Mathematics of Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Heavy-Traffic Analysis of Queueing Systems with No Complete Resource Pooling

Daniela Andrea Hurtado Lange, Siva Theja Maguluri

MATHEMATICS OF OPERATIONS RESEARCH
Vol. 47, No. 4, November 2022, pp. 3129–3155
ISSN 0364-765X (print), ISSN 1526-5471 (online)
https://pubsonline.informs.org/journal/moor

Heavy-Traffic Analysis of Queueing Systems with No Complete Resource Pooling

Daniela Andrea Hurtado Lange,[a,*] Siva Theja Maguluri[b]

[a] The College of William and Mary, Williamsburg, Virginia 23187; [b] Georgia Institute of Technology, Atlanta, Georgia 30332
*Corresponding author

Contact: dahurtadolange@wm.edu, https://orcid.org/0000-0002-1073-4716 (DAHL); siva.theja@gatech.edu, https://orcid.org/0000-0002-5797-1639 (STM)

Received: November 12, 2019
Revised: March 1, 2021; September 22, 2021
Accepted: December 2, 2021
Published Online in Articles in Advance: March 2, 2022

MSC2020 Subject Classification: Primary: 60K25, 68M20, 90B22, 60H99

https://doi.org/10.1287/moor.2021.1248

Abstract. We study the heavy-traffic limit of the generalized switch operating under Max-Weight, without assuming that the complete resource pooling condition is satisfied and allowing for correlated arrivals. The main contribution of this paper is the steady-state mean of linear combinations of queue lengths in heavy traffic. We showcase the generality of our result by presenting various stochastic networks as corollaries, each of which is a contribution by itself. In particular, we study the input-queued switch with correlated arrivals, and we show that, if the state space collapses to a full-dimensional subspace, the correlation among the arrival processes does not matter in heavy traffic. We exemplify this last case with a parallel-server system, an $\mathcal{N}$-system, and an ad hoc wireless network. Whereas these results are obtained using the drift method, we additionally present a negative result showing a limitation of the drift method. We show that it is not possible to obtain the individual queue lengths using the drift method with polynomial test functions. We do this by presenting an alternate view of the drift method in terms of a system of linear equations, and we use this system of equations to obtain bounds on arbitrary linear combinations of the queue lengths.

Funding: This work was supported by ARC-TRIAD (Algorithms and Randomness Center-Transdisciplinary Research Institute for Advancing Data Science) Student Fellowships, Georgia Tech [Grant 3601410], National Science Foundation [Grant CCF-1850439], and CONICYT PFCHA (Comisión Nacional de Investigación Científica y Tecnológica-Programa Formación Capital Humano Avanzado)/DOCTORADO BECAS CHILE/2018 [Grant 72190413].
Supplemental Material: The online appendix is available at https://doi.org/10.1287/moor.2021.1248.

Keywords: drift method • state space collapse • generalized switch • input-queued switch • $\mathcal{N}$-system

1. Introduction

Resource allocation problems arise in a variety of settings, such as wireless networks, wired networks, data centers, cloud computing, ride-hailing systems, call centers, etc. One way to analyze them is to study the delay, using tools from queueing theory. In such cases, they are modeled as stochastic processing networks (SPNs). A major challenge is that the analysis of such queueing models is usually not tractable in general settings, and so asymptotic analysis is a popular methodology. Heavy-traffic analysis is an asymptotic approach by which the system is loaded very close to its capacity, and the corresponding queueing (delay) behavior of various resource allocation algorithms is studied.

Heavy-traffic limits are obtained in a wide variety of systems using a program based on fluid limits, diffusion limits, and reflected Brownian motion (RBM) processes as shown by Harrison [13]. In this approach, the queueing process is scaled appropriately, and the limiting fluid or diffusion process is studied. The limit of a diffusion-scaled process is shown to converge to an RBM process. Typically, this RBM lives in a lower dimensional subspace. This phenomenon is known as state space collapse (SSC), and it makes the heavy-traffic analysis tractable because one can study a lower dimensional RBM. Several systems in which the state space collapses to a line (i.e., to a one-dimensional subspace) are extensively studied in the literature using this approach (see the survey by Williams [30] for a rigorous list). Typically, this happens when there is a unique outer normal vector to the point of the boundary of the capacity region that is being approached in heavy traffic. Such systems are said to satisfy the complete resource pooling (CRP) condition. For a formal definition of the CRP condition, the reader is referred to Stolyar's [26] work. As shown by Harrison and López [14] and Dai and Lin [5], the intuitive meaning of the CRP is that, in the heavy-traffic limit, there is a single bottleneck resource, and hence, the queueing system behaves as a single server queue. Under the CRP condition, in the diffusion limit, one obtains an RBM on a line, which is well-understood. However, a major challenge is in using this program for SPNs in which the CRP condition is not satisfied (i.e., when there are multiple resources that are simultaneously in heavy traffic). In such cases, one needs to solve for the steady-state distribution of an RBM in a multidimensional subset of $\mathbb{R}^n$, and this is not known in general as shown by Kang and Williams [16]. The focus of this paper is to study systems that do not satisfy the CRP condition.

According to Kang and Williams [16] and Shah et al. [24], one of the simplest queueing systems in which the CRP condition is not satisfied is an input-queued switch, and Williams [30] identifies it as a focus of study in the SPN literature because it serves as a guiding example to study more general systems that do not satisfy CRP. Recently, the drift method was developed by Eryilmaz and Srikant [7], as an alternate way to study heavy-traffic limits of queueing systems based on a generalization of Kingman's [18] bound in a $G/G/1$ queue. The drift method was used to characterize the heavy-traffic scaled sum of queue lengths in input-queued switches by Maguluri and Srikant [20], and Maguluri et al. [21].

In order to study SPNs when the CRP condition is not met, in this paper, we consider a very general queueing model, called *generalized switch*, that subsumes several (single-hop) SPNs with control in the service process and was first proposed by Stolyar [26]. A detailed description of the model is provided in Section 2. Particular cases of the generalized switch are ad hoc wireless networks, wireless networks in the presence of fading, input-queued switches, and parallel-server systems.

In this paper, we study the generalized switch operating under a MaxWeight scheduling algorithm, which we describe in detail in Section 2. The MaxWeight algorithm was first proposed by Tassiulas and Ephremides [27] in the context of a down-link in wireless base stations and is used in a variety of queueing systems, for example, in the work by Stolyar [26], Gupta and Shroff [10], and Meyn [23]. Some of its advantages are that it is a throughput optimal algorithm (i.e., it keeps the system stable for all arrival rates in the capacity region), and it only requires information about the state of the system (and not parameters such as the arrival rates).

The generalized switch is studied under the CRP condition and independent arrivals assumption using both the diffusion limits approach by Stolyar [26] and the drift method by Eryilmaz and Srikant [7]. In this paper, we focus on the case when the CRP condition is not necessarily met, and so SSC may occur to a multidimensional subspace. Also, we assume that the arrival process to each queue is a sequence of independent and identically distributed (i.i.d.) random variables, but we do not require that these sequences are independent of each other. The main contributions of this paper are

i. In Theorem 1, we characterize the heavy-traffic scaled mean of certain linear combinations of the queue lengths in steady state under the MaxWeight algorithm. Moreover, we obtain lower and upper bounds that are valid in all regimes (not necessarily heavy traffic) but are tight in the heavy-traffic regime. This result is immediately applicable in several systems as we showcase in Section 4, and it includes both the CRP and the non-CRP cases. Little is known about SPNs that do not satisfy CRP because the most common approach in the literature is the use of diffusion limits, and solving a multidimensional RBM is an open question. In this paper, we contribute to understanding the heavy-traffic behavior of non-CRP systems by providing the mean of some linear combinations of the queue lengths.

ii. In Corollary 1, we compute the heavy-traffic limit of the total queue length in an input-queued switch with correlated arrivals. As mentioned, the input-queued switch has had considerable attention in the literature. However, it has only been studied under independent arrival processes by Maguluri and Srikant [20] and Maguluri et al. [21]. The input-queued switch is a model for an ideal data center network, and independent arrivals is an unrealistic assumption in this setting. In fact, as shown by Benson et al. [2] and Kandula et al. [15], for example, data centers experience hot spots, and hence, the arrivals to different queues are highly correlated.

iii. We illustrate how Theorem 1 can be immediately applied to a variety of systems. Specifically, we show how to apply it to parallel-server systems (Corollaries 4 and 5), the so-called $\mathcal{N}$-system (Corollary 6), and ad hoc wireless networks (Corollary 7).

iv. In Section 4.2, we show that, if SSC is full-dimensional, then the heavy-traffic limit of the mean queue lengths does not depend on the correlation among arrival processes. In other words, if the systems experience full-dimensional SSC, the expected linear combination of queue lengths behaves as if the queues were independent. This result is rather surprising, and it was not known.

v. In Theorem 2, we show that, using the drift method with polynomial test functions, it is impossible to obtain the moments of all linear combinations of the queue lengths. We prove this result by presenting an alternate way of thinking of the drift method. Traditionally, the key step in using this approach is to design the correct test function to obtain all the moments. However, it is not clear a priori if there are test functions that give all these moments. Instead of trying to guess the right test function, this point of view shows that one can think about solving a set of linear equations. This system of linear equations turns out to be underdetermined, and the major challenge is to obtain more equations using the constraints in the system in order to solve for all the unknowns and obtain the complete joint distribution of the queue lengths when the CRP condition is not satisfied.

vi. In Theorem 3, we obtain lower and upper bounds on the steady-state mean of an arbitrary linear combination of queue lengths. We do this by formulating a linear program (LP) using the underdetermined system of equations from Theorem 2. We present numerical results in the case of Bernoulli arrivals for different values of the traffic intensity. For simplicity of exposition, we do this only in the special case of an input-queued switch, and the same approach can be used for the generalized switch.

The second, third, and fourth contributions described are proved as corollaries of the main theorem (Theorem 1). However, they answer questions that, to the best of our knowledge, were open, and thus, they are contributions by themselves. This shows the versatility and power of Theorem 1.

The system of equations we propose in the fifth contribution presents an alternate view of the drift method, and it explains its success. Whereas it is known that it is notoriously hard to solve the stationary distribution of a multidimensional RBM, it is a little surprising that simple drift-based arguments give the mean of the sum of the queue lengths in several systems, such as the ones studied by Maguluri and Srikant [20], Maguluri et al. [21], and Wang et al. [28]. The system of equations shows that, because of the difficulty of the underlying problem, it is not possible to get all the mean queue lengths individually. However, because of the structure of the system of equations, it is possible to obtain certain linear combinations. In the case of input-queued switch and the bandwidth sharing system, the sum of the queue lengths is one of the linear combinations that is easy to obtain.

In the proof of Theorem 1, we use the drift method, and we work directly with the original queueing system without any fluid or diffusion scaling. The drift method consists of two main steps: (1) prove SSC and (2) compute asymptotically tight bounds. We additionally compute a universal lower bound (ULB) for a linear combination of the queue lengths, that is, a lower bound on the queue lengths that does not depend on the scheduling policy. We establish SSC in terms of certain moment bounds using a Lyapunov drift argument. By definition of steady state, the drift of any function with finite expectation is zero. We pick a quadratic test function and set its drift to zero in steady state to obtain the result. The choice of this test function is important in the drift method. We use the norm of the projection of the queue length vector into the space of the SSC as our test function, which was also used by Eryilmaz and Srikant [7], Maguluri and Srikant [20], and Maguluri et al. [21]. In this paper, we use this method in a discrete time system, but it can be also used in continuous time systems. For example, Wang et al. [28] use the drift method in the context of a bandwidth sharing network, which operates in continuous time.

The organization of the rest of this paper is as follows. We start in Section 1.1 establishing the main notation that is used in this paper. In Section 2, we describe the generalized switch model. In Section 3, we present the main result of this paper along with SSC and the ULB. Specifically, in Section 3.1, we present the ULB; in Section 3.2, we show SSC; and in Section 3.3, we present the main result of this paper (Theorem 1) together with remarks that help its interpretation. Then, in Section 4, we present relevant applications of our result, including the study of the input-queued switch under correlated arrivals (Corollary 1), full-dimensional SSC (Corollary 3), and the $\mathcal{N}$-system (Corollary 6). In Section 5, we present the alternate view of the drift method in the context of an input-queued switch (Theorem 2) and the linear programs to obtain bounds (Theorem 3). In Section 6, we present the proof of Theorems 1 and 2, and in Section 7, we present our conclusions and future work.

