$

Because $V_M \rightarrow V$ and $\Delta V_M \rightarrow \Delta V$ pointwise as $M \rightarrow \infty$, and because $\|\Delta V_M\| \leq \lambda + \|\phi\|$, we see by applying dominated convergence on the left-hand side, and monotone convergence on the right, that

$$\int_0^t E_x \Delta V(X(s)) ds = E_x V(X(t)) - V(x).$$

Because V is positive, this implies that

$$\limsup_{t \rightarrow \infty} t^{-1} \int_0^t E_x \Delta V(X(s)) ds \geq 0,$$

and consequently,

$$\limsup_{t \rightarrow \infty} E_x \Delta V(X(t)) \geq 0.$$

Now assume that the process X is unstable. Then Proposition 2 implies that $X(t) \rightarrow \infty$ in probability. Because X^n is stable, it follows that $X_{n+1}(t) \rightarrow \infty$ in probability, and we may conclude by virtue of Lemma 2 that

$$\sum_{x^n \in \mathbb{Z}_+^n} \left\{ \lambda_{n+1} - \liminf_{x_{n+1} \rightarrow \infty} \phi_{n+1}(x^n, x_{n+1}) \right\} \pi(x^n) \geq \limsup_{t \rightarrow \infty} E_x \Delta V(X(t)) \geq 0.$$

□

To prove the following converse of Proposition 5, we make the additional assumption that ϕ_{n+1} only depends on the first n input arguments. Hence (X^n, X_{n+1}) is a Markov additive process, where the state space of the modulating process X^n is countably infinite. For Markov additive processes where the modulating process only takes on finitely many values, the following kind of result is well-known (Asmussen 2003, Proposition XI.2.11).

Proposition 6 *Let $X = (X_1, \dots, X_{n+1})$ be an $(n + 1)$ -class birth and death process with strictly positive birth rates λ_i and bounded death rates $\phi_i(x)$ such that $\phi_i(x) = \phi_i(x_1, \dots, x_n)$ only depends on the first n input arguments for all $i = 1, \dots, n + 1$. Assume that:*

1. *The Markov process $X^n = (X_1, \dots, X_n)$ is stable and has the stationary distribution π .*
2. *The birth rate of X_{n+1} satisfies the condition*

$$\lambda_{n+1} > \sum_{x^n \in \mathbb{Z}_+^n} \phi_{n+1}(x^n) \pi(x^n).$$

Then the process X_{n+1} is transient, regardless of the initial state.

Proof Let us define the free process (X^n, X_{n+1}^f) as the Markov process with values in $\mathbb{Z}_+^n \times \mathbb{Z}$, so that $X^n = (X_1, \dots, X_n)$ is as before, and conditional on X^n , X_{n+1}^f is a birth and death process on \mathbb{Z} with birth rates λ_{n+1} and time-varying death rates $\phi_{n+1}(X^n(t))$. Note first that conditional on X^n , the process X_{n+1}^f can be represented as a difference of two independent Poisson processes with intensity functions λ_{n+1} and $\phi_{n+1}(X(t))$, respectively. Hence, by conditioning on X^n , we find that

$$E_x \left[X_{n+1}^f(t) - X_{n+1}^f(0) \right] = E_x \int_0^t \left[\lambda_{n+1} - \phi_{n+1}(X^n(s)) \right] ds.$$

The above equation together with the strong Markov property shows that

$$M(t) = X_{n+1}^f(t) - X_{n+1}^f(0) - \int_0^t \left[\lambda_{n+1} - \phi_{n+1}(X^n(s)) \right] ds \tag{3.4}$$

is a martingale. Moreover, by conditioning on X again, one may verify that

$$E_x M(t)^2 = E_x \int_0^t \left[\lambda_{n+1} + \phi_{n+1}(X^n(s)) \right] ds \leq (\lambda_{n+1} + \|\phi_{n+1}\|)t.$$

This shows that $(1+t)^{-1}M(t)$ is a supermartingale that converges to zero in L^2 . Hence, $(1+t)^{-1}M(t)$ converges almost surely to zero as $t \rightarrow \infty$ (Rogers and Williams 1994, Theorem 69.1). On the other hand, by the ergodic theorem for positive recurrent Markov processes (Kallenberg 2002, Theorem 20.21),

$$t^{-1} \int_0^t \phi_{n+1}(X^n(s)) ds \rightarrow \sum_{x^n \in \mathbb{Z}_+^n} \phi_{n+1}(x^n) \pi(x^n) \quad \text{a.s.}$$

Hence, by dividing both sides of Eq. 3.4 by t and taking $t \rightarrow \infty$, we see that regardless of the initial state, $X_{n+1}^f(t) \rightarrow \infty$ almost surely.

Finally, observe that whenever (X^n, X_{n+1}) and (X^n, X_{n+1}^f) are started in the same initial state, X_{n+1} can be represented in terms of X_{n+1}^f via the Skorokhod map (Robert 2003, Theorem D.1)

$$X_{n+1}(t) = X_{n+1}^f(t) + \sup_{s \leq t} \left[-X_{n+1}^f(s) \right]^+.$$

Because $X_{n+1}^f(t) \rightarrow \infty$ almost surely, it follows that X_{n+1} is transient. □

4 Stability results for queueing systems

4.1 General service allocations

Let us return to the queueing system described in Section 2.1, so from now on $X = (X_1, \dots, X_N)$ describes the queue lengths of the system, and $\phi_i(x)$ is the service rate for queue i . The following result gives stability conditions valid for an arbitrary bounded service allocation $\phi = (\phi_1, \dots, \phi_N)$. Although these conditions are not sharp in general, they may provide useful inner and outer bounds for stability regions of complex systems that are not easy to analyze exactly.

Theorem 1 Let $X = (X_1, \dots, X_N)$ be the queue length process of the system with strictly positive arrival rates λ_i and bounded service rates $\phi_i(x)$. Then, regardless of the initial state, the process X_i is stable if

$$\lambda_i < \liminf_{x_i \rightarrow \infty} \inf_{x_j: j \neq i} \phi_i(x), \tag{4.1}$$

and transient, if

$$\lambda_i > \limsup_{x_i \rightarrow \infty} \sup_{x_j: j \neq i} \phi_i(x). \tag{4.2}$$

In particular, X is stable if Eq. 4.1 holds for all i .

Proof Assume that Eq. 4.1 holds for some i . By relabeling the queues if necessary, we may assume without loss of generality that $i = 1$. Let Y_1 be the one-class birth and death process with constant birth rate λ_1 and death rates $\psi_1(x_1) = \inf_{x_2, \dots, x_N} \phi_1(x)$. Because $\lambda_1 < \liminf_{x_1 \rightarrow \infty} \psi_1(x_1)$, it then follows from Proposition 5 that Y_1 is stable, regardless of the initial state. Moreover, because $\phi_1(x) \geq \psi_1(x_1)$ for all x , it follows from Lemma 1 that $X_1[x] \leq_{st} Y_1[x_1]$ for any initial state x . Hence, $X_1[x]$ is stable by Proposition 4.

Analogously, if Eq. 4.2 holds, then it again suffices to consider $i = 1$. In that case, we let Z_1 be the one-class birth and death process with birth rate λ_1 and death rates $\chi_1(x_1) = \sup_{x_2, \dots, x_N} \phi_1(x) + \epsilon$, where $\epsilon > 0$ is such that $\lambda_1 - \epsilon$ is strictly larger than the right-hand side of Eq. 4.2. Then Z_1 is irreducible and $\lambda_1 > \limsup_{x_1 \rightarrow \infty} \chi_1(x_1)$, so it follows from the classical theory of ordinary birth and death processes (Asmussen 2003, Proposition III.2.1) that Z_1 is transient. Applying Lemma 1 once more, we see that $Z_1[x_1] \leq_{st} X_1[x]$ for any $x \in \mathbb{Z}_+^N$. Hence, $X_1[x]$ is transient by Proposition 4.

Finally, if we assume that Eq. 4.1 holds for all i , then all $X_i[x]$ are stable, regardless of the initial state x , hence X is stable by Proposition 3. □

4.2 Partially decreasing service allocations

For the service allocation ϕ , the following notion of monotonicity is fundamental in comparing multiclass birth and death processes. A function $\phi = (\phi_1, \dots, \phi_N)$ from \mathbb{Z}_+^N into \mathbb{R}_+^N is said to be *partially decreasing* if for all i ,

$$\phi_i(x) \geq \phi_i(y) \quad \text{for all } x \geq y \text{ such that } x_i = y_i. \tag{4.3}$$

Note that a function $\phi = (\phi_1, \dots, \phi_N)$ is partially decreasing, if each of the coordinate functions ϕ_i is decreasing. Moreover, any function $\phi = (\phi_1, \dots, \phi_N)$ such that ϕ_i only depends on x_i , is partially decreasing. Recall also that a continuous-time Markov process X is said to be *monotone*, if the map $x \mapsto E_x f(X_t)$ is increasing for all t and for any bounded increasing function f . Using a result of Massey (1987, Theorem 5.2), it can be checked that a multiclass birth and death process X with constant birth rates and bounded state-dependent death rates ϕ_i is monotone if and only if ϕ is partially decreasing.

We define Y^n as the n -class birth and death process with birth rates λ_i and death rates $\ell^n \phi_i$ given by the lower partial limits

$$\ell^n \phi_i(x_1, \dots, x_n) = \lim_{r \rightarrow \infty} \inf_{x_{n+1}, \dots, x_N > r} \phi_i(x_1, \dots, x_N). \tag{4.4}$$

The process Y^n may intuitively be viewed as a partially saturated version of the queue length process $X = (X_1, \dots, X_N)$, where queues $1, \dots, n$ are allocated the asymptotically worst-case service rates as X_{n+1}, \dots, X_N tend to infinity. We also define

$$L_i^n(\lambda_1, \dots, \lambda_n; \phi) = \sum_{x \in \mathbb{Z}_+^n} \ell^n \phi_i(x) \pi^n(x), \tag{4.5}$$

if Y^n has a unique stationary distribution π^n , and set $L_i^n(\lambda_1, \dots, \lambda_n; \phi) = 0$ otherwise. Thus, the quantity $L_i^n(\lambda_1, \dots, \lambda_n; \phi)$ can be interpreted as a worst-case average service rate dedicated to queue i in a partially saturated system where the numbers of customers in queues $n + 1, \dots, N$ tend to infinity. For notational convenience, we define $L_i^n(\lambda_1, \dots, \lambda_n; \phi) = \ell^0 \phi_i$ for $n = 0$.

Theorem 2 *Let $X = (X_1, \dots, X_N)$ be the queue length process of the system with strictly positive arrival rates λ_i and a bounded partially decreasing service allocation $\phi = (\phi_1, \dots, \phi_N)$. Assume that there exists an n such that*

$$\lambda_i < L_i^{i-1}(\lambda_1, \dots, \lambda_{i-1}; \phi) \tag{4.6}$$

for all $i = 1, \dots, n$. Then the processes X_1, \dots, X_n are stable, regardless of the initial state.

Before proving Theorem 2, we establish some auxiliary results, the latter of which will also be used in Section 4.3.

Lemma 3 *Let $1 \leq n \leq N$. Then for all i and all x ,*

$$\ell^{n-1} \phi_i(x_1, \dots, x_{n-1}) \leq \liminf_{x_n \rightarrow \infty} \ell^n \phi_i(x_1, \dots, x_n).$$

Proof Fix an $x \in \mathbb{Z}_+^n$ and denote $\alpha = \ell^{n-1} \phi_i(x_1, \dots, x_{n-1})$. Then for any $\epsilon > 0$ there exists an r such that

$$\alpha - \epsilon \leq \inf_{y_n, \dots, y_N \geq r} \phi_i(x_1, \dots, x_{n-1}, y_n, \dots, y_N).$$

Hence, it follows that for all $y_n \geq r$ and for all $s \geq r$,

$$\begin{aligned} \alpha - \epsilon &\leq \inf_{y_{n+1}, \dots, y_N \geq r} \phi_i(x_1, \dots, x_{n-1}, y_n, \dots, y_N) \\ &\leq \inf_{y_{n+1}, \dots, y_N \geq s} \phi_i(x_1, \dots, x_{n-1}, y_n, \dots, y_N). \end{aligned}$$

By taking $s \rightarrow \infty$, it follows that $\alpha - \epsilon \leq \ell^n \phi_i(x_1, \dots, x_{n-1}, y_n)$ for all $y_n \geq r$, so the claim follows because ϵ is arbitrary. □

Lemma 4 *Assume that the function $\phi = (\phi_1, \dots, \phi_N)$ is bounded and partially decreasing, and assume that λ_i are strictly positive and satisfy inequalities (4.6) for $i = 1, \dots, n$. Then the n -class birth and death process $Y^n = (Y_1^n, \dots, Y_n^n)$ with birth rates λ_i and death rates $\ell^n \phi_i$ is stable.*

Proof To prove the claim for $n = 1$, let us assume that $\lambda_1 < \ell^0\phi_1$. Then Y^1 is a one-class birth and death process with birth rate λ_1 and state-dependent death rates $\ell^1\phi_1(x_1)$. The stability of Y_1 follows from Proposition 5, and Lemma 3,

$$\lambda_1 < \ell^0\phi_1 \leq \liminf_{x_1 \rightarrow \infty} \ell^1\phi_1(x_1).$$

To proceed by induction, suppose that the claim is true for $n - 1$, and assume that the inequalities (4.6) hold for $i = 1, \dots, n$. Then Y^{n-1} is stable by the induction assumption, and Proposition 2 shows that Y^{n-1} has a unique stationary distribution π^{n-1} . Let $Z^n = (Z^n_1, \dots, Z^n_n)$ be the n -class birth and death process with birth rates λ_i and death rates

$$\psi_i(x_1, \dots, x_n) = \begin{cases} \ell^{n-1}\phi_i(x_1, \dots, x_{n-1}), & i < n, \\ \ell^n\phi_i(x_1, \dots, x_n), & i = n. \end{cases}$$

This choice of rates implies that the marginal process $(Z^n_1, \dots, Z^n_{n-1})$ is Markov and coincides with Y^{n-1} . Now observe that the condition $\lambda_n < L^{n-1}_n(\lambda_1, \dots, \lambda_{n-1}; \phi)$ is equivalent to

$$\lambda_n < \sum_{x \in \mathbb{Z}^{n-1}_+} \ell^{n-1}\phi_n(x_1, \dots, x_{n-1}) \pi^{n-1}(x_1, \dots, x_{n-1}),$$

so using Lemma 3, we see that

$$\lambda_n < \sum_{x \in \mathbb{Z}^{n-1}_+} \left\{ \liminf_{x_n \rightarrow \infty} \psi_n(x_1, \dots, x_{n-1}, x_n) \right\} \pi^{n-1}(x_1, \dots, x_{n-1}).$$

Hence, Proposition 5 shows that Z^n is stable.