1.1. Notation

In this section, we introduce the notation that we use along the paper. We use $[n]$ to denote the set of integer numbers between one and n, both included. We use $\mathbb{R}$ to denote the set of real numbers and $\mathbb{Z}$ to denote the set of integer numbers. We add a subscript + to denote nonnegativity, and a superscript with a number to denote the dimension. We use bold letters to denote vectors and nonbold letters with a subscript to denote their elements. We write $\boldsymbol{x}=(x_1,x_2,\dots,x_n)$ for convenience, but we treat vectors as column vectors unless otherwise stated. Given two vectors $\boldsymbol{x}$ and $\boldsymbol{y}$, we write $\langle \boldsymbol{x},\boldsymbol{y}\rangle$ to denote the dot product between $\boldsymbol{x}$ and $\boldsymbol{y}$ and $\|\boldsymbol{x}\|$ to denote the Euclidean norm. For a matrix A, we write A^T to denote its transpose. Given two matrices A and B, we write $A\circ B$ to denote the Hadamard's product between A and B, that is, the matrix that results from multiplying term by term the elements of A and B. Let $\mathbb{I}_n$ be the identity matrix of $n\times n$. We use $\boldsymbol{e}^{(i,n)}$ to denote the i^{th} canonical vector in $\mathbb{R}^n$, that is, a vector with a one in the i^{th} element and zeros in all other entries. If the dimension is clear from the context, we omit the n.

For an irreducible and aperiodic Markov chain $\{X(k):k\in\mathbb{Z}_+\}$ over a countable state space $\mathcal{X}$, suppose $Z:\mathcal{X}\longrightarrow\mathbb{R}_+$ is a Lyapunov function. Define the drift of Z at x as

$$\Delta Z(x)\triangleq[Z(X(k+1))-Z(X(k))]\mathbb{1}_{\{X(k)=x\}}. \tag{1}$$

Thus, $\Delta Z(x)$ is a random variable that measures the amount of change in the value of Z in one time slot, starting from the state x.

2. Generalized Switch Model

In this section, we describe the model in detail, and we state known stability results. Consider n queues operating in discrete time, and let $\boldsymbol{q}(k)$ be the vector of queue lengths at the beginning of time slot k for each $k \in \mathbb{Z}_+$. For each $i \in [n]$, let $\{a_i(k) : k \in \mathbb{Z}_+\}$ be a sequence of i.i.d. random variables such that $a_i(k)$ is the number of arrivals to the i^{th} queue in time slot k. For each $i \in [n]$ let $\lambda_i \triangleq \mathbb{E}[a_i(1)]$, and $A_{\max}$ be a finite constant such that $a_i(1) \leq A_{\max}$ with probability one for all $i \in [n]$. Let Σ_a be the covariance matrix of the vector $\boldsymbol{a}(1)$.

The servers interfere with each other, so in each time slot, a set of interference constraints must be satisfied. Additionally, there are conditions of the environment that affect these constraints, which we group in a single random variable called the channel state. In other words, given the channel state, the set of interference constraints is known, and it may change if the channel state changes. Let $\{M(k) : k \in \mathbb{Z}_+\}$ be a sequence of i.i.d. random variables such that $M(k)$ is the channel state in time slot k. Let $\mathcal{M}$ be the state space of the channel state and ψ be the probability mass function of $M(1)$; that is, for each $m \in \mathcal{M}$, we define $\psi_m \triangleq P[M(1) = m]$. For each $m \in \mathcal{M}$, define $\mathcal{S}^{(m)}$ as the set of feasible service rate vectors in channel state m, that is, the set of vectors that satisfy the interference constraints in channel state m. For each $m \in \mathcal{M}$, the set $\mathcal{S}^{(m)}$ contains the projection on the coordinate axes of all its vectors. Formally, we assume that, if $\boldsymbol{x} \in \mathcal{S}^{(m)}$ for some $m \in \mathcal{M}$, then $\boldsymbol{x} - x_i\boldsymbol{e}^{(i)} \in \mathcal{S}^{(m)}$ for all $i \in [n]$. We assume that $\mathcal{M}$ is a finite set and that, for each $m \in \mathcal{M}$, the set $\mathcal{S}^{(m)}$ is finite. Therefore, the potential service offered to each queue in each time slot is bounded. Let $S_{\max}$ be a finite upper bound.

Let $\boldsymbol{s}(\boldsymbol{q}(k), M(k))$ be the vector of potential service in time slot k. In other words, for each $i \in [n]$, $s_i(\boldsymbol{q}(k), M(k))$ is the number of jobs from queue i that would be processed if there were enough jobs in line. Let $\boldsymbol{u}(\boldsymbol{q}(k), M(k), \boldsymbol{a}(k))$ be the vector of unused service in time slot k; that is, for each $i \in [n]$, $u_i(\boldsymbol{q}(k), M(k), \boldsymbol{a}(k))$ is the difference between the potential service and the number of packets that are actually processed from queue i in time slot k. For ease of exposition and with a slight abuse of notation, from now on, we use $\boldsymbol{s}(k)$ and $\boldsymbol{u}(k)$ to denote $\boldsymbol{s}(\boldsymbol{q}(k), M(k))$ and $\boldsymbol{u}(\boldsymbol{q}(k), M(k), \boldsymbol{a}(k))$, respectively.

In each time slot, a scheduling problem must be solved to decide which queues are served and at which rate. If queue i is not scheduled to receive service in time slot k, then $s_i(k) = 0$. In our model, the order of events in one time slot is as follows. First, the channel state is observed; second, a schedule is selected; third, arrivals occur; and at the end of each time slot, the jobs are processed according to the selected schedule. Hence, the dynamics of the queues are as follows. For each $k \in \mathbb{Z}_+$ and $i \in [n]$, we have

$$q_i(k+1) = q_i(k) + a_i(k) - s_i(k) + u_i(k). \tag{2}$$

In each queue, the unused service is nonzero only when the respective potential service is greater than the number of packets available (packets in line and arrivals). In such a case, the queue is empty in the next time slot. Therefore, the following equation is satisfied with probability one:

$$q_i(k+1)u_i(k) = 0 \qquad \forall k \in \mathbb{Z}_+, \ \ \forall i \in [n]. \tag{3}$$

However, if $i \neq j$, then $q_i(k+1)u_j(k)$ is not necessarily zero.

In this paper, we consider the generalized switch operating under the MaxWeight algorithm, which means that, in each time slot, the schedule with the longest total weighted queue length is selected for which the possible weight vectors are the feasible service rate vectors. Formally, if $M(k) = m$, then

$$\boldsymbol{s}(k) \in \arg\max_{\boldsymbol{x} \in \mathcal{S}^{(m)}} \langle \boldsymbol{q}(k), \boldsymbol{x} \rangle, \tag{4}$$

and ties are broken at random.

It is proved by Eryilmaz and Srikant [7] that the capacity region of the generalized switch is

$$\mathcal{C} = \sum_{m \in \mathcal{M}} \psi_m \, ConvexHull\,(\mathcal{S}^{(m)}) \tag{5}$$

$$= ConvexHull\left(\left\{\sum_{m \in \mathcal{M}} \psi_m \boldsymbol{x}^{(m)} : \boldsymbol{x}^{(m)} \in \mathcal{S}^{(m)} \ \ \forall m \in \mathcal{M}\right\}\right), \tag{6}$$

and the MaxWeight algorithm is throughput optimal. Then, the generalized switch operating under MaxWeight is positive recurrent for all $\boldsymbol{\lambda}$ in the interior of $\mathcal{C}$.

Because the sets $\mathcal{M}$ and $\mathcal{S}^{(m)}$ are finite for all $m \in \mathcal{M}$, the capacity region $\mathcal{C}$ is a polytope (bounded polyhedron) in $\mathbb{R}^n_+$. Then, we can describe it as the intersection of finitely many half-spaces. Let L be the minimum number of half-spaces that is needed to describe $\mathcal{C}$ and, for each $\ell \in [L]$, let $\boldsymbol{c}^{(\ell)}$ and $b^{(\ell)}$ be the parameters that define the ℓ^{th}

half-space. Then,

$$\mathcal{C} = \{\boldsymbol{x} \in \mathbb{R}^n_+ : \langle \boldsymbol{c}^{(\ell)}, \boldsymbol{x}\rangle \le b^{(\ell)},\ \ell = 1, \ldots, L\}. \tag{7}$$

Because the sets $\mathcal{S}^{(m)}$ contain the projection of their elements on the coordinate axes, the capacity region $\mathcal{C}$ is coordinate convex. Then, without loss of generality, we assume $\boldsymbol{c}^{(\ell)} \ge \boldsymbol{0}$ and $b^{(\ell)} > 0$ for all $\ell \in [L]$. For ease of exposition, we also assume $\|\boldsymbol{c}^{(\ell)}\| = 1$ for all $\ell \in [L]$. For each $\ell \in [L]$, let $\mathcal{F}^{(\ell)}$ be the ℓ^{th} facet of $\mathcal{C}$; that is, we define $\mathcal{F}^{(\ell)} \triangleq \{\boldsymbol{x} \in \mathcal{C} : \langle \boldsymbol{c}^{(\ell)}, \boldsymbol{x}\rangle = b^{(\ell)}\}$.

Observe that the schedules selected by the MaxWeight algorithm do not necessarily belong to the capacity region $\mathcal{C}$. This can be seen from (4) and (5) because $\psi_m \le 1$ for all $m \in \mathcal{M}$. However, the expected service rate vector does belong to the capacity region. We prove this result formally in Lemma 1.

Lemma 1. *Consider a generalized switch operating under MaxWeight as described, and let* $\mathbb{E}_{\boldsymbol{q}}[\cdot] \triangleq \mathbb{E}[\cdot \,|\, \boldsymbol{q}(k) = \boldsymbol{q}]$. *Then,* $\mathbb{E}_{\boldsymbol{q}}[\langle \boldsymbol{q}(k), \boldsymbol{s}(k)\rangle] = \max_{\boldsymbol{x} \in \mathcal{C}} \langle \boldsymbol{q}, \boldsymbol{x}\rangle$.

Proof of Lemma 1. Because $\boldsymbol{s}(k)$ is selected using the MaxWeight algorithm (see (4)), we have

$$\begin{aligned}\mathbb{E}_{\boldsymbol{q}}[\langle \boldsymbol{q}(k), \boldsymbol{s}(k)\rangle] &= \mathbb{E}_{\boldsymbol{q}}\left[\max_{\boldsymbol{x} \in \mathcal{S}^{(M(k))}} \langle \boldsymbol{q}(k), \boldsymbol{x}\rangle\right] \\ &\overset{(a)}{=} \sum_{m \in \mathcal{M}} \psi_m \max_{\boldsymbol{x} \in \mathcal{S}^{(m)}} \langle \boldsymbol{q}, \boldsymbol{x}\rangle \overset{(b)}{=} \max_{\boldsymbol{x} \in \mathcal{C}} \langle \boldsymbol{q}, \boldsymbol{x}\rangle,\end{aligned}$$

where (a) holds because the channel state process is independent from the queue lengths process, and (b) holds by definition of the capacity region $\mathcal{C}$ presented in (5). □

For technical reasons that are apparent in Section 3.3, we introduce the following definition. For each $\ell \in [L]$ and $m \in \mathcal{M}$ define the maximum $\boldsymbol{c}^{(\ell)}$-weighted service rate available when channel state is m as

$$b^{(m,\ell)} = \max_{\boldsymbol{x} \in \mathcal{S}^{(m)}} \langle \boldsymbol{c}^{(\ell)}, \boldsymbol{x}\rangle. \tag{8}$$

Observe that $\boldsymbol{c}^{(\ell)}$ and $b^{(m,\ell)}$ define a half-space that passes through the boundary of *ConvexHull*$(\mathcal{S}^{(m)})$, but this half-space does not necessarily define a facet of *ConvexHull*$(\mathcal{S}^{(m)})$. For each $\ell \in [L]$ and $k \in \mathbb{Z}_+$, let $B_\ell(k) \triangleq b^{(M(k),\ell)}$. Notice that $B_\ell(k)$ is an i.i.d. sequence of random variables that satisfies $P[B_\ell(1) = b^{(m,\ell)}] = \psi_m$ for each $m \in \mathcal{M}$. Let Σ_B be the covariance matrix of the vector $\boldsymbol{B}(1) \triangleq (B_1(1), \ldots, B_L(1))$; that is, for each $\ell_1, \ell_2 \in [L]$, we have

$$(\Sigma_B)_{\ell_1,\ell_2} \triangleq \mathbb{E}\,[B_{\ell_1}(k) B_{\ell_2}(k)] - \mathbb{E}[B_{\ell_1}(k)]\mathbb{E}[B_{\ell_2}(k)].$$

We model heavy traffic as follows. We fix a vector $\boldsymbol{\nu}$ in the boundary of $\mathcal{C}$, and we consider a set of generalized switches operating under MaxWeight as described, parameterized by $\epsilon \in (0,1)$. The heavy-traffic limit is the limit as $\epsilon \downarrow 0$, and as ϵ gets small, the vector of mean arrival rates approaches $\boldsymbol{\nu}$. Formally, we parameterize the queueing system in the following way. We let $\boldsymbol{q}^{(\epsilon)}(k)$, $\boldsymbol{a}^{(\epsilon)}(k)$, $\boldsymbol{s}^{(\epsilon)}(k)$, and $\boldsymbol{u}^{(\epsilon)}(k)$ be the vectors of queue lengths, arrivals, potential service, and unused service, respectively, in time slot k in the system parameterized by ϵ. The parameterization is such that the vector of mean arrival rate is $\boldsymbol{\lambda}^{(\epsilon)} \triangleq \mathbb{E}[\boldsymbol{a}^{(\epsilon)}(1)] = (1-\epsilon)\boldsymbol{\nu}$. Therefore, $\boldsymbol{\lambda}^{(\epsilon)}$ belongs to the interior of $\mathcal{C}$ for each $\epsilon \in (0,1)$, and as $\epsilon \downarrow 0$, the arrival rate vector $\boldsymbol{\lambda}^{(\epsilon)}$ approaches the boundary of the capacity region at the point $\boldsymbol{\nu}$.