Because ϕ is partially decreasing, we have $\ell^{n-1}\phi_i(x_1, \dots, x_{n-1}) \leq \phi_i(x)$ for all $i < n$, and thus $\ell^{n-1}\phi_i(x_1, \dots, x_{n-1}) \leq \ell^n\phi_i(x_1, \dots, x_n)$ for all $i < n$. In particular, it follows that $\psi_i(x) \leq \ell^n\phi_i(x)$ for all x and all $i \leq n$. Further, because $\ell^n\phi$ is also partially decreasing, it follows that $\ell^n\phi_i(x) \geq \psi_i(y)$ for all x and y in \mathbb{Z}^n_+ such that $x \leq y$ and $x_i = y_i$. Thus, Lemma 1 implies that $Y^n \leq_{st} Z^n$, whenever Y^n and Z^n are started in the same initial state. Hence, Y^n is stable by Proposition 4. □

Proof of Theorem 2 Assuming the inequalities (4.6) are valid for $i = 1, \dots, n$, we see using Lemma 4 that Y^n is stable. Moreover, $\phi_i(x_1, \dots, x_N) \geq \ell^n\phi_i(x_1, \dots, x_n)$ for all $i \leq n$ and all $x \in \mathbb{Z}^N_+$, because ϕ is partially decreasing. Because $\ell^n\phi$ is also partially decreasing, it follows that $\phi_i(x) \geq \ell^n\phi_i(y)$ for all $x \in \mathbb{Z}^N_+$ and $y \in \mathbb{Z}^n_+$ such that $(x_1, \dots, x_n) \leq (y_1, \dots, y_n)$ and $x_i = y_i$. Hence, Lemma 1 shows that $(X_1[x], \dots, X_n[x]) \leq_{st} (Y^n_1[y], \dots, Y^n_n[y])$ for all initial states $x \in \mathbb{Z}^N_+$ and $y \in \mathbb{Z}^n_+$ such that $(x_1, \dots, x_n) \leq (y_1, \dots, y_n)$. Proposition 4 thus implies that $(X_1[x], \dots, X_n[x])$ is stable. □

4.3 Partially decreasing service allocations with uniform limits

In the following, we restrict ourselves to queueing systems where the service allocation $\phi = (\phi_1, \dots, \phi_N)$ is such that each coordinate function ϕ_i has a uniform limit as

some of the input variables tend to infinity. More precisely, we say that a function $f : \mathbb{Z}_+^N \rightarrow \mathbb{R}$ has *uniform limits at infinity*, if the following conditions hold:

1. There exists a constant f^0 such that

$$\sup_{x \in \mathbb{Z}_+^N : x_1, \dots, x_N > r} |f(x) - f^0| \rightarrow 0, \quad \text{as } r \rightarrow \infty.$$

2. For any $n = 1, \dots, N - 1$ and any permutation σ on $\{1, \dots, N\}$, there exists a function $f^{n,\sigma} : \mathbb{Z}_+^n \rightarrow \mathbb{R}$ such that as $r \rightarrow \infty$,

$$\sup_{x \in \mathbb{Z}_+^N : x_{\sigma(n+1)}, \dots, x_{\sigma(N)} > r} |f(x) - f^{n,\sigma}(x_{\sigma(1)}, \dots, x_{\sigma(n)})| \rightarrow 0.$$

We will next show that the class of functions having uniform limits at infinity is rich enough to be of interest. For example, assume that the allocation function ϕ is of the form

$$\phi_i(x) = g_i(x_i)h(x)$$

where g_i has a limit at infinity and h is decreasing. Then g_i obviously has uniform limits at infinity, and hence the same applies for ϕ_i using the result below.

Proposition 7 *Let f and g be bounded functions on \mathbb{Z}_+^N .*

1. *If f is positive and decreasing, then it has uniform limits at infinity.*
2. *If f and g have uniform limits at infinity, then so do the functions $f + g$ and fg .*

Proof See [Appendix](#). □

If ϕ has uniform limits, then the partial lower limits $\ell^n \phi$ defined in Eq. 4.4 become true limits in the sense that

$$\phi(x_1, \dots, x_N) \rightarrow \ell^n \phi(x_1, \dots, x_n)$$

uniformly over x_1, \dots, x_n , as $\min(x_{n+1}, \dots, x_N) \rightarrow \infty$.

Theorem 3 *Let $X = (X_1, \dots, X_N)$ be the queue length process of the system with strictly positive arrival rates λ_i and a partially decreasing service allocation $\phi = (\phi_1, \dots, \phi_N)$ having uniform limits at infinity. Assume that there is an index n such that*

$$\lambda_i < L_i^{i-1}(\lambda_1, \dots, \lambda_{i-1}; \phi) \quad \text{for all } i \leq n, \tag{4.7}$$

$$\lambda_i > L_i^n(\lambda_1, \dots, \lambda_n; \phi) \quad \text{for all } i > n. \tag{4.8}$$

Then the process (X_{n+1}, \dots, X_N) is unstable, regardless of the initial state.

Proof Given $\epsilon \geq 0$, let $Y^{n,\epsilon}$ be the n -class birth and death process with birth rates λ_i and death rates $\ell^n \phi_i(x_1, \dots, x_n) + \epsilon$, and let $Y^n = Y^{n,0}$. Then Y^n is stable by Lemma 4. Further, because $\ell^n \phi$ is partially decreasing, it follows by Lemma 1 that $Y^{n,\epsilon} \leq_{\text{st}} Y^n$, whenever $Y^{n,\epsilon}$ and Y^n are started in the same state. Hence, Lemma 4 implies that $Y^{n,\epsilon}$ is stable for all $\epsilon > 0$, and moreover $\pi^{n,\epsilon} \leq_{\text{st}} \pi^n$, where $\pi^{n,\epsilon}$ and π^n denote the stationary distributions of $Y^{n,\epsilon}$ and Y^n , respectively. In particular, the

family $\{\pi^{n,\epsilon}\}_{\epsilon>0}$ is tight, and so by Lemma 6 (see [Appendix](#)), $\pi^{n,\epsilon} \rightarrow \pi^n$ weakly, as $\epsilon \rightarrow 0$. Because ϕ is bounded, it follows that for all i ,

$$\lim_{\epsilon \rightarrow 0} \sum_{x \in \mathbb{Z}_+^n} \ell^n \phi_i(x) \pi^{n,\epsilon}(x) = L_i^n(\lambda_1, \dots, \lambda_n; \phi).$$

Hence, by Eq. 4.8, we can choose an $\epsilon > 0$ such that

$$\lambda_i > \sum_{x \in \mathbb{Z}_+^n} \ell^n \phi_i(x) \pi^{n,\epsilon}(x) \quad \text{for all } i > n. \tag{4.9}$$

Let Z be the N -class birth and death process with birth rates λ_i and death rates $\psi_i(x_1, \dots, x_N) = \ell^n \phi_i(x_1, \dots, x_n) + \epsilon$. Then (Z_1, \dots, Z_n) is Markov and coincides with $Y^{n,\epsilon}$ as defined above. Moreover, inequality (4.9) implies that for all $i > n$, the mean drift of Z_i is strictly positive, so Proposition 6 implies that Z_i is transient for each $i > n$.

Next, observe that because ϕ_i has uniform limits, it follows that there exists an r_0 such that $\psi_i(x) \geq \phi_i(x)$ for all $i = 1, \dots, N$ and for all x such that $x_{n+1}, \dots, x_N \geq r_0$. Let us choose an $r \geq r_0$ and define A_r to be the complement of the set

$$B_r = \{x \in \mathbb{Z}_+^N : x_{n+1}, \dots, x_N \geq r\},$$

Because each of the processes Z_{n+1}, \dots, Z_N is transient, it follows that we can choose an $x \in B_r$ such that

$$P_x(\tau_{A_r}(Z) = \infty) > 0, \tag{4.10}$$

where $\tau_{A_r}(Z) = \inf\{t > 0 : Z(t) \in A_r\}$ denotes the hitting time of Z into A_r . Because ϕ is partially decreasing, it follows that $\psi_i(x) \geq \phi_i(y)$ for all x and y in B_r such that $x \leq y$ and $x_i = y_i$. Using a similar coupling construction as in Lemma 1, it is then straightforward to verify that $\tau_{A_r}(Z[x]) \leq_{st} \tau_{A_r}(X[x])$ for all $x \in B_r$. Hence using Eq. 4.10 we may conclude that there exists an $x \in B_r$ such that

$$P_x(\tau_{A_r}(X) = \infty) > 0. \tag{4.11}$$

We now assume that X is stable and derive a contradiction. Using Proposition 2, we know that there exists a set C such that $X[x]$ is irreducible and positive recurrent for all $x \in C$. Let us now choose an $r \geq r_0$ such that $C \cap A_r$ is nonempty. Using Lemma 7 and standard properties of irreducible Markov processes (Kallenberg 2002, Proposition 8.13), we then know that the hitting time of $X[x]$ into $C \cap A_r$ is finite almost surely for all x , which contradicts Eq. 4.11. Hence, X must be unstable. In particular, because (X_1, \dots, X_n) is stable, it follows from Proposition 3 that (X_{n+1}, \dots, X_N) is unstable. □

We are now ready to present our main theorem. For any permutation σ on $\{1, \dots, N\}$, we define $\lambda_i^\sigma = \lambda_{\sigma(i)}$ and $\phi_i^\sigma(x) = \phi_{\sigma(i)}(x_{\sigma^{-1}(1)}, \dots, x_{\sigma^{-1}(N)})$. The vector λ^σ and the function ϕ^σ will then correspond to the system where queues are relabeled according to σ . Denote

$$S_N(\phi) = \left\{ \lambda \in (0, \infty)^N : \lambda_i < L_i^{i-1}(\lambda_1, \dots, \lambda_{i-1}; \phi) \ \forall i = 1, \dots, N \right\},$$

where $L_i^{i-1}(\lambda_1, \dots, \lambda_{i-1}; \phi)$ are given by Eq. 4.5, and define

$$\mathcal{S}(\phi) = \bigcup_{\sigma} \{ \lambda \in (0, \infty)^N : \lambda^\sigma \in S_N(\phi^\sigma) \}, \tag{4.12}$$

where the union is taken over all permutations on $\{1, \dots, N\}$.

Theorem 4 *Assume $\phi = (\phi_1, \dots, \phi_N)$ is partially decreasing and has uniform limits at infinity. Then the set $\mathcal{S}(\phi)$ defined by Eq. 4.12 is open, and the queue length process $X = (X_1, \dots, X_N)$ of the system with arrival rates $\lambda = (\lambda_1, \dots, \lambda_N)$ and service allocation ϕ is stable for all $\lambda \in \mathcal{S}(\phi)$, and unstable for all λ outside the closure of $\mathcal{S}(\phi)$.*

Note that it is not possible to characterize the stability of the system for arrival rate vectors belonging to the boundary of $\mathcal{S}(\phi)$ without more detailed information on the allocation function ϕ . Consider for example the one-server system where $\phi_1(x_1) = (1 + 1/x_1)^\alpha$ for some $\alpha \geq 0$. Then $\mathcal{S}(\phi) = \{ \lambda_1 : \lambda_1 < 1 \}$, and the system with $\lambda_1 = 1$ is positive recurrent if $\alpha > 1$ and null recurrent otherwise (Asmussen 2003, Section III.2).

Lemma 5 *Let $\phi = (\phi_1, \dots, \phi_N)$ be bounded and partially decreasing. Then for all $n = 1, \dots, N$, the set*

$$S_n(\phi) = \{ \lambda \in (0, \infty)^N : \lambda_i < L_i^{i-1}(\lambda_1, \dots, \lambda_{i-1}; \phi) \ \forall i = 1, \dots, n \}$$

is open, and the function $f^n = (f_1^n, \dots, f_n^n)$ on $(0, \infty)^N$ defined by

$$f_i^n(\lambda_1, \dots, \lambda_N) = L_i^{i-1}(\lambda_1, \dots, \lambda_{i-1}; \phi) \tag{4.13}$$

is continuous on $S_{n-1}(\phi)$.

Proof First observe that the function f^1 , being a constant, is continuous on the open set $S_0(\phi) = (0, \infty)^N$. To proceed by induction, let us next assume that $S_{n-1}(\phi)$ is open and f^n is continuous in $S_{n-1}(\phi)$ for some n . To show that $S_n(\phi)$ is open, assume that $S_n(\phi)$ is nonempty, and choose a vector $\lambda \in S_n(\phi)$. Then

$$\lambda_n < L_n^{n-1}(\lambda_1, \dots, \lambda_{n-1}; \phi). \tag{4.14}$$

Moreover, because $S_n(\phi) \subseteq S_{n-1}(\phi)$, it follows that f^n is continuous at λ , so in particular the map $\eta \mapsto L_n^{n-1}(\eta_1, \dots, \eta_{n-1}; \phi)$ is continuous at λ . This implies that Eq. 4.14 remains valid in some open neighborhood B_λ of λ . Further, because $S_{n-1}(\phi)$ is open, there is another open neighborhood B'_λ of λ such that $B'_\lambda \subseteq S_{n-1}(\phi)$. It follows that $B_\lambda \cap B'_\lambda \subseteq S_n(\phi)$, so we may conclude that $S_n(\phi)$ is open.

To complete the induction, we next prove the continuity of f^{n+1} on $S_n(\phi)$, under the assumptions that $S_n(\phi)$ is open and f^n is continuous on $S_{n-1}(\phi)$. If $\lambda \in S_n(\phi)$, then f^n is continuous at λ , because $S_n(\phi) \subseteq S_{n-1}(\phi)$. Because the first n coordinate functions of f^{n+1} coincide with f^n , it is sufficient to prove that the function $\eta \mapsto L_{n+1}^n(\eta_1, \dots, \eta_n; \phi)$ is continuous at λ . Thus, let λ^k be a sequence converging to λ . Because $S_n(\phi)$ is open, we may assume without loss of generality that there exists a vector $\lambda' \in S_n(\phi)$ such that $\lambda^k \leq \lambda'$ for all k . Let Y be the n -class birth and death process with birth rates λ_i and death rates $\ell^n \phi_i(x)$, $i = 1, \dots, n$, and let Y^k and Y' be the corresponding processes with λ replaced by λ^k and λ' , respectively. Then

by Lemma 4, all of the processes Y , Y^k , and Y' are stable, so we denote their stationary distributions by π , π^k , and π' , respectively. Moreover, because $\ell^n \phi$ is partially decreasing, Lemma 1 shows that $\pi^k \leq_{st} \pi'$ for all k , so the family $\{\pi^k\}_{k \geq 0}$ is tight. Hence, we can apply Lemma 6 to see that $\pi^k \rightarrow \pi$ weakly, so it follows from the boundedness of ϕ that $L_{n+1}^n(\lambda_1^k, \dots, \lambda_n^k, \phi) \rightarrow L_{n+1}^n(\lambda_1, \dots, \lambda_n; \phi)$. This shows that f^{n+1} is continuous at λ . \square

Proof of Theorem 4 First, let us note that by using Lemma 5 and relabeling the classes if necessary, we see that the set $\{\lambda : \lambda^\sigma \in S_N(\phi^\sigma)\}$ is open for all σ . Hence, the set $\mathcal{S}(\phi)$ is open. Moreover, by again relabeling the classes if necessary, we find using Theorem 2 that $\lambda \in \mathcal{S}(\phi)$ implies the stability of the queueing system.