Heavy-traffic analysis of the generalized switch has been performed in the past, using the diffusion limits approach by Stolyar [26], and the drift method by Eryilmaz and Srikant [7]. However, in both cases, the analysis is under the assumption that SSC occurs into a one-dimensional subspace (CRP condition), that is, when the vector $\boldsymbol{\nu}$ is in the interior of a facet of the capacity region $\mathcal{C}$. In this paper, we focus on cases in which the vector $\boldsymbol{\nu}$ may live at the intersection of facets. Define $P \triangleq \{\ell \in [L] : \boldsymbol{\nu} \in \mathcal{F}^{(\ell)}\}$; that is, P is the set of indices of all the facets that intersect at $\boldsymbol{\nu}$. Observe that, if P has only one element, we are under the CRP condition, and our results in this case agree with the results proved by Eryilmaz and Srikant [7]. In this paper, we focus on the case in which P is allowed to have more than one element.

For each $\epsilon \in (0,1)$, let $\overline{\boldsymbol{q}}^{(\epsilon)}$ be a steady-state random vector such that the Markov chain $\{\boldsymbol{q}^{(\epsilon)}(k) : k \in \mathbb{Z}_+\}$ converges in distribution to $\overline{\boldsymbol{q}}^{(\epsilon)}$ as $k \uparrow \infty$. Because MaxWeight is throughput optimal, the Markov chain $\{\boldsymbol{q}^{(\epsilon)}(k) : k \in \mathbb{Z}_+\}$ is positive recurrent for each $\epsilon \in (0,1)$, so $\overline{\boldsymbol{q}}^{(\epsilon)}$ is well-defined. Let $\overline{\boldsymbol{a}}^{(\epsilon)}$ be a steady-state vector that is equal in distribution to $\boldsymbol{a}^{(\epsilon)}(1)$. Then, $\mathbb{E}[\overline{\boldsymbol{a}}^{(\epsilon)}] = \boldsymbol{\lambda}^{(\epsilon)}$, and for each $i \in [n]$, we have $\overline{a}_i^{(\epsilon)} \le A_{\max}$ with probability one. Let $\Sigma_a^{(\epsilon)}$ be the covariance matrix of the vector $\overline{\boldsymbol{a}}^{(\epsilon)}$. Let $\overline{M}$ and $\overline{B}_\ell$ be steady-state random variables that are equal in

Figure 1. (Color online) Example of capacity region $\mathcal{C}$ and cone $\mathcal{K}$.

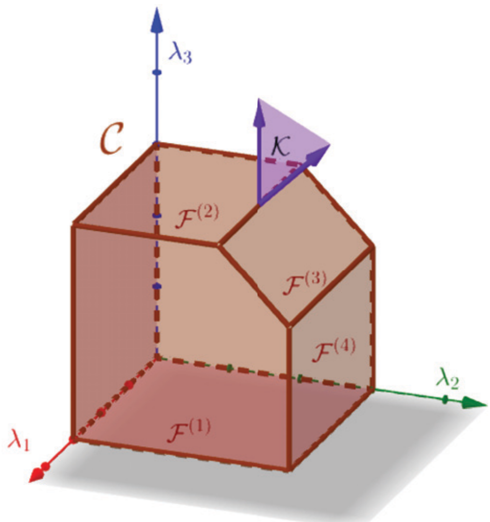

distribution to $M(1)$ and $B_\ell(1)$ for each $\ell \in [L]$, respectively. Let $\overline{\boldsymbol{s}}^{(\epsilon)} \triangleq \boldsymbol{s}(\overline{\boldsymbol{q}}^{(\epsilon)}, \overline{M})$ be the vector of potential service in steady-state, and $\overline{\boldsymbol{u}}^{(\epsilon)} \triangleq \boldsymbol{u}(\overline{\boldsymbol{q}}^{(\epsilon)}, \overline{M}, \overline{\boldsymbol{a}}^{(\epsilon)})$ be the vector of unused service. Define $(\overline{\boldsymbol{q}}^{(\epsilon)})^+ \triangleq \overline{\boldsymbol{q}}^{(\epsilon)} + \overline{\boldsymbol{a}}^{(\epsilon)} - \overline{\boldsymbol{s}}^{(\epsilon)} + \overline{\boldsymbol{u}}^{(\epsilon)}$ as the vector of queue lengths one time slot after $\overline{\boldsymbol{q}}^{(\epsilon)}$ is observed given that the vectors of arrivals and potential service are $\overline{\boldsymbol{a}}^{(\epsilon)}$ and $\overline{\boldsymbol{s}}^{(\epsilon)}$, respectively.

In Section 3.2, we prove that the state space collapses into the cone $\mathcal{K}$ described as follows. In other words, we show that the vector of queue lengths can be approximated by a vector in $\mathcal{K}$ in heavy traffic. Let $\mathcal{K}$ be the cone generated by $\{\boldsymbol{c}^{(\ell)} : \ell \in P\}$ and $\mathcal{H}$ be the subspace generated by the same set of vectors. Formally,

$$\mathcal{K} = \left\{ \boldsymbol{x} \in \mathbb{R}^n_+ : \boldsymbol{x} = \sum_{\ell \in P} \xi_\ell \boldsymbol{c}^{(\ell)},\ \xi_\ell \geq 0\ \ \forall \ell \in P \right\}. \tag{9}$$

A pictorial example of the capacity region $\mathcal{C}$ and the cone $\mathcal{K}$ when $n = 3$ is presented in Figure 1. Let $\tilde{P} \subset P$ be a set of indices such that the set $\{\boldsymbol{c}^{(\ell)} : \ell \in \tilde{P}\}$ is linearly independent, and let $C = [\boldsymbol{c}^{(\ell)}]_{\ell \in \tilde{P}}$ be a matrix in which the columns are a linearly independent subset of the vectors that generate the cone $\mathcal{K}$. Observe that the column space of the matrix C is exactly the subspace $\mathcal{H}$.

3. Heavy-Traffic Analysis of the Generalized Switch

In this section, we perform heavy-traffic analysis of the generalized switch. In Section 3.1, we present a ULB that is independent of the scheduling policy; in Section 3.2, we present the SSC result formally; and in Section 3.3, we present the main result of this paper (Theorem 1), in which we compute asymptotically tight bounds on linear combinations of the queue lengths.

3.1. Universal Lower Bound

In this section, we compute a ULB for certain linear combinations of the vector of queue lengths. The bound is universal in the sense that it remains valid for all scheduling policies.

Proposition 1. *Consider a generalized switch parameterized by* $\epsilon \in (0,1)$ *as described in Section 2. Let* $\boldsymbol{z} \in \mathcal{K}$ *with* $\boldsymbol{z} \neq \boldsymbol{0}$ *and for each* $\ell \in P$ *let* $r_\ell \geq 0$ *be such that* $\boldsymbol{z} = \sum_{\ell \in P} \boldsymbol{r}_\ell \boldsymbol{c}^{(\ell)}$. *Then, for each* $\epsilon \in (0,1)$, *we have*

$$\mathbb{E}\left[\langle \boldsymbol{z}, \overline{\boldsymbol{q}}^{(\epsilon)} \rangle\right] \geq \frac{1}{2\epsilon \langle \boldsymbol{z}, \boldsymbol{\nu} \rangle}\left(\boldsymbol{z}^T \Sigma_a^{(\epsilon)} \boldsymbol{z} + \boldsymbol{r}^T \Sigma_B \boldsymbol{r} \right) - f(\epsilon),$$

where $f(\epsilon) = \frac{b_{\max}\langle \boldsymbol{1}, \boldsymbol{r} \rangle}{2} - \frac{\epsilon \langle \boldsymbol{z}, \boldsymbol{\nu} \rangle}{2}$ *is* $o\left(\frac{1}{\epsilon}\right)$ *(i.e.,* $\lim_{\epsilon \downarrow 0} \epsilon f(\epsilon) = 0$*) and* $b_{\max} = \max_{m \in \mathcal{M}, \ell \in P} b^{(m,\ell)}$.

The proposition is proved by coupling the queue length vector of the generalized switch with a single-server queue $\{\Phi^{(\epsilon)}(k) : k \in \mathbb{Z}_+\}$ constructed as follows. We let $\alpha^{(\epsilon)}(k) \triangleq \langle \boldsymbol{z}, \boldsymbol{a}^{(\epsilon)}(k) \rangle$ be the number of arrivals in time slot k

and $\beta(k)$ be the potential service, in which $P[\beta(k)=\sum_{\ell\in P} r_\ell b^{(m,\ell)}]=\psi_m$ for each $m\in\mathcal{M}$. Then, it is easy to see that $\Phi^{(\epsilon)}(k)$ is stochastically smaller than $\langle \boldsymbol{z}, \boldsymbol{q}^{(\epsilon)}(k)\rangle$ (by definition of $b^{(m,\ell)}$ in (8)). Therefore, a lower bound to the expected value of $\Phi^{(\epsilon)}(k)$ in the steady state is also a lower bound to $\mathbb{E}[\langle \boldsymbol{z}, \overline{\boldsymbol{q}}^{(\epsilon)}\rangle]$. The last step in the proof is to compute such a lower bound, which we do by setting to zero the drift of $V_{ULB}(\Phi)=\Phi^2$. Note that it is essential that the weights r_ℓ are nonnegative to obtain a lower bound in the proof. This is the reason why $\boldsymbol{z}\in\mathcal{H}$ is not enough, and we require $\boldsymbol{z}\in\mathcal{K}$. The rest of the proof is presented in Online Appendix A.

3.2. State Space Collapse

We prove SSC into the cone $\mathcal{K}$ defined in (9) in heavy traffic. We start introducing the notation. For each $\epsilon\in(0,1)$, let $\boldsymbol{q}_{\|\mathcal{K}}^{(\epsilon)}(k)$ be the projection of $\boldsymbol{q}^{(\epsilon)}(k)$ on $\mathcal{K}$ and $\boldsymbol{q}_{\perp\mathcal{K}}^{(\epsilon)}(k)\triangleq\boldsymbol{q}^{(\epsilon)}(k)-\boldsymbol{q}_{\|\mathcal{K}}^{(\epsilon)}(k)$. Similarly, define $\boldsymbol{q}_{\|\mathcal{H}}^{(\epsilon)}(k)$ as the projection of $\boldsymbol{q}^{(\epsilon)}(k)$ on $\mathcal{H}$ and $\boldsymbol{q}_{\perp\mathcal{H}}^{(\epsilon)}(k)\triangleq\boldsymbol{q}^{(\epsilon)}(k)-\boldsymbol{q}_{\|\mathcal{H}}^{(\epsilon)}(k)$. We know the Markov chain $\{\boldsymbol{q}^{(\epsilon)}(k):k\in\mathbb{Z}_+\}$ is positive recurrent for each $\epsilon\in(0,1)$, so by definition of projection, we also have that $\{\boldsymbol{q}_{\|\mathcal{K}}^{(\epsilon)}(k):k\in\mathbb{Z}_+\}$, $\{\boldsymbol{q}_{\perp\mathcal{K}}^{(\epsilon)}(k):k\in\mathbb{Z}_+\}$, $\{\boldsymbol{q}_{\|\mathcal{H}}^{(\epsilon)}(k):k\in\mathbb{Z}_+\}$, and $\{\boldsymbol{q}_{\perp\mathcal{H}}^{(\epsilon)}(k):k\in\mathbb{Z}_+\}$ are positive recurrent for each $\epsilon\in(0,1)$. Then, we define $\overline{\boldsymbol{q}}_{\|\mathcal{K}}^{(\epsilon)}$, $\overline{\boldsymbol{q}}_{\perp\mathcal{K}}^{(\epsilon)}$, $\overline{\boldsymbol{q}}_{\|\mathcal{H}}^{(\epsilon)}$ and $\overline{\boldsymbol{q}}_{\perp\mathcal{H}}^{(\epsilon)}$ as steady-state vectors that are limited in the distribution of each of them, respectively. In the next proposition, we state SSC formally.

Proposition 2. *Given a vector $\boldsymbol{\nu}$ in the boundary of $\mathcal{C}$ and $\epsilon\in(0,1)$, consider a generalized switch operating under MaxWeight, parameterized by ϵ as described in Section 2, and let P be defined as in Section 2 as well. Let $\delta>0$ be such that $\delta\le b^{(\ell)}-\langle \boldsymbol{c}^{(\ell)},\boldsymbol{\nu}\rangle$ for all $\ell\in[L]\setminus P$ if $[L]\setminus P\neq\emptyset$, and $\delta=1$ if $[L]\setminus P=\emptyset$. If $\epsilon<\delta/(2\|\boldsymbol{\nu}\|)$, then for each $t=1,2,\ldots$, there exists a constant T_t such that $\mathbb{E}[\|\overline{\boldsymbol{q}}_{\perp\mathcal{H}}^{(\epsilon)}\|^t]\le\mathbb{E}[\|\overline{\boldsymbol{q}}_{\perp\mathcal{K}}^{(\epsilon)}\|^t]\le T_t$.*

We provide the proof of Proposition 2 in Online Appendix B. We adopt the technique introduced by Eryilmaz and Srikant [7], which is based on the bounds proved by Hajek [11]. Our proofs are similar, so we omit it for brevity. The challenges in obtaining our result arise in the second step of the drift method, which corresponds to Theorem 1.

SSC is a consequence of Proposition 2 for the following reason. As $\epsilon\downarrow 0$, $\|\overline{\boldsymbol{q}}^{(\epsilon)}\|$ goes to infinity (this can be easily concluded from Theorem 1). Therefore, Proposition 2 implies that, as ϵ gets small, we can approximate $\overline{\boldsymbol{q}}^{(\epsilon)}\approx\overline{\boldsymbol{q}}_{\|\mathcal{K}}^{(\epsilon)}$ because all the moments of $\|\overline{\boldsymbol{q}}_{\perp\mathcal{K}}^{(\epsilon)}\|$ are bounded.

Observe that the cone $\mathcal{K}$ is determined by the facets that intersect at $\boldsymbol{\nu}$. Moreover, the dimension of the cone is $n-d_{\boldsymbol{\nu}}$, where $d_{\boldsymbol{\nu}}$ is the dimension of the face of $\mathcal{C}$ where $\boldsymbol{\nu}$ is. For example, if $\boldsymbol{\nu}$ is in the relative interior of a facet, then $d_{\boldsymbol{\nu}}=n-1$, and this implies that $\mathcal{K}$ is one-dimensional. This is the CRP case, which was studied by Eryilmaz and Srikant [7] and Stolyar [26]. Similarly, if $\boldsymbol{\nu}$ is a vertex of $\mathcal{C}$, then $d_{\boldsymbol{\nu}}=0$, and hence, $\mathcal{K}$ is n-dimensional. In the last case, we say that SSC is full-dimensional. We study the full-dimensional case in Section 4.2.

3.3. Asymptotically Tight Bounds

In Section 3.2, we show SSC into the cone $\mathcal{K}$, which implies SSC into the subspace $\mathcal{H}$. In this section, we present the main result of this paper (Theorem 1), in which we provide asymptotically tight bounds to the expected value of certain linear combinations of the queue lengths in steady state. After the statement of the theorem, we present some remarks and applications, and we delay the proof to Section 6.1.

Theorem 1. *Given a vector $\boldsymbol{\nu}$ in the boundary of $\mathcal{C}$, let P be defined as in Section 2. Consider a set of generalized switches operating under MaxWeight, indexed by the heavy-traffic parameter $\epsilon\in(0,1)$ as described in Section 2. Then, for any vector $\boldsymbol{w}\in\cap_{\ell\in P}\mathcal{F}^{(\ell)}$, we have*

$$\left|\mathbb{E}\left[\left\langle\overline{\boldsymbol{q}}^{(\epsilon)},\boldsymbol{w}\right\rangle\right]-\frac{1}{2\epsilon}\mathbf{1}^T(H\circ\Sigma_a^{(\epsilon)})\mathbf{1}-\frac{1}{2\epsilon}\mathbf{1}^T((C^TC)^{-1}\circ\Sigma_B)\mathbf{1}\right|\le K(\epsilon), \tag{10}$$

where $\epsilon K(\epsilon)$ converges to zero as $\epsilon\downarrow 0$ and $H\triangleq C(C^TC)^{-1}C^T$ is the projection matrix into $\mathcal{H}$. Further, suppose $\lim_{\epsilon\downarrow 0}\Sigma_a^{(\epsilon)}=\Sigma_a$ component-wise. Then,

$$\lim_{\epsilon\downarrow 0}\epsilon\,\mathbb{E}\left[\left\langle\overline{\boldsymbol{q}}^{(\epsilon)},\boldsymbol{w}\right\rangle\right]=\frac{1}{2}(\mathbf{1}^T(H\circ\Sigma_a)\mathbf{1}+\mathbf{1}^T((C^TC)^{-1}\circ\Sigma_B)\mathbf{1}). \tag{11}$$

First, observe that (10) gives bounds that are valid for all regimes, not necessarily heavy traffic. Additionally, it shows that the queue lengths grow to infinity as the traffic intensity grows (i.e., as $\epsilon\downarrow 0$).

In (11), observe that the right-hand side has two terms: one corresponding to randomness in the arrival process and the other one to randomness in the service process. The first term is a linear combination of the covariance matrix of the arrival process, and the weights of the linear combination are determined by the projection matrix on the subspace $\mathcal{H}$, which is where SSC occurs. The second term is a linear combination of the elements of a covariance matrix that is related to the channel state. Because the potential service rate vector is selected using the MaxWeight algorithm (see (4)), it is not actually random once queue lengths and the channel state are observed. However, the channel state is a random variable that defines the feasible set in which MaxWeight is solved. Hence, the second term in (11), which includes a covariance matrix related to the channel state, represents the randomness on the service process.

A third observation is that, in order to project on the subspace $\mathcal{H}$ generated by the cone $\mathcal{K}$, we had to drop the vectors $\boldsymbol{c}^{(\ell)}$ with $\ell \in P$ that are linearly dependent (recall that the columns of the matrix C are a linearly independent subset of the vectors that generate $\mathcal{K}$). Clearly, the cone generated by the columns of C is not equal to $\mathcal{K}$. However, projecting on the subspace $\mathcal{H}$ is sufficient, and we do not need to worry about these linearly dependent vectors that we dropped.

In the next remark, we write (11) in different ways to facilitate interpretation of the result.

Remark 1. Equation (11) can be also written as

$$\lim_{\epsilon\downarrow 0} \epsilon\mathbb{E}\left[\left\langle \overline{\boldsymbol{q}}^{(\epsilon)}, \boldsymbol{w}\right\rangle\right] = \frac{1}{2}\left(\sum_{i=1}^{n}\sum_{j=1}^{n}\left\langle \boldsymbol{e}^{(i)}, \boldsymbol{e}^{(j)}_{\|\mathcal{H}}\right\rangle(\Sigma_a)_{i,j} + \sum_{\ell_1\in\tilde{P}}\sum_{\ell_2\in\tilde{P}}(C^TC)^{-1}_{\ell_1,\ell_2}(\Sigma_B)_{\ell_1,\ell_2}\right) \tag{12}$$

$$= \frac{1}{2}\left(Trace\,(H\Sigma_a^T) + Trace\left((C^TC)^{-1}\Sigma_B^T\right)\right), \tag{13}$$

where the subscript $\|\mathcal{H}$ denotes projection on the subspace $\mathcal{H}$, $(\Sigma_a)_{i,j}$ is the element (i,j) of the covariance matrix Σ_a for each $i,j\in[n]$, and $(\Sigma_B)_{\ell_1,\ell_2}$ is the element (ℓ_1,ℓ_2) of Σ_B for each $\ell_1,\ell_2\in\tilde{P}$.

In some cases, the projection of a vector on $\mathcal{H}$ is known in closed form, and it is simpler to work with than the projection matrix. For example, in the case of a completely saturated input-queued switch, Maguluri and Srikant [20] directly compute the projections, but writing down the projection matrix is more involved.

We present the proof of Remark 1 as follows.

Proof of Remark 1. If we expand the products on the right-hand side of (11), we obtain

$$\begin{aligned}
&\frac{1}{2}\left(\mathbf{1}^T(H\circ\Sigma_a)\mathbf{1} + \mathbf{1}^T\left((C^TC)^{-1}\circ\Sigma_B\right)\mathbf{1}\right)\\
&\overset{(a)}{=} \frac{1}{2}\left(\sum_{i=1}^{n}\sum_{j=1}^{n}h_{i,j}(\Sigma_a)_{i,j} + \sum_{\ell_1\in\tilde{P}}\sum_{\ell_2\in\tilde{P}}(C^TC)^{-1}_{\ell_1,\ell_2}(\Sigma_B)_{\ell_1,\ell_2}\right)\\
&\overset{(b)}{=} \frac{1}{2}\left(\sum_{i=1}^{n}\sum_{j=1}^{n}(\boldsymbol{e}^{(i)})^TH\boldsymbol{e}^{(j)}(\Sigma_a)_{i,j} + \sum_{\ell_1\in\tilde{P}}\sum_{\ell_2\in\tilde{P}}(C^TC)^{-1}_{\ell_1,\ell_2}(\Sigma_B)_{\ell_1,\ell_2}\right)\\
&\overset{(c)}{=} \frac{1}{2}\left(\sum_{i=1}^{n}\sum_{j=1}^{n}\left\langle \boldsymbol{e}^{(i)}, \boldsymbol{e}^{(j)}_{\|\mathcal{H}}\right\rangle(\Sigma_a)_{i,j} + \sum_{\ell_1\in\tilde{P}}\sum_{\ell_2\in\tilde{P}}(C^TC)^{-1}_{\ell_1,\ell_2}(\Sigma_B)_{\ell_1,\ell_2}\right),
\end{aligned}$$

where (*a*) holds by definition of Hadamard's product, (*b*) holds by definition of the canonical vectors $\boldsymbol{e}^{(i)}$ and by definition of the matrix product, and (*c*) holds by definition of the inner product and because $H\boldsymbol{e}^{(j)}$ is the projection of $\boldsymbol{e}^{(j)}$ on the subspace $\mathcal{H}$.

The proof of (13) holds by properties of Hadamard's product and trace, and we omit it. □

Observe that the bounds presented in Proposition 1 and Theorem 1 may be for different linear combinations of the vector of queue lengths. In Proposition 1, the vector of weights is $\boldsymbol{z}\in\mathcal{K}$, and in Theorem 1, it is $\boldsymbol{w}\in\cap_{\ell\in P}\mathcal{F}^{(\ell)}$. In the next remark, we give sufficient conditions under which these bounds correspond to the same linear combination of the queue lengths.

Remark 2. Let A be a matrix with columns $\boldsymbol{c}^{(\ell)}$ for $\ell\in P$ and $\boldsymbol{b}_P$ be a vector with elements $b^{(\ell)}$ for $\ell\in P$. Observe that the column space of A is equal to the column space of C, but the columns of A may not be linearly independent. In fact, if the columns of A are linearly independent, then $A=C$. Then, Proposition 1 and Theorem 1 give bounds to the same linear combination of the queue lengths if the set $\mathcal{A}\triangleq\{\boldsymbol{x}\in\mathbb{R}^{|P|} : \boldsymbol{x}^TA^TA\geq 0, \boldsymbol{x}^T\boldsymbol{b}_P<0\}$ is empty.