To study the instability of the system, let us define analogously to Eq. 4.12 the set

$$\mathcal{U}(\phi) = \bigcup_{\sigma} \{\lambda \in (0, \infty)^N : \lambda^\sigma \in U_N(\phi^\sigma)\},$$

where the union is taken over all permutations on $\{1, \dots, N\}$, and $U_N(\phi)$ is defined as the set of $\lambda \in (0, \infty)^N$ such that the inequalities (4.7) and (4.8) are valid for some $n \in \{0, \dots, N - 1\}$. Then by first relabeling the classes if necessary, and then using Theorem 3, we see that the system is instable for all $\lambda \in \mathcal{U}(\phi)$.

It remains to show the instability of the system under the assumption that λ belongs to the complement of the closure of $\mathcal{S}(\phi)$ in $(0, \infty)^N$, which we denote by $\text{ext}(\mathcal{S})$. We prove this by showing that for each $\lambda \in \text{ext}(\mathcal{S})$ there exists a $\tilde{\lambda} \in \mathcal{U}(\phi)$ such that $\tilde{\lambda} \leq \lambda$ and then applying Proposition 4.

Let $\lambda \in \text{ext}(\mathcal{S})$. Given a permutation σ , let us define $n(\lambda, \sigma)$ as the largest integer n such that

$$\lambda_i^\sigma < L_i^{i-1}(\lambda_1^\sigma, \dots, \lambda_{i-1}^\sigma; \phi^\sigma) \quad \text{for all } i < n. \tag{4.15}$$

Because $\lambda \notin \mathcal{S}(\phi)$, we know that $n(\lambda, \sigma) < N$ for each σ , and

$$\lambda_i^\sigma \geq L_i^{i-1}(\lambda_1^\sigma, \dots, \lambda_{i-1}^\sigma; \phi^\sigma) \quad \text{for } i = n(\lambda, \sigma). \tag{4.16}$$

Let D be the set of σ for which Eq. 4.16 is strict, and assume $\sigma \in D$. Then by Lemma 5 we see that the function

$$(\lambda_1^\sigma, \dots, \lambda_{n-1}^\sigma) \mapsto (L_1^0(\phi^\sigma), \dots, L_n^{n-1}(\lambda_1^\sigma, \dots, \lambda_{n-1}^\sigma; \phi^\sigma)) \tag{4.17}$$

is continuous in a neighborhood of $(\lambda_1^\sigma, \dots, \lambda_{n-1}^\sigma)$. Thus we can choose an $\epsilon_\sigma \in (0, \min_j \lambda_j)$ such that the inequalities (4.15) remain valid and the inequality (4.16) remains strict when λ is replaced by $\tilde{\lambda}$ such that $\|\tilde{\lambda} - \lambda\| < \epsilon_\sigma$.

On the other hand, if $\sigma \notin D$, then again by the continuity of the function in Eq. 4.17, there exists an $\epsilon_\sigma \in (0, \min_j \lambda_j)$ such that the inequalities (4.15) remain valid when λ is replaced by $\tilde{\lambda}$ such that $\|\tilde{\lambda} - \lambda\| < \epsilon_\sigma$. Moreover, if we let $\tilde{\lambda} = \lambda - re$ for some $r \in (0, \epsilon_\sigma)$, then it can be checked using Lemma 1 that $L_n^{n-1}(\lambda_1^\sigma, \dots, \lambda_{n-1}^\sigma; \phi^\sigma) \leq L_n^{n-1}(\tilde{\lambda}_1^\sigma, \dots, \tilde{\lambda}_{n-1}^\sigma; \phi^\sigma)$, because the function $\ell^{n-1} \phi^\sigma$ is partially decreasing, and in particular, $\ell^{n-1} \phi_n^\sigma$ is decreasing. Thus we see that for $n = n(\lambda, \sigma)$,

$$\begin{aligned} \tilde{\lambda}_n^\sigma &< \lambda_n^\sigma = L_n^{n-1}(\lambda_1^\sigma, \dots, \lambda_{n-1}^\sigma; \phi^\sigma) \\ &\leq L_n^{n-1}(\tilde{\lambda}_1^\sigma, \dots, \tilde{\lambda}_{n-1}^\sigma; \phi^\sigma). \end{aligned}$$

Let $\epsilon = \min_{\sigma} \epsilon_{\sigma}$ and choose a small enough $r \in (0, \epsilon)$ such that $\tilde{\lambda} = \lambda - re$ belongs to $\text{ext}(\mathcal{S})$. Then it follows from the above observations that for all σ ,

$$\begin{aligned} \tilde{\lambda}_i^{\sigma} &< L_i^{i-1} \left(\tilde{\lambda}_1^{\sigma}, \dots, \tilde{\lambda}_{i-1}^{\sigma}; \phi^{\sigma} \right), & i < n(\lambda, \sigma), \\ \tilde{\lambda}_i^{\sigma} &\neq L_i^{i-1} \left(\tilde{\lambda}_1^{\sigma}, \dots, \tilde{\lambda}_{i-1}^{\sigma}; \phi^{\sigma} \right), & i = n(\lambda, \sigma). \end{aligned}$$

Hence, either $\tilde{\lambda}_n^{\sigma} > L_n^{n-1} \left(\tilde{\lambda}_1^{\sigma}, \dots, \tilde{\lambda}_{n-1}^{\sigma}; \phi^{\sigma} \right)$ for all σ and all $n = n(\lambda, \sigma)$, which implies that $\tilde{\lambda} \in \mathcal{U}(\phi)$; or else,

$$n(\tilde{\lambda}, \sigma) \geq n(\lambda, \sigma) \quad \text{for all } \sigma, \tag{4.18}$$

and $\tilde{\lambda}_n^{\sigma} < L_n^{n-1} \left(\tilde{\lambda}_1^{\sigma}, \dots, \tilde{\lambda}_{n-1}^{\sigma}; \phi^{\sigma} \right)$ for at least some σ and $n = n(\lambda, \sigma)$, which implies that Eq. 4.18 holds strictly for at least one σ .

Assuming $\lambda^k \in \text{ext}(\mathcal{S}) \setminus \mathcal{U}(\phi)$, we can apply the above procedure to λ^k to find a $\lambda^{k+1} \in \text{ext}(\mathcal{S})$ such that $\lambda^{k+1} \leq \lambda^k$, and so that either $\lambda^{k+1} \in \mathcal{U}(\phi)$, or else

$$n(\lambda^{k+1}, \sigma) \geq n(\lambda^k, \sigma) \quad \text{for all } \sigma, \tag{4.19}$$

where Eq. 4.19 holds strictly for at least one σ . The sequence λ^k started in $\lambda^1 = \lambda$ must hit $\mathcal{U}(\phi)$ for some k , because otherwise $n(\lambda^k, \sigma) = N + 1$ eventually for some k and some σ , which would imply that $\lambda^k \in \mathcal{S}(\phi)$. Hence, $\lambda^k \in \mathcal{U}$ for some k , and $\lambda^k \leq \lambda$. Now by Proposition 4, it follows that the system is unstable. \square

5 Applications

5.1 Three weakly coupled queues

Consider a system of three queues where the service rates at each queue only depend on whether the other queues are empty or not, so that for all $i \neq j \neq k$,

$$\phi_i(x) = \begin{cases} a_i, & x_j = 0, x_k = 0, \\ a_{ij}, & x_j > 0, x_k = 0, \\ 1, & x_j > 0, x_k > 0. \end{cases}$$

Let us assume $a_i \geq a_{ij} \geq 1$, so that $\phi = (\phi_1, \phi_2, \phi_3)$ is partially decreasing.