Proof of Remark 2. We can obtain bounds to the same linear combination of the queue lengths if there exists a vector $\boldsymbol{y} \in \mathcal{K} \cap (\cap_{\ell \in P} \mathcal{F}^{(\ell)})$—in other words, if the set $\mathcal{Y} \triangleq \{\boldsymbol{y} \in \mathbb{R}_+^{|P|} : AA^T\boldsymbol{y} = \boldsymbol{b}_P\}$ is nonempty. By Farkas' lemma (Bertsimas and Tsitsiklis [3, theorem 4.6]), proving that $\mathcal{Y} \neq \emptyset$ is equivalent to proving that $\mathcal{A} = \emptyset$. □

In the proof of Theorem 1, we use the drift method, which is a two-step procedure to compute bounds on linear combinations of the queue lengths that are tight in heavy traffic. The first step is to prove SSC, which we do in Proposition 2, and the second step is to set to zero the drift of $V(\boldsymbol{q}) = \|\boldsymbol{q}_{\|\mathcal{H}}\|^2$. Whereas these steps are standard for the drift method as developed by Eryilmaz and Srikant [7], Maguluri and Srikant [20], Maguluri et al. [21], and Wang et al. [28], different challenges arise in each case depending on the system one is studying. In this case, we are working with the generalized switch, which is a very general model. Hence, we overcome difficulties that are not part of the work listed. We summarize these:

a. Because the effective capacity region is the average of several individual capacity regions (see the definition of the capacity region in (5)), the vector of potential service does not necessarily belong to the effective capacity region. Then, it is not obvious how to deal with the terms that involve the service vector.

b. In the case of an input-queued switch as studied by Maguluri and Srikant [20], the projected service vector $\overline{\boldsymbol{s}}_{\|\mathcal{H}}$ is constant because of the structure of the system. In the case of the generalized switch, this is not the case, and this leads to significant challenges. In particular, the computation of the term $\mathbb{E}[\langle \overline{\boldsymbol{q}}_{\|\mathcal{H}}, \overline{\boldsymbol{s}}_{\|\mathcal{H}} \rangle]$ is not trivial. We use the properties of the system and the MaxWeight algorithm to bound this term.

c. The final closed-form expression that we obtain for the steady-state expectation of the queue lengths is novel and is a contribution in itself. To compute this expression (the term $\mathcal{T}_2$ in the proof), we use the least squares problem to obtain an expression that is valid for any generalized switch. Eryilmaz and Srikant [7], Maguluri and Srikant [20], and Maguluri et al. [21] explicitly use the underlying symmetry of the specific systems that are studied, and therefore, it is not clear how to generalize.

Challenge (a) is addressed in Lemmas 2 and 3. These lemmas form an important part of the entire proof and are used repeatedly. Challenge (b) is addressed in Claim 1. Finally, overcoming challenge (c) using the least squares problem gives us the closed-form expression for the right-hand side in Theorem 1.

4. Applications of Theorem 1

The generalized switch is a model that subsumes several SPNs, such as ad hoc wireless networks, the input-queued switch, down-link base stations, and the parallel-server system. In this section, we elaborate on a few applications to give examples of the use of Theorem 1, and it is by no means an exhaustive list. We start with an input-queued switch in Section 4.1, and then, in Section 4.2, we present examples in which full-dimensional SSC is observed.

4.1. Input-Queued Switch

The drift method is used to perform heavy-traffic analysis of the input-queued switch operating under MaxWeight in both completely and incompletely saturated cases by Maguluri and Srikant [20] and Maguluri et al. [21], respectively. In both scenarios, the analysis is performed under the assumption that the arrivals to different queues are independent. However, this is an unrealistic assumption in data center networks. Indeed, it is shown that the traffic exhibits hot spots; that is, there are subsets of queues that simultaneously perceive a surge on traffic as shown by Benson et al. [2] and Kandula et al. [15]. This implies that the arrival processes are highly correlated. In this section, we focus on the completely saturated input-queued switch, and we obtain the heavy-traffic limit of the scaled total queue length when the arrivals are correlated as a corollary of Theorem 1. Corollary 1 generalizes the main result proved by Maguluri and Srikant [20], and it is of special interest by itself given the nature of the arrival processes to data center networks observed in reality. We start specifying the model.

Consider a system with N^2 queues operating in discrete time. There are N input ports, N output ports, and a different queue for each input/output pair. Each of these pairs has its own arrival process, and all the arriving packets have the same size, which is equal to one time slot. The service process must satisfy the following feasibility constraints. In each time slot, at most one packet can be transmitted from each input port, and each output port can process at most one packet. We can think of this system as a matrix of input/output pairs, in which rows represent inputs and columns represent outputs. Then, the constraint described can be also stated as follows. In each time slot, at most one queue can be active (i.e., processing jobs) in each row and each column.

This model corresponds to a generalized switch with $n = N^2$ queues, in which the channel state is constant over time. As mentioned, the input-queued switch has a natural matrix-shape interpretation. Maguluri and Srikant [20] and Maguluri et al. [21] represent the vectors of queue lengths, arrivals, and services by $N \times N$ matrices, but they are treated as vectors. Specifically, dot products and norms are computed as if these matrices were column vectors. In this paper, however, we write them as column vectors to be consistent with the notation we introduced in Section 2. We enumerate the elements of the vectors row by row. For each $i \in [n]$, we have that $q_i(k)$ is the number of packets in line in input port $\lceil i/N \rceil$, waiting for service from output $(i \bmod N)$ if i is not a multiple of N and output N otherwise and similarly, for the vectors of arrivals, potential service, and unused service. In Figure 2, we show how to build the vectors in the case of $n = 2$ (Figure 2(a)) and $n = 3$ (Figure 2(b)).

For ease of exposition, we introduce the following notation. For each $i \in [N^2]$, let

$$\begin{aligned} row(i) &\triangleq \left\{ \left(\left\lceil \frac{i}{N} \right\rceil - 1 \right) N + j : j \in [N] \right\} \setminus \{i\} \\ col(i) &\triangleq \{ j \in [N^2] : i \bmod N = j \bmod N \} \setminus \{i\} \\ other(i) &\triangleq [N^2] \setminus (row(i) \cup col(i) \cup \{i\}). \end{aligned}$$

In words, the set $row(i)$ contains the index of all elements in the same row as i except by i, $col(i)$ contains the index of the elements in the same column as i except by i, and $other(i)$ contains all indexes that do not correspond to the same row or column as i or i itself.

We explicitly know the feasibility constraints in the input-queued switch. Then, we can compute the set of feasible service rate vectors $\mathcal{S}$ and the capacity region $\mathcal{C}$. We obtain

$$\mathcal{S} = \left\{ \boldsymbol{x} \in \{0,1\}^{N^2} : \sum_{i=1}^{N} x_{N(j-1)+i} \le 1 \ \ \forall j \in [N] \text{ and } \sum_{j=1}^{N} x_{N(j-1)+i} \le 1 \ \ \forall i \in [N] \right\},$$

and

$$\begin{aligned} \mathcal{C} &= ConvexHull(\mathcal{S}) \\ &= \left\{ \boldsymbol{x} \in \mathbb{R}_+^{N^2} : \sum_{i=1}^{N} x_{N(j-1)+i} \le 1 \ \ \forall j \in [N] \text{ and } \sum_{j=1}^{N} x_{N(j-1)+i} \le 1 \ \ \forall i \in [N] \right\}. \end{aligned} \tag{14}$$

Then, the number of hyperplanes that define the capacity region is $L = 2N$, the right-hand side parameters are $b^{(\ell)} = 1$ for all $\ell \in [2N]$, and the left-hand side vectors $\boldsymbol{c}^{(\ell)}$ are defined as follows:

$$\boldsymbol{c}^{(\ell)} = \begin{cases} \displaystyle\sum_{i=N(\ell-1)+1}^{N\ell} \boldsymbol{e}^{(i)}, & \text{if } \ell \in [N] \\ \displaystyle\sum_{i \in \{i' : i' \bmod N = \ell \bmod N\}} \boldsymbol{e}^{(i)}, & \text{if } \ell \in [2N] \setminus [N]. \end{cases} \tag{15}$$

A completely saturated switch means that the vector $\boldsymbol{\nu}$ that we approach in the heavy-traffic limit satisfies all the inequalities in (14) at equality. Formally, $\boldsymbol{\nu}$ satisfies $\langle \boldsymbol{c}^{(\ell)}, \boldsymbol{\nu} \rangle = b^{(\ell)}$ for all $\ell \in [2N]$. Then, $P = [2N]$. If $\boldsymbol{\nu}$ does not satisfy all the inequalities at equality, it is said that the switch is incompletely saturated. We do not study the incompletely saturated case in this paper.

Recall that the cone $\mathcal{K}$ in which SSC occurs is the cone generated by the vectors $\boldsymbol{c}^{(\ell)}$ with $\ell \in P$. In this case, because $P = [2N]$ and because we explicitly know the vectors $\boldsymbol{c}^{(\ell)}$, it can be easily proved that the cone $\mathcal{K}$ can be described as

$$\mathcal{K} = \left\{ \boldsymbol{x} \in \mathbb{R}_+^{N^2} : x_i = \frac{1}{N} \sum_{j \in row(i) \cup \{i\}} x_j + \frac{1}{N} \sum_{j \in col(i) \cup \{i\}} x_j - \frac{1}{N^2} \sum_{j=1}^{N^2} x_j \right\}. \tag{16}$$

Figure 2. Diagram of the queue length vector for the input-queued switch. (a) 2×2 switch. (b) 3×3 switch.

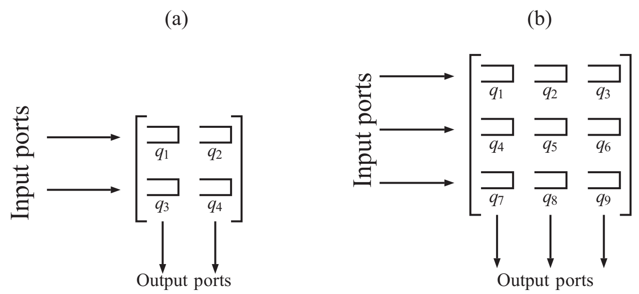

The proof of this claim is just algebra, and we omit it for brevity. In this case, it can be also proved that the subspace $\mathcal{H}$ generated by the cone $\mathcal{K}$ satisfies $\mathcal{K} = \mathcal{H} \cap \mathbb{R}_+^{N^2}$.

Now we present the heavy-traffic limit of the scaled sum of the queue lengths in a completely saturated switch with correlated arrival processes as a corollary of Theorem 1. This corollary by itself is a contribution of this paper because, to the best of our knowledge, the input-queued switch has been studied only under an independent arrivals assumption. However, it is known that, in data centers, this is not satisfied, and in fact, hot spots are frequently observed.

Corollary 1. *Let $\boldsymbol{\nu}$ be an N^2-dimensional vector that satisfies $\langle \boldsymbol{c}^{(\ell)}, \boldsymbol{\nu}\rangle = b^{(\ell)}$ for all $\ell \in [2N]$, for $\boldsymbol{c}^{(\ell)}$ as defined in (15) and $b^{(\ell)} = 1$ for all $\ell \in [2N]$. Consider a set of $N \times N$ input-queued switches as described, parameterized by $\epsilon \in (0,1)$ as described in Theorem 1. For each $i \in [N^2]$, let $\sigma_{a_i}^2 = (\Sigma_a)_{i,i}$. Then,*

$$\lim_{\epsilon \downarrow 0} \epsilon \mathbb{E}\left[\sum_{i=1}^{N^2} \overline{q}_i^{(\epsilon)}\right] = \frac{1}{2N}\sum_{i=1}^{N^2}\left((2N-1)\sigma_{a_i}^2 + (N-1)\sum_{j \in row(i)\cup col(i)} (\Sigma_a)_{i,j} - \sum_{j \in other(i)} (\Sigma_a)_{i,j}\right).$$

Proof of Corollary 1. We use Remark 1. We first compute $\boldsymbol{e}_{\|\mathcal{H}}^{(i)}$ for each $i \in [N^2]$. For any vector $\boldsymbol{y} \in \mathbb{R}_+^{N^2}$, we have that $\boldsymbol{y}_{\|\mathcal{H}}$ has elements

$$\boldsymbol{y}_{\|\mathcal{H}j} = \frac{1}{N}\sum_{j' \in row(j)\cup\{j\}} y_{j'} + \frac{1}{N}\sum_{j' \in col(j)\cup\{j\}} y_{j'} - \frac{1}{N^2}\sum_{j'=1}^{N^2} y_{j'} \qquad \forall j \in [N^2].$$

Then, for each $i \in [N^2]$, the vector $\boldsymbol{e}_{\|\mathcal{H}}^{(i)}$ has elements

$$\boldsymbol{e}_{\|\mathcal{H}j}^{(i)} = \begin{cases} \dfrac{2N-1}{N^2}, & \text{if } j = i \\ \dfrac{N-1}{N^2}, & \text{if } j \in row(i) \text{ or } j \in col(i) \\ -\dfrac{1}{N^2}, & \text{if } j \in other(i) \end{cases} \qquad \forall j \in [N^2].$$

Using this expression in Remark 1 we immediately obtain the result. □

Corollary 2. *Consider a set of $N \times N$ input-queued switches operating under MaxWeight, parameterized by $\epsilon \in (0,1)$ as described in Corollary 1. Further, assume that the arrival processes to different queues are independent. Then,*

$$\lim_{\epsilon \downarrow 0} \epsilon \mathbb{E}\left[\sum_{i=1}^{N^2} \overline{q}_i^{(\epsilon)}\right] = \left(1 - \frac{1}{2N}\right)\sum_{i=1}^{N^2} \sigma_{a_i}^2.$$

The proof of Corollary 2 is easy after considering Corollary 1 because $(\Sigma_a)_{i,j} = 0$ for all $i \neq j$ under the independent arrivals assumption. Corollary 2 recovers the main result presented by Maguluri and Srikant [20], in which they explicitly set to zero the drift of $V_{\|\mathcal{H}}(\boldsymbol{q}) = \|\boldsymbol{q}_{\|\mathcal{H}}\|^2$ (similarly to our approach in the proof of Theorem 1).

4.2. Full-Dimensional SSC

As mentioned in Section 3.2, if the point $\boldsymbol{\nu}$ is a vertex of the capacity region $\mathcal{C}$, the cone $\mathcal{K}$ is n-dimensional. In other words, $\mathcal{K}$ is full-dimensional. In this section, we explore this situation, and we present examples of SPNs in which this phenomenon is observed. In particular, we present the case of a parallel-server system operating in discrete time in Section 4.2.1, an $\mathcal{N}$-system in Section 4.2.2, and an ad hoc wireless network in Section 4.2.3. We first present the result in a general case.

Corollary 3. *Consider a set of generalized switches operating under MaxWeight, parameterized by $\epsilon \in (0,1)$ as described in Theorem 1. Let P, $\tilde{P}$, and $\boldsymbol{\nu}$ be as in Theorem 1 and suppose the cone $\mathcal{K}$ is n-dimensional. Let $\sigma_{a_i}^2 \triangleq (\Sigma_a)_{i,i}$ for each $i \in [n]$ and $\boldsymbol{\sigma}_a$ be a vector with elements σ_{a_i}. Then,*

$$\lim_{\epsilon \downarrow 0} \epsilon \mathbb{E}\left[\left\langle \overline{\boldsymbol{q}}^{(\epsilon)}, \boldsymbol{w}\right\rangle\right] = \frac{1}{2}\left(\|\boldsymbol{\sigma}_a\|^2 + \boldsymbol{1}^T\left((C^TC)^{-1} \circ \Sigma_B\right)\boldsymbol{1}\right).$$

Observe that Corollary 3 gives a rather surprising result. The right-hand side of the limit does not depend on the correlation among arrivals to different queues. In other words, in the heavy-traffic limit, these linear combinations of the queue lengths behave as if the arrival processes were independent if SSC is full-dimensional.

The proof of Corollary 3 follows immediately from Theorem 1 because, if the cone $\mathcal{K}$ is full-dimensional, then the subspace $\mathcal{H} = \mathbb{R}^n$, and therefore, the projection matrix on $\mathcal{H}$ satisfies $H = \mathbb{I}$.

In the rest of this section, we present examples of SPNs that experience full-dimensional SSC.

4.2.1. Parallel-Server System. Consider a parallel-server system as follows. There are n types of jobs that arrive according to arrival processes as described in Section 2. Each job type can be processed by a subset of servers, and these subsets are modeled by a compatibility graph. In Figure 3, we present three examples of parallel-server systems, in which the dotted lines represent the compatibility of the job types with the servers. In Figure 3(a), all jobs can be served by all servers (fully flexible system); in Figure 3(b), each job can be processed by only one server (dedicated system); and in Figure 3(c), the jobs from the first queue can be processed by any server, and the jobs from the second queue can only be processed by the second server ($\mathcal{N}$-system studied in Section 4.2.2). The parallel-server systems (also called process flexibility) receive plenty of attention in the literature. For example, see the work by Bell and Williams [1], Garnett and Mandelbaum [8], Harrison [12], Shi et al. [25], and Williams [29] and the survey paper by Williams [30]. However, most of the prior work is under the CRP condition. In this section, we show that the parallel-server system can be studied as an immediate application of Theorem 1 regardless of the CRP condition being satisfied.

To model a parallel-server system as a generalized switch, we assume that the service rate offered by each server in each time slot is a random variable that may depend on the service rate of other servers, but it is independent of the arrival and queueing processes. The joint distribution of the offered service rates is known, and we assume its state space is finite. Hence, the joint distribution of the offered service can be modeled as the channel state, and the compatibility graph determines the feasible service rate vectors in each time slot. Because we need the set of feasible service rate vectors in each channel state to be finite, we only consider the maximal vectors and their projection on the coordinate axes. Once the offered service rates are observed, the scheduler follows the MaxWeight algorithm to decide which job types are served and at which rate. We obtain the following result.

Corollary 4. *Consider a set of parallel-server systems as described, parameterized by ϵ as described in Theorem 1. Suppose the capacity region $\mathcal{C}$ has vertices that do not lie on the coordinate axes and that $\boldsymbol{w}$ is one of them. Let Σ_B be as in Theorem 1 and $\boldsymbol{\sigma}_a$ be as in Corollary 3. Then,*

$$\lim_{\epsilon\downarrow 0}\epsilon\mathbb{E}\left[\left\langle\overline{\boldsymbol{q}}^{(\epsilon)},\boldsymbol{w}\right\rangle\right] = \frac{1}{2}\left(\|\boldsymbol{\sigma}_a\|^2 + \mathbf{1}^T\left((C^TC)^{-1}\circ\Sigma_B\right)\mathbf{1}\right).$$

The proof of Corollary 4 only requires modeling the parallel-server system as a generalized switch as we show, so we omit it.

Remark 3. In Corollary 4, we consider a vector $\boldsymbol{w}$ in a vertex of the capacity region. However, Theorem 1 is immediately applicable for any $\boldsymbol{w}$ in the boundary of the capacity region. Here, we focus on a special case to illustrate the full-dimensional SSC result.

Before finishing this section, we present one of the simplest parallel-server systems to illustrate the result in Corollary 4. Specifically, we work with a dedicated system, in which every job type can be processed by exactly one server. A diagram with three job types and three servers is presented in Figure 3(b).

Consider an SPN with n servers, each with its own queue. Let $\{\hat{\boldsymbol{s}}(k) : k\in\mathbb{Z}_+\}$ be a sequence of i.i.d. random vectors such that $\hat{s}_i(k)$ is the potential service in queue i in time slot k. Let $\boldsymbol{\mu} = \mathbb{E}[\hat{\boldsymbol{s}}(1)]$ and Σ_s be the covariance matrix of $\hat{\boldsymbol{s}}(1)$. Suppose the vector $\hat{\boldsymbol{s}}(1)$ has finite state space and that $\hat{s}_i(1)\le S_{\max}$ with probability one for all $i\in[n]$. Suppose $\min_{i\in[n]}\mu_i > 0$.

Figure 3. (Color online) Diagrams of examples of parallel-server systems. The dotted lines represent the compatibility between job types and servers. (a) Fully flexible servers. (b) Dedicated servers. (c) $\mathcal{N}$-system.

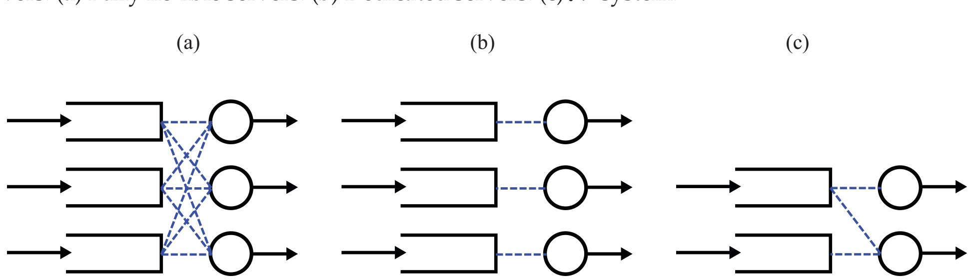

The arrival process is defined as in Section 2, and we model heavy traffic as described there as well. Specifically, let $\epsilon \in (0,1)$ be the heavy-traffic parameter. Then, for each $\epsilon \in (0,1)$ and each $i \in [n]$, let the arrival process to the system be $\{\boldsymbol{a}^{(\epsilon)}(k) : k \in \mathbb{Z}_+\}$, which is a sequence of i.i.d. random vectors with mean $\boldsymbol{\lambda}^{(\epsilon)} = \mathbb{E}[\boldsymbol{a}^{(\epsilon)}(1)] = (1-\epsilon)\boldsymbol{\mu}$ and covariance matrix $\Sigma_a^{(\epsilon)}$.

Corollary 5. *Consider a set of dedicated parallel-server systems as described, parameterized by $\epsilon \in (0,1)$ as described in Theorem 1. Suppose $\lim_{\epsilon\downarrow 0}\Sigma_a^{(\epsilon)} = \Sigma_a$ component-wise. Let $\sigma_{a_i}^2 = (\Sigma_a)_{i,i}$ and $\sigma_{s_i}^2 = (\Sigma_s)_{i,i}$ for each $i \in [n]$. Then,*

$$\lim_{\epsilon\downarrow 0} \epsilon \mathbb{E}\left[\sum_{i=1}^n \mu_i \overline{q}_i^{(\epsilon)}\right] = \frac{1}{2}\sum_{i=1}^n \left(\sigma_{a_i}^2 + \sigma_{s_i}^2\right).$$

From the discussion after Corollary 4, we expect that the correlation among the arrival processes would not be part of the right-hand side of the limit. However, observe that the correlation among the service processes does not appear in the answer either. Then, even though the arrival and potential service processes are correlated among queues, the mean linear combination of queue lengths behaves as if the queues were independent. Moreover, Corollary 5 recovers Kingman's bound. We present the proof of Corollary 5 as follows.

Proof of Corollary 5. The capacity region of this queueing system is

$$\mathcal{C} = \{\boldsymbol{x} \in \mathbb{R}_+^n : x_i \le \mu_i,\ i \in [n]\}.$$

To write it in the form of (7), we set $L = n$, and for each $i \in [n]$, we set $\boldsymbol{c}^{(i)} = \boldsymbol{e}^{(i)}$ and $b^{(i)} = \mu_i$. Therefore, the matrix C is the identity matrix, which implies that the projection matrix H is also the identity matrix.

Let $P = [n]$. Then, $\cap_{\ell \in P}\mathcal{F}^{(\ell)} = \{\boldsymbol{\mu}\}$, and the left-hand side of (11) yields

$$\lim_{\epsilon\downarrow 0} \epsilon \mathbb{E}\left[\sum_{i=1}^n \mu_i \overline{q}_i^{(\epsilon)}\right].$$

Because the projection matrix satisfies $H = \mathbb{I}$, the first term on the right-hand side of (11) yields

$$\frac{1}{2}\mathbf{1}^T(H \circ \Sigma_a)\mathbf{1} = \frac{1}{2}\mathbf{1}^T(\mathbb{I} \circ \Sigma_a)\mathbf{1} = \frac{1}{2}\sum_{i=1}^n \sigma_{a_i}^2.$$

To compute the second term of the right-hand side of (11), we consider the following interpretation of the channel state. Let $\mathcal{M}$ be an enumeration of the elements of the state space of $\hat{\boldsymbol{s}}(1)$ and $\boldsymbol{s}^{(m)}$ be its m^{th} element for each $m \in \mathcal{M}$. For each $m \in \mathcal{M}$, let the set of feasible service rate vectors in channel state m be

$$\mathcal{S}^{(m)} = \left\{\boldsymbol{s}^{(m)}\right\} \cup \left\{\boldsymbol{s}^{(m)} - s_i^{(m)}\boldsymbol{e}^{(i)} : i \in [n]\right\},$$

that is, the set $\mathcal{S}^{(m)}$ contains $\boldsymbol{s}^{(m)}$ and its projection on the coordinate axes. We assume that MaxWeight breaks ties by choosing maximal schedules. Then, if the channel state is m, then the service rates vector is always $\boldsymbol{s}^{(m)}$. With this assumption, we lose some generality because arrivals occur after deciding the optimal schedule. However, we are interested in heavy-traffic analysis, so this slight loss of generality does not affect our result. Then, the probability mass function of the channel state $\boldsymbol{\psi}$ satisfies $\psi_m \triangleq P[\hat{\boldsymbol{s}}(1) = \boldsymbol{s}^{(m)}]$ for each $m \in \mathcal{M}$.