Theorem 4 shows that the stability region is a union of six regions corresponding to the six possible permutations of the queues. The first of these regions corresponding to the identity permutation is the set of $(\lambda_1, \lambda_2, \lambda_3)$ such that

$$\lambda_1 < 1, \tag{5.1}$$

$$\lambda_2 < \lambda_1 + a_{23}(1 - \lambda_1), \tag{5.2}$$

$$\lambda_3 < a_{31}\pi_{00} + a_{31}\pi_{10} + a_{32}\pi_{01} + \pi_{11}, \tag{5.3}$$

where

$$\begin{aligned} \pi_{00} &= P(Y_1 = 0, Y_2 = 0), \\ \pi_{01} &= P(Y_1 = 0, Y_2 > 0), \\ \pi_{10} &= P(Y_1 > 0, Y_2 = 0), \\ \pi_{11} &= P(Y_1 > 0, Y_2 > 0), \end{aligned}$$

and $Y = (Y_1, Y_2)$ is a random vector distributed according to the stationary number of customers in queues 1 and 2 given that the length of queue 3 is infinite, which is well-defined when inequalities (5.1) and (5.2) hold. To the best of our knowledge, there are no closed-form expressions available for the probabilities $\pi_{00}, \pi_{01}, \pi_{10}, \pi_{11}$, so they must be evaluated numerically.

The above result coincides with the stability characterization derived earlier by Fayolle, Malyshev, and Menshikov using Foster–Lyapunov criteria for deflected random walks in \mathbb{Z}_+^3 (Fayolle et al. 1995, Theorem 4.4.4).

5.2 Two interfering wireless base stations with channel-aware scheduling

We will now turn our attention to a system of two queues served at the state-dependent rates

$$\phi_i(x) = g_i(x_i)h_i(x),$$

where g_i is a bounded increasing function on \mathbb{Z}_+ , and h_i is a decreasing function on \mathbb{Z}_+^2 . This particular form of allocation function arises as a model for two interfering wireless base stations operating according to a channel-aware scheduling discipline, where the functions h_i model the interference between the base stations, and the functions g_i represent the scheduling gain, which increases in the number of customers due to multiuser diversity (Liu et al. 2003). Denote

$$\begin{aligned} g_i^* &= \lim_{x_i \rightarrow \infty} g_i(x_i), \\ h_i^* &= \lim_{r \rightarrow \infty} \sup_{x_1, x_2 \geq r} h_i(x), \\ h_i^j(x_i) &= \limsup_{x_j \rightarrow \infty} h_i(x), \quad j \neq i, \end{aligned}$$

and let

$$\bar{h}_j^i(\lambda_i, \phi) = \sum_{x_i} h_j^i(x_i) \pi^i(x_i), \quad j \neq i,$$

where π^i is the probability distribution (c is a normalization constant) given by

$$\pi^i(x_i) = c \prod_{z=1}^{x_i} \frac{\lambda_i}{g_i(z)h_i^j(z)}, \quad j \neq i.$$




Fig. 1 Stability regions for interference functions of the form (5.5) with $\gamma = 0.05$ (left) and $\gamma = 2.0$ (right)

The stability region can now be described using Theorem 4 as the set of (λ_1, λ_2) such that either

$$\lambda_1 < g_1^* h_1^* \text{ and } \lambda_2 < g_2^* \bar{h}_2^1(\lambda_1, \phi),$$

$$\text{or } \lambda_2 < g_2^* h_2^* \text{ and } \lambda_1 < g_1^* \bar{h}_1^2(\lambda_2, \phi).$$

Figure 1 represents the full and partial stability regions of the system, where the scheduling gains are given by

$$g_i(x_i) = \min(3, \log(1 + x_i)), \tag{5.4}$$

and the interference functions are of the form

$$h_i(x_j) = \frac{1}{6 - 4e^{-\gamma x_j}}, \quad j \neq i. \tag{5.5}$$

The area \mathcal{S} is the set of arrival rates such that both queues are stable, while the areas \mathcal{S}_i correspond to the set of arrival rates so that only queue i is stable, and \mathcal{U} is the set




Fig. 2 Stability regions for interference functions of the form (5.6) with $\gamma = 0.4$ (left) and $\gamma = 2.0$ (right)

of arrival rates where both queues are unstable. Figure 2 illustrates the corresponding stability regions when g_i are as in Eq. 5.4 and

$$h_i(x_j) = \frac{1}{6 - 4(1 + x_j)^{-\gamma}}, \quad j \neq i. \tag{5.6}$$

6 Conclusion

We provided sufficient and necessary conditions for the stability of a parallel queueing system with coupled service rates, and showed that these conditions are sharp when the service rate at each queue is decreasing in the number of customers in other queues, and has uniform limits as the queue lengths tend to infinity. Moreover, we presented conditions for partial stability, where only some of the queues are stable. The most general stability conditions, although not sharp, may yield useful inner and outer bounds for the stability region of systems that are too complex to characterize exactly. An interesting and important direction for future research is to study whether the given results extend to the case where the service allocation does not have uniform limits, and the service times distributions are nonexponential.

Appendix 1

A1.1 Small perturbations of transition rates

Lemma 6 *Let X and X^n be continuous-time Markov processes on a countable state space having transition rates $q(x, y)$ and $q^n(x, y)$, and unique stationary distributions π and π^n , respectively. Assume that*

1. $q^n(x, y) \rightarrow q(x, y)$ as $n \rightarrow \infty$ for all x and y ,
2. the set $\{x : q^n(x, y) \neq 0 \text{ for some } n\}$ is finite for all y ,
3. $\{\pi^n\}_{n \geq 0}$ is a tight family of probability measures.

Then $\pi^n(x) \rightarrow \pi(x)$ for all x .

Proof Let us assume that $\pi^n(z)$ does not converge to $\pi(z)$ for some z . Then there exists an $\epsilon > 0$ and a subsequence $\mathbb{Z}'_+ \subseteq \mathbb{Z}_+$ such that $|\pi^n(z) - \pi(z)| \geq \epsilon$ for all $n \in \mathbb{Z}'_+$. Because $\{\pi^n\}_{n \in \mathbb{Z}'_+}$ is tight, there exists a further subsequence $\mathbb{Z}''_+ \subseteq \mathbb{Z}'_+$ such that π^n converges weakly to a probability measure $\tilde{\pi}$ as $n \rightarrow \infty$ along \mathbb{Z}''_+ (Kallenberg 2002, Proposition 5.21). Observe that for all y and for all n ,

$$\sum_x \pi^n(x)q^n(x, y) = 0.$$

By virtue of assumption 2, we can take $n \rightarrow \infty$ along \mathbb{Z}''_+ on both sides of the above equation, and bring the limit inside the sum, which shows that

$$\sum_x \tilde{\pi}(x)q(x, y) = 0.$$

Because we assumed the stationary distribution of X to be unique, it follows that $\tilde{\pi} = \pi$, and hence $\pi^n \rightarrow \pi$ weakly along \mathbb{Z}'_+ . This is a contradiction, because $|\pi^n(z) - \pi(z)| \geq \epsilon$ for all $n \in \mathbb{Z}'_+$. \square

A1.2 Stable multiclass birth and death processes with strictly positive birth rates