By the definition of $b^{(m,\ell)}$ in (8) and definition of the sets $\mathcal{S}^{(m)}$ and the vectors $\boldsymbol{c}^{(\ell)}$, we obtain that, for each $\ell \in [n]$, we have

$$b^{(m,\ell)} = \left\langle \boldsymbol{c}^{(\ell)}, \boldsymbol{s}^{(m)} \right\rangle = \left\langle \boldsymbol{e}^{(\ell)}, \boldsymbol{s}^{(\ell)} \right\rangle = s_\ell^{(m)}.$$

Then, for each $\ell \in [n]$, the random variable $B_\ell(1)$ is such that $P[B_\ell(1) = s_\ell^{(m)}] = \psi_m$ and $\mathbb{E}[B_\ell(1)] = \mu_\ell$. Therefore, the vectors $(B_1(1), \ldots, B_n(1))$ and $\hat{\boldsymbol{s}}(1)$ have the same distribution. Hence, $(\Sigma_B)_{i,j} = \mathrm{Cov}[\hat{\boldsymbol{s}}_i, \hat{\boldsymbol{s}}_j]$, and the second term in the right-hand side of (11) becomes

$$\frac{1}{2}\mathbf{1}^T\left((C^TC)^{-1} \circ \Sigma_B\right)\mathbf{1} \overset{(a)}{=} \frac{1}{2}\mathbf{1}^T(\mathbb{I} \circ \Sigma_B)\mathbf{1} \overset{(b)}{=} \sum_{i=1}^n \sigma_{s_i}^2,$$

where (a) holds because $C = \mathbb{I}$, and (b) holds by definition of Hadamard's product and because the diagonal of Σ_B contains the variance of $\hat{s}_i(1)$ for each $i \in [n]$. □

4.2.2. $\mathcal{N}$-System. The $\mathcal{N}$-system model is a parallel-server system with two servers and two job types. One of the servers exclusively serves job type 1, and the other server can process both. A diagram of the $\mathcal{N}$-system is presented in Figure 3(c). According to Ghamami and Ward [9, p. 1], "The $\mathcal{N}$-system is one of the simplest parallel server system models that retains much of the complexity inherent in more general models." Consequently, it has received plenty of attention over the years, and there is vast literature that only focuses on its performance under the CRP condition. Examples can be found in the work by Bell and Williams [1], Garnett and Mandelbaum [8], Harrison [12], and Shi et al. [25]. Theorem 1 is immediately applicable to this system and gives information about the mean queue lengths in both the CRP and non-CRP cases. In this section, we focus on the non-CRP case.

Let the arrival processes be as described in Section 2 and suppose that each server processes jobs at rate 1. Then, the capacity region of this system is $\mathcal{C} = \{\boldsymbol{x} \in \mathbb{R}^2_+ : x_1 \le 1, x_2 \le 1\}$. We consider the heavy-traffic parameterization $\boldsymbol{\lambda}^{(\epsilon)} = (1-\epsilon)\mathbf{1}$ for $\epsilon \in (0,1)$. Then, as $\epsilon \downarrow 0$, the arrival rate vector approaches a vertex of the capacity region, and hence, the $\mathcal{N}$-system experiences full-dimensional SSC. We now present the result.

Corollary 6. *Consider a set of $\mathcal{N}$-systems parameterized by $\epsilon \in (0,1)$ as described. Let $\boldsymbol{\sigma}_a$ be as in Corollary 4. Then,*

$$\lim_{\epsilon\downarrow 0} \epsilon\mathbb{E}\left[\overline{q}_1^{(\epsilon)} + \overline{q}_2^{(\epsilon)}\right] = \frac{\sigma_{a_1}^2 + \sigma_{a_2}^2}{2}.$$

The proof is an immediate application of Corollary 3, so we omit it.

4.2.3. Ad Hoc Wireless Network. An ad hoc wireless network is composed of a set of nodes with no infrastructure for central coordination, and packets are transmitted between nodes (a transmitter and a receiver) if there is a link. The links interfere with each other, and therefore, not all of them can be active at the same time. These interference constraints are frequently represented with a graph, in which the vertices represent links and an edge between two links represents interference. In Figure 4(a), we present an example of the interference graph of an ad hoc wireless network with four links, in which all links interfere with each other. The packets to be transmitted arrive to each of the links and wait in line until they can be processed. This model is studied in a long line of literature, including but not limited to the work by Dimakis and Walrand [6], Eryilmaz and Srikant [7], Kang et al. [17], and Maguluri et al. [22], but in most of the cases, the focus is on studying stability or optimality of the scheduling policy. Here, we provide the heavy-traffic limit of linear combinations of the queue lengths under the MaxWeight algorithm. A particular case of our results is the results obtained by Eryilmaz and Srikant [7].

An ad hoc wireless network can be modeled as a generalized switch with a fixed channel state. Then, Theorem 1 can be immediately applied. In this section, we provide an example of an ad hoc wireless network that experiences full-dimensional SSC. We focus on a network with two links to illustrate the geometry of the capacity region and the cone in which SSC occurs, but similar work can be done for larger networks.

Let $\sigma_1^2 \triangleq \mathrm{Var}[\overline{a}_1^{(\epsilon)}]$, $\sigma_2^2 \triangleq \mathrm{Var}[\overline{a}_2^{(\epsilon)}]$, and $\varphi \triangleq \mathrm{Cov}[\overline{a}_1^{(\epsilon)}, \overline{a}_2^{(\epsilon)}]$, where these three parameters do not depend on ϵ. Suppose the set of feasible service rate vectors is $\mathcal{S} = \{(1,0), (0,1), (2/3, 2/3)\}$. Then, the capacity region is $\mathcal{C} = \{\boldsymbol{x} \in \mathbb{R}^2_+ : x_1 + 2x_2 \le 2, 2x_1 + x_2 \le 2\}$. Applying Theorem 1, we obtain the following corollary.

Figure 4. (Color online) Diagram of ad hoc wireless networks. (a) Example of an interference graph for an ad hoc wireless network with four links. (b) Capacity region and cone for SPN in Section 4.2.3.

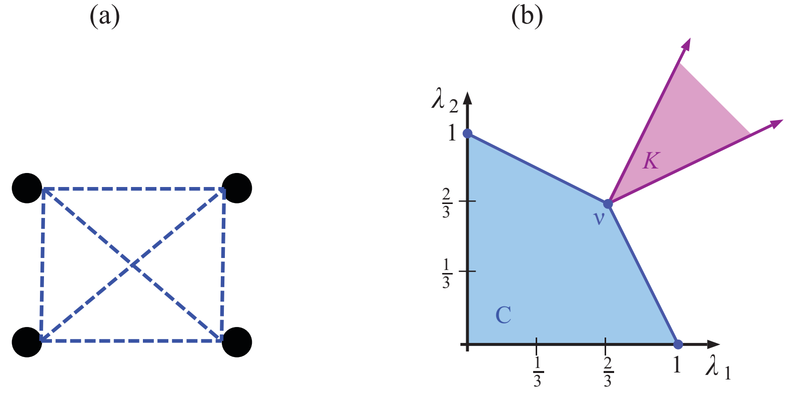

Corollary 7. *Consider an ad hoc wireless network as described. Then,*

$$\lim_{\epsilon \downarrow 0} \epsilon \mathbb{E}\left[\overline{q}_1^{(\epsilon)} + \overline{q}_2^{(\epsilon)}\right] = \frac{3}{4}(\sigma_1^2 + \sigma_2^2).$$

In the proof of Corollary 7, we take the heavy-traffic limit as the vector of arrival rate approaches the point $\boldsymbol{\nu} = (2/3)(1,1)$ in the boundary of $\mathcal{C}$. The proof is simple, so we omit it. In Figure 4(b), we plot the capacity region, the point $\boldsymbol{\nu}$, and the cone $\mathcal{K}$ in which SSC occurs. Observe that the cone $\mathcal{K}$ is two-dimensional, and therefore, this ad hoc wireless network experiences full-dimensional SSC.

Remark 4. The input-queued switch can be modeled similarly to an ad hoc wireless network. However, the input-queued switch cannot experience full-dimensional SSC because all the vertices of its capacity region are on the coordinate axes. In other words, all the vertices in the capacity region of the input-queued switch require the arrival rate to (at least) one queue to be zero. This is equivalent to considering a queueing system in which the zero-arrival rate queue does not exist, which already has a lower dimensional state space.

5. Individual Queue Lengths and Higher Moments in the Input-Queued Switch

In this section, we show that the drift method with polynomial test functions does not provide all the information that is necessary to compute the moments of all the linear combinations of the scaled queue lengths in systems that do not satisfy the CRP condition. We do this by presenting an alternate view of the drift method.

In the proof of Theorem 1, we use $V(\boldsymbol{q}) = \|\boldsymbol{q}_{\|\mathcal{H}}\|^2$ as a test function to obtain bounds on certain linear combinations of the queue lengths in a generalized switch. This choice of test function is first proposed by Maguluri and Srikant [20], and the main reason to use it is that the term $\mathcal{T}_4$ consisting of the "qu" terms (i.e., cross-terms between the queue length and the unused service) converge to zero in the heavy-traffic limit. All of queueing theory in some sense is to get a handle on the unused service terms, and the drift method handles these terms by making sure that they "cancel out" in heavy traffic, using SSC and our choice of the test function. In this section, instead of trying to cancel out the qu terms, we consider them as unknowns and try to solve for them along with the mean queue lengths. We see that this is impossible even if we use all possible quadratic test functions.

For simplicity of exposition, we present this result in the context of an input-queued switch, which is one of the simplest queueing systems that experience multidimensional SSC, and it is a special case of the generalized switch as shown in Section 4.1. The organization of this section is as follows. In Section 5.1, we present the main result; in Section 5.2, we use this result to compute bounds on the first moment of linear combinations of the scaled queue lengths; and in Section 5.3, we discuss how to generalize this approach to other queueing systems that experience multidimensional SSC.

5.1. System of Equations to Compute Linear Combinations of the First Moment of Scaled Queue Lengths

In this section, we prove that the drift method with polynomial test functions is not sufficient to compute all the linear combinations of the first moment of the scaled queue lengths in queueing systems that do not satisfy the CRP condition. Specifically, we show that the use of polynomial test functions yields an underdetermined system of equations.

In the drift method, one of the key challenges is to get a handle on the unused service. In general, when one sets to zero the drift of a polynomial test function in steady state, terms of the form $q_i(k+1)u_j(k)$ arise. The idea is to use a test function that captures the geometry of SSC so that we can show that all these cross-terms are small. Therefore, the choice of the test function is important, and the region into which SSC happens must be used in this choice. The quadratic test function, $V(\boldsymbol{q}) = \|\boldsymbol{q}_{\|\mathcal{H}}\|^2$ is successfully used by Eryilmaz and Srikant [7], Maguluri and Srikant [20], Maguluri et al. [21], and Wang et al. [28] to obtain the mean sum of the queue lengths, similarly to Theorem 1. Typically, one uses polynomial test functions of degree $(m+1)$ to get bounds on the expected value of the m^{th} power of the queue lengths. Therefore, in order to obtain bounds on the mean queue lengths, one must use quadratic test functions. In order to get all the linear combinations of the queue lengths, one can search through all the quadratic test functions, and this is equivalent to searching through all the quadratic monomials. The following theorem presents the result of using all the quadratic monomial test functions.

For ease of exposition, in this section, we prove our result in the case of $N = 2$ and independent arrivals, that is, in the case of a 2×2 input-queued switch with independent arrivals. We present generalizations to this result in Online Appendix C. Specifically, we present the case of a 2×2 input-queued switch with correlated arrivals in Online Appendix C.1 and the case of an $N \times N$ input-queued switch with independent arrivals in Online

Appendix C.2. The latter result can be easily generalized to the case of correlated arrivals, but we do not present the result here for ease of exposition.