Lemma 7 *Let X be a N -class birth and death process with birth rates λ_i and bounded death rates $\phi_i(x)$, and assume that $\lambda_i > 0$ for all i . Then the hitting time of X into an arbitrary increasing set A is almost surely finite, regardless of the initial state.*

Proof Let A be an increasing set. Then $re \in A$ for some positive integer r , where $e = (1, \dots, 1) \in \mathbb{Z}_+^N$. Let \hat{X} be the discrete-time jump chain of X , with $\hat{X}(n)$ being the value of X at its n -th jump. Then for all x ,

$$P_x(\hat{X}(1) = x + e_i) = \frac{\lambda_i}{\sum_j \lambda_j + \sum_{j:x_j>0} \phi_j(x)},$$

and because \hat{X} can reach A from any state x by taking r upward jumps into all coordinate directions, we see that $P_x(\hat{X}(r^N) \in A) \geq \delta$, where

$$\delta = \left(\frac{\min_j \lambda_j}{\sum_j \lambda_j + N\|\phi\|} \right)^{rN} > 0.$$

By induction, it then follows that for all x and all M ,

$$P_x(\hat{X}(mr^N) \notin A \forall m = 1, \dots, M) \leq (1 - \delta)^M.$$

Thus, by taking $M \rightarrow \infty$, we may conclude that $P_x(\hat{\tau}_A < \infty) = 1$, which is equivalent to $P_x(\tau_A < \infty) = 1$. \square

Proof of Proposition 2 We prove that (1) \Rightarrow (2) \Rightarrow (3), the reverse direction being clear. Let us denote $x \rightarrow y$ if the process X started in x can reach y , and let $C(x) = \{y : x \rightarrow y \text{ and } y \rightarrow x\}$ be the communicating class associated with x . Recall that a set C is said to be absorbing if $x \rightarrow y$ implies $y \in C$ for all $x \in C$. Observe first that because all birth rates of X are strictly positive, it follows that for all x and y ,

$$x \leq y \implies x \rightarrow y. \tag{6.1}$$

From Eq. 6.1 we see that all absorbing sets are increasing. Moreover, if a communicating class C is increasing, and if $x \rightarrow y$ for some $x \in C$, then there exists a z such that $x \leq z$ and $y \leq z$. Hence, $y \rightarrow z$ by Eq. 6.1 and $z \in C$, because C is increasing. Because C is a communicating class, it follows that $y \in C$. We may thus conclude that any communicating class C is absorbing if and only if it is increasing.

We next show that X has a unique absorbing communicating class. Assume first that all communicating classes are nonabsorbing. Then none of the communicating classes $C(x)$ is increasing, and Lemma 7 implies that $P_x(\tau_{C(x)^c} < \infty)$ for all x . Because $P_y(\tau_{C(x)} = \infty) = 1$ for all $y \notin C$, it follows that with probability one, X eventually leaves any finite set without ever returning, regardless of the initial state. Thus, $X(t) \rightarrow \infty$ almost surely, which contradicts the assumption that X started in some

initial state, say x^0 , is stable. Hence, X must have at least one absorbing communicating class. To see that there is no more than one such class, it suffices to observe that if $C(x)$ and $C(y)$ are disjoint sets, then they can not both be increasing.

Now let C be the unique absorbing communicating class of X , and assume that $X[x]$ is unstable for all $x \in C$. Then for any finite set K and any $0 \leq s \leq t$ it follows that

$$P_{x^0}(X(t) \in K, \tau_C \leq s) = \sum_{x \in C} P_{x^0}(X(s) = x, \tau_C \leq s) P_x(X(t-s) \in K), \tag{6.2}$$

because $X(s)$ belongs to the absorbing set C on the event $\{\tau_C \leq s\}$. Hence, by dominated convergence, $\lim_{t \rightarrow \infty} P_{x^0}(X(t) \in K, \tau_C \leq s) = 0$ for all s . Furthermore, because

$$P_{x^0}(X(t) \in K) \leq P_{x^0}(\tau_C > s) + P_{x^0}(X(t) \in K, \tau_C \leq s),$$

and because τ_C is finite almost surely, we see by taking first $t \rightarrow \infty$ and then $s \rightarrow \infty$ that $P_{x^0}(X(t) \in K) \rightarrow 0$, which contradicts the stability of $X[x^0]$. Hence, we may conclude that $X[y]$ is stable for some $y \in C$. Moreover, because $P_x(X(1) = y) > 0$ for all x , we see that for all finite sets K ,

$$\limsup_{t \rightarrow \infty} P_x(X(t+1) \in K) \geq \limsup_{t \rightarrow \infty} P_x(X(1) = y) P_y(X(t) \in K) > 0,$$

so $X[x]$ is stable for all initial states $x \in \mathbb{Z}_+^N$.

Finally, let X^C be the Markov process on the state space C with the same transition rates as X in C . Then X^C is irreducible and stable, regardless of the initial state. Hence, it follows (Kallenberg 2002, Theorem 12.25) that X^C is positive recurrent, and thus has a unique stationary distribution π^C on C such that the distribution of $X^C(t)$ converges to π^C in total variation. By defining $\pi(B) = \pi^C(B \cap C)$, it follows that π is stationary for the unrestricted version of X , and because $P_x(\tau_C < \infty)$ for all x , one can verify using Eq. 6.2 that the distribution of $X(t)$ converges to π in total variation, regardless of the initial state.

Having proved the equivalence of (1)–(3), let now assume that $X[x]$ is unstable for all x and show that $X(t) \rightarrow \infty$ in probability regardless of the initial state. We saw above that if all communicating classes of X are nonabsorbing, then $X(t)$ is transient, so let us assume that X has the unique absorbing class C . Then $X[x]$ is irreducible and positive recurrent for all $x \in C$, so it follows from standard theory (Kallenberg 2002, Theorem 12.25) that for any finite set A , the function $h_A(x, t) = P_x(X(t) \in A)$ tends to zero as $t \rightarrow \infty$ for all $x \in C$. Denoting the hitting time of X into C by τ_C , the strong Markov property implies that for all x ,

$$\begin{aligned} P_x(X(t) \in A) &= P(X(t) \in A, \tau_C \leq t) + P(X(t) \in A, \tau_C > t) \\ &= E_x h_A(X(\tau_C), t - \tau_C) 1(\tau_C \leq t) + P_x(X(t) \in A, \tau_C > t). \end{aligned}$$

Because $P_x(\tau_C < \infty) = 1$, it then follows from dominated convergence that the right-hand side in the above equality converges to zero as $t \rightarrow \infty$. Thus, $X(t) \rightarrow \infty$ in probability.

To see that $X(t) \rightarrow \infty$ implies the instability of X , let us assume that X is stable. Then by choosing a finite set A such that $\pi(A) > 0$, we see that $P_x(X(t) \in A) > 0$ for large t . This contradicts the fact that $X(t) \rightarrow \infty$ in probability, so X must be unstable. □

A1.3 Uniform limits of monotone functions

Proof of Proposition 7 (1) We show that when f is decreasing in all its input variables, the uniform limits of f are given by

$$f^0 = \inf_x f(x),$$

$$f^{n,\sigma}(x_{\sigma(1)}, \dots, x_{\sigma(n)}) = \inf_{x_{\sigma(n+1)}, \dots, x_{\sigma(N)}} f(x_1, \dots, x_N).$$

Observe first that given $\epsilon > 0$, there exists y such that $|f(y) - f^0| \leq \epsilon$. Hence, by defining $r = \max(y_1, \dots, y_N)$, it follows from the monotonicity of f that

$$\sup_{x: x_1, \dots, x_N > r} |f(x) - f^0| \leq \epsilon. \tag{6.3}$$

This shows that the assertion holds for $N = 1$.

To proceed by induction, let us assume that the claim holds for all positive decreasing functions on \mathbb{Z}_+^{N-1} . Let f be a positive and decreasing function on \mathbb{Z}_+^N , let $\epsilon > 0$, and choose a permutation σ . By symmetry, we assume without loss of generality that σ is the identity permutation, and denote $f^n = f^{n,\sigma}$. Using Eq. 6.3, we can first choose an r_0 such that $|f(x) - f^0| \leq \epsilon/2$ when $x_1, \dots, x_N > r_0$. Then by the monotonicity of f , it follows that $|f^n(x_1, \dots, x_n) - f^0| \leq \epsilon/2$ for all $x_1, \dots, x_n > r_0$. Thus,

$$|f(x) - f^n(x_1, \dots, x_n)| \leq \epsilon \tag{6.4}$$

for all x such that $x_1, \dots, x_N > r_0$.

Let us next choose an $i \in \{1, \dots, n\}$ and $y_i \in \{0, \dots, r_0\}$. Then the function

$$(x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_N) \mapsto f(x_1, \dots, x_{i-1}, y_i, x_{i+1}, \dots, x_N)$$

is decreasing on \mathbb{Z}_+^{N-1} . Hence, by the induction assumption we can choose a number $r_i(y_i)$ such that

$$|f(x) - \inf_{x_{n+1}, \dots, x_N} f(x)| \leq \epsilon$$

for all x such that $x_i = y_i$ and $x_{n+1}, \dots, x_N > r(y_i)$. In particular, by defining $r_i = \max(r_i(0), \dots, r_i(r_0))$, it follows that

$$|f(x) - f^n(x_1, \dots, x_n)| \leq \epsilon \tag{6.5}$$

for all x such that $x_i \leq r_0$ and $x_{n+1}, \dots, x_N > r_i$. Finally, by defining $r = \max(r_0, r_1, \dots, r_n)$, we see by combining Eqs. 6.4 and 6.5 that

$$\sup_{x \in \mathbb{Z}_+^N: x_{n+1}, \dots, x_N > r} |f(x) - f^n(x_1, \dots, x_n)| \leq \epsilon,$$

which completes the induction step.

(2) To see that fg has uniform limits at infinity, it suffices to note that for any n and any σ (omitting the arguments of the functions),

$$|fg - f^{n,\sigma}g^{n,\sigma}| \leq \|f\| \|g - g^{n,\sigma}\| + \|g\| \|f - f^{n,\sigma}\|,$$

and the same obviously holds for f^0 and g^0 in place of $f^{n,\sigma}$ and $g^{n,\sigma}$. Hence, the claim for fg follows by taking the supremum over $x \in \mathbb{Z}_+^N$ such that $x_{\sigma(n+1)}, \dots, x_{\sigma(N)} > r$ on both sides of the above inequality, and then letting r tend to infinity. The proof for $f + g$ is analogous. \square

References

- Asmussen S (2003) Applied probability and queues, 2nd edn. Springer
- Bonald T, Borst S, Hegde N, Proutière A (2004) Wireless data performance in multi-cell scenarios. In: Proc. ACM Sigmetrics/Performance 2004, pp 378–388
- Bonald T, Massoulié L, Proutière A, Virtamo J (2006) A queueing analysis of max-min fairness, proportional fairness and balanced fairness. Queueing Syst 53(1–2):65–84
- Cohen J, Boxma OJ (1983) Boundary value problems in queueing system analysis. North-Holland, Amsterdam
- Dai JG (1995) On positive Harris recurrence of multiclass queueing networks: a unified approach via fluid limit models. Ann Appl Probab 5(1):49–77
- de Veciana G, Lee T-J, Konstantopoulos T (2001) Stability and performance analysis of networks supporting elastic services. IEEE/ACM Trans Netw 9(1):2–14
- Fayolle G, Iasnogorodski R (1979) Two coupled processors: the reduction to a Riemann–Hilbert problem. Z Wahrsch Verw Gebiete 47(3):325–351
- Fayolle G, Malyshev VA, Menshikov MV (1995) Topics in the constructive theory of countable Markov chains. Cambridge University Press
- Jonckheere M, Borst SC (2006) Stability of multi-class queueing systems with state-dependent service rates. In: Proc. Valuetools'06
- Kallenberg O (2002) Foundations of modern probability, 2nd edn. Springer
- Kamae T, Krengel U, O'Brien GL (1977) Stochastic inequalities on partially ordered spaces. Ann Probab 5(6):899–912
- Liu X, Chong E, Shroff N (2003) A framework for opportunistic scheduling in wireless networks. Comp Netw 41:451–474
- Massey WA (1987) Stochastic orderings for Markov processes on partially ordered spaces. Math Oper Res 12(2):350–367
- Meyn SP (1995) Transience of multiclass queueing networks with via fluid limit models. Ann Appl Probab 5(4):946–957
- Meyn SP, Tweedie RL (1993) Markov chains and stochastic stability. Springer
- Neuts MF (1978) Markov chains with applications in queueing theory, which have a matrix-geometric invariant probability vector. Adv Appl Probab 10:185–212
- Rao RR, Ephremides A (1988) On the stability of interacting queues in a multiple-access system. IEEE Trans Inf Theory 34(5):918–930
- Robert P (2003) Stochastic networks and queues. Springer
- Rogers LCG, Williams D (1994) Diffusions, Markov processes, and martingales, vol I, 2nd edn. Wiley
- Szpankowski W (1988) Stability conditions for multidimensional queueing systems with computer applications. Oper Res 36(6):944–957
- Szpankowski W (1994) Stability conditions for some distributed systems: buffered random access systems. Adv Appl Probab 26(2):498–515
- Tweedie RL (1982) Operator-geometric stationary distributions for Markov chains, with application to queueing models. Adv Appl Probab 14:368–391



Sem Borst received the MSc degree in applied mathematics from the University of Twente, The Netherlands, in 1990, and the PhD degree from the University of Tilburg, The Netherlands, in 1994. During the fall of 1994, he was a visiting scholar at the Statistical Laboratory of the University of Cambridge, England. In 1995, he joined the Mathematics of Networks and Systems department of Bell Laboratories in Murray Hill, USA. He also holds a part-time appointment as a professor of Stochastic Operations Research at Eindhoven University of Technology. Sem Borst is a member of IFIP Working Group 7.3, and serves or has served as a member of several program committees and editorial boards. His main research interests are in the area of performance evaluation and resource allocation algorithms for communication networks.



Matthieu Jonckheere received the MSc degree from ENST Paris in 2002 and the PhD degree in applied mathematics from the Ecole Polytechnique, France, in 2005. During 2006-2007, he was a visiting scholar at Centrum voor Wiskunde en Informatica (CWI), Amsterdam and from January 2007 onwards, he has been appointed as an assistant professor in the mathematics and computer science department at Eindhoven University of Technology. His main research interests are in the area of applied probability and performance evaluation of communication networks.



Lasse Leskelä received the M.Sc. and D.Sc. degrees in mathematics from Helsinki University of Technology, Finland, in 1999 and 2005, respectively. Afterwards he spent one and half years visiting Columbia University, New York, USA, and Centrum voor Wiskunde en Informatica, Amsterdam, The Netherlands. In 2007 he was appointed assistant professor in the Department of Mathematics and Computer Science at the Eindhoven University of Technology, The Netherlands. His research interests include the analysis, stability, and control of stochastic processes, dynamical systems, and networks.