Theorem 2. *Consider a set of* 2×2 *input-queued switches operating under MaxWeight, indexed by* $\epsilon\in(0,1)$ *as described in Corollary 2. Let* $\left(\Sigma_a^{(\epsilon)}\right)_{i,i}=\sigma_{a_i}^{(\epsilon)}$ *and suppose* $\lim_{\epsilon\downarrow 0}\sigma_{a_i}^{(\epsilon)}=\sigma_{a_i}$ *for all* $i\in[4]$*. Then, the following system of equations is satisfied:*

$$\lim_{\epsilon\downarrow 0}\epsilon\mathbb{E}[\overline{q}_1] = \frac{9\sigma_{a_1}^2+\sigma_{a_2}^2+\sigma_{a_3}^2+\sigma_{a_4}^2}{16}+\frac{1}{2}\lim_{\epsilon\downarrow 0}\mathbb{E}[\overline{q}_1^+(\overline{u}_2+\overline{u}_3)]-\frac{1}{2}\lim_{\epsilon\downarrow 0}\mathbb{E}[(\overline{q}_2^++\overline{q}_3^+)\overline{u}_4] \tag{17}$$

$$\lim_{\epsilon\downarrow 0}\epsilon\mathbb{E}[\overline{q}_2] = \frac{\sigma_{a_1}^2+9\sigma_{a_2}^2+\sigma_{a_3}^2+\sigma_{a_4}^2}{16}+\frac{1}{2}\lim_{\epsilon\downarrow 0}\mathbb{E}[\overline{q}_2^+(\overline{u}_1-\overline{u}_3+\overline{u}_4)] \tag{18}$$

$$\lim_{\epsilon\downarrow 0}\epsilon\mathbb{E}[\overline{q}_3] = \frac{\sigma_{a_1}^2+\sigma_{a_2}^2+9\sigma_{a_3}^2+\sigma_{a_4}^2}{16}+\frac{1}{2}\lim_{\epsilon\downarrow 0}\mathbb{E}[\overline{q}_3^+(\overline{u}_1-\overline{u}_2+\overline{u}_4)] \tag{19}$$

$$\lim_{\epsilon\downarrow 0}\epsilon\mathbb{E}[\overline{q}_1+\overline{q}_2] = \frac{3\sigma_{a_1}^2+3\sigma_{a_2}^2-\sigma_{a_3}^2-\sigma_{a_4}^2}{8}+\frac{1}{2}\lim_{\epsilon\downarrow 0}\mathbb{E}[\overline{q}_1^+(3\overline{u}_2-\overline{u}_3)]+\frac{1}{2}\lim_{\epsilon\downarrow 0}\mathbb{E}[\overline{q}_2^+(3\overline{u}_1+\overline{u}_3)]+\frac{1}{2}\lim_{\epsilon\downarrow 0}\mathbb{E}[\overline{q}_3^+\overline{u}_4] \tag{20}$$

$$\lim_{\epsilon\downarrow 0}\epsilon\mathbb{E}[\overline{q}_1+\overline{q}_3] = \frac{3\sigma_{a_1}^2-\sigma_{a_2}^2+3\sigma_{a_3}^2-\sigma_{a_4}^2}{8}+\frac{1}{2}\lim_{\epsilon\downarrow 0}\mathbb{E}[\overline{q}_1^+(-\overline{u}_2+3\overline{u}_3)]+\frac{1}{2}\lim_{\epsilon\downarrow 0}\mathbb{E}[\overline{q}_2^+\overline{u}_4]+\frac{1}{2}\lim_{\epsilon\downarrow 0}\mathbb{E}[\overline{q}_3^+(3\overline{u}_1+\overline{u}_2)] \tag{21}$$

$$\lim_{\epsilon\downarrow 0}\epsilon\mathbb{E}[\overline{q}_2+\overline{q}_3] = \frac{\sigma_{a_1}^2-3\sigma_{a_2}^2-3\sigma_{a_3}^2+\sigma_{a_4}^2}{8}+\frac{1}{2}\lim_{\epsilon\downarrow 0}\mathbb{E}[\overline{q}_2^+(\overline{u}_1+3\overline{u}_3+\overline{u}_4)]+\frac{1}{2}\lim_{\epsilon\downarrow 0}\mathbb{E}[\overline{q}_3^+(\overline{u}_1+3\overline{u}_2+\overline{u}_4)], \tag{22}$$

where we omit the dependence on ϵ *of the variables for ease of exposition.*

The proof of Theorem 2 is presented in Section 6.2. Observe that, in Theorem 2, we have system of six equations and 11 variables, in which the variables are

$$\begin{aligned}
&\lim_{\epsilon\downarrow 0}\epsilon\mathbb{E}[\overline{q}_1],\ \lim_{\epsilon\downarrow 0}\epsilon\mathbb{E}[\overline{q}_2],\ \lim_{\epsilon\downarrow 0}\epsilon\mathbb{E}[\overline{q}_3],\\
&\lim_{\epsilon\downarrow 0}\mathbb{E}[\overline{q}_1^+\overline{u}_2],\ \lim_{\epsilon\downarrow 0}\mathbb{E}[\overline{q}_1^+\overline{u}_3],\\
&\lim_{\epsilon\downarrow 0}\mathbb{E}[\overline{q}_2^+\overline{u}_1],\ \lim_{\epsilon\downarrow 0}\mathbb{E}[\overline{q}_2^+\overline{u}_3],\ \lim_{\epsilon\downarrow 0}\mathbb{E}[\overline{q}_2^+\overline{u}_4],\\
&\lim_{\epsilon\downarrow 0}\mathbb{E}[\overline{q}_3^+\overline{u}_1],\ \lim_{\epsilon\downarrow 0}\mathbb{E}[\overline{q}_3^+\overline{u}_2],\ \lim_{\epsilon\downarrow 0}\mathbb{E}[\overline{q}_3^+\overline{u}_4].
\end{aligned}$$

Therefore, it cannot be solved uniquely. However, a specific linear combination of the scaled queue lengths can be obtained as shown in the next corollary.

Corollary 8. *Consider a set of* 2×2 *input-queued switches as described in Theorem 2. Then,*

$$\lim_{\epsilon\downarrow 0}\epsilon\mathbb{E}[\overline{q}_2+\overline{q}_3]=\frac{3}{8}\left(\sigma_{a_1}^2+\sigma_{a_2}^2+\sigma_{a_3}^2+\sigma_{a_4}^2\right).$$

Proof of Corollary 8. Consider the following linear combination of the equations in Theorem 2:

$$(17)+(18)+(19)-\frac{1}{2}(20)-\frac{1}{2}(21)+\frac{1}{2}(22).$$

Then, reorganizing terms, we obtain the result. □

Corollary 8 can be also obtained as a consequence of Corollary 2 in the following way.

Alternative Proof of Corollary 8. From Corollary 2, for $N = 2$, we know

$$\lim_{\epsilon\downarrow 0} \epsilon\mathbb{E}\left[\overline{q}_1 + \overline{q}_2 + \overline{q}_3 + \overline{q}_4\right] = \frac{3}{4}\left(\sigma_{a_1}^2 + \sigma_{a_2}^2 + \sigma_{a_3}^2 + \sigma_{a_4}^2\right). \tag{23}$$

From SSC as proved in Proposition 2 and by definition of the cone $\mathcal{K}$ in (16), we also know that, for all $i \in [4]$, we have

$$\lim_{\epsilon\downarrow 0} \epsilon\mathbb{E}\left[\overline{q}_{\|\mathcal{H}i}\right] = \lim_{\epsilon\downarrow 0} \epsilon\mathbb{E}\left[\overline{q}_i\right],$$

where $\overline{q}_{\|\mathcal{H}i}$ is the i^{th} element of $\overline{\boldsymbol{q}}_{\|\mathcal{H}}$. Also, one interpretation of the cone $\mathcal{K}$ presented by Maguluri and Srikant [20] is that, for each vector in $\mathcal{K}$, all schedules have the same weight in the MaxWeight algorithm. This can be easily verified by definition of the cone $\mathcal{K}$ in (16). Then,

$$\overline{q}_{\|\mathcal{H}1} + \overline{q}_{\|\mathcal{H}4} = \overline{q}_{\|\mathcal{H}2} + \overline{q}_{\|\mathcal{H}3}. \tag{24}$$

Putting everything together, we obtain the result in Corollary 8. □

A special case in which we can solve for each of the expected individual queue lengths is the symmetric case, that is, when all the arrival processes have the same distribution. We present the result as follows.

Corollary 9. *Consider a set of 2×2 input-queued switches as described in Theorem 2, in which all the arrival processes have the same distribution. Let $\sigma_a^{(\epsilon)} \triangleq \sigma_{a_i}^{(\epsilon)}$ for all $i \in [4]$ and suppose $\lim_{\epsilon\downarrow 0}\sigma_a^{(\epsilon)} = \sigma_a$. Then, for each $i \in [4]$, we have*

$$\lim_{\epsilon\downarrow 0} \epsilon\mathbb{E}\left[\overline{q}_i\right] = \frac{3}{4}\sigma_a^2.$$

Proof of Corollary 9. In this case, because the arrivals are symmetric, all the queue lengths have the same expectation. Using this fact in Corollary 8, we obtain the result. □

In Theorem 2, we prove that setting to zero the drift of all monomials of degree two leads to a system of six equations and 11 variables. Therefore, the solution is not unique. However, Maguluri and Srikant [20] and Maguluri et al. [21] obtain the limit of specific linear combinations of the scaled queue lengths. These linear combinations can be obtained because some of the variables cancel out as shown in the first proof of Corollary 8. However, to obtain other linear combinations of the expected heavy-traffic scaled queue lengths, we need to actually work with all the variables of the system of equations. Therefore, we need additional equations.

To better understand this argument, consider a tandem queue system with memoryless interarrival and service times in any (not necessarily heavy) traffic. We know that the steady-state joint distribution is a product of two geometrics and can be obtained using reversibility arguments. Using the drift approach described, we get three equations and four unknowns. However, in addition to the drift arguments, if we use reversibility to separately prove that the queues are independent in the steady state and impose it as an additional condition, we can solve for all the unknowns.

5.2. Bounds on Linear Combinations of the Scaled Queue Lengths in Heavy Traffic

In Section 5.1, we present a linear system of equations that the vector of queue lengths must satisfy in heavy traffic. In this section, we use this system of equations to obtain bounds on linear combinations of the expected scaled queue lengths in heavy traffic. A similar approach is studied by Berstimas et al. [4] and Kumar and Kumar [19] in which an underdetermined set of linear systems of equations is obtained and linear programming is used to obtain bounds. However, the focus in those papers is on queueing networks under fixed arrival and service rates as opposed to the heavy-traffic analysis in the current paper.

In the next theorem, we provide upper and lower bounds for the heavy-traffic limit of the expected value of any linear combination of the queue lengths in a 2×2 input-queued switch.

Theorem 3. *Consider the equations*

$$v_1 - \frac{w_1 - w_2 + w_5 + w_8}{2} = \frac{9\sigma_{a_1}^2 + \sigma_{a_2}^2 + \sigma_{a_3}^2 + \sigma_{a_4}^2}{16}, \tag{25}$$

$$v_2 - \frac{w_3 + w_4 - w_5}{2} = \frac{\sigma_{a_1}^2 + 9\sigma_{a_2}^2 + \sigma_{a_3}^2 + \sigma_{a_4}^2}{16}, \tag{26}$$

$$v_3 - \frac{w_6 + w_7 - w_8}{2} = \frac{\sigma_{a_1}^2 + \sigma_{a_2}^2 + 9\sigma_{a_3}^2 + \sigma_{a_4}^2}{16}, \tag{27}$$

$$v_1 + v_2 - \frac{3w_+ w_2 - 3w_3 - w_4 - w_8}{2} = \frac{3\sigma_{a_1}^2 + 3\sigma_{a_2}^2 - \sigma_{a_3}^2 - \sigma_{a_4}^2}{8}, \tag{28}$$

$$v_1 + v_3 + \frac{w_1 - 3w_2 - w_5 - 3w_6 - w_7}{2} = \frac{3\sigma_{a_1}^2 - \sigma_{a_2}^2 + 3\sigma_{a_3}^2 - \sigma_{a_4}^2}{8}, \tag{29}$$

$$v_2 + v_3 - \frac{w_3 - 3w_4 - w_5 - w_6 - 3w_7 - w_8}{2} = \frac{\sigma_{a_1}^2 - 3\sigma_{a_2}^2 - 3\sigma_{a_3}^2 + \sigma_{a_4}^2}{8}, \tag{30}$$

$$-v_1 + v_2 + v_3 \geq 0, \tag{31}$$

$$-w_1 + w_7 \geq 0, \tag{32}$$

$$-w_2 + w_4 \geq 0, \tag{33}$$

and define $\mathcal{P} \triangleq \{(\boldsymbol{v}, \boldsymbol{w}) \in \mathbb{R}_+^3 \times \mathbb{R}_+^8 :$ *Equations* (25)–(33) *are satisfied*$\}$. *For* $\boldsymbol{\alpha} \in \mathbb{R}^3$, *define*

$$\underline{f}(\boldsymbol{\alpha}) \triangleq \min\{\langle \boldsymbol{\alpha}, \boldsymbol{v}\rangle : \exists \boldsymbol{w} \text{ such that } (\boldsymbol{v}, \boldsymbol{w}) \in \mathcal{P}\}$$

$$\text{and} \quad \overline{f}(\boldsymbol{\alpha}) \triangleq \max\{\langle \boldsymbol{\alpha}, \boldsymbol{v}\rangle : \exists \boldsymbol{w} \text{ such that } (\boldsymbol{v}, \boldsymbol{w}) \in \mathcal{P}\}.$$

Then,

$$\underline{f}(\boldsymbol{\alpha}) \leq \lim_{\epsilon \downarrow 0} \epsilon \mathbb{E}\left[\left\langle \boldsymbol{\alpha}, \overline{\boldsymbol{q}}^{(\epsilon)} \right\rangle\right] \leq \overline{f}(\boldsymbol{\alpha}), \tag{34}$$

where ϵ *and* $\overline{\boldsymbol{q}}^{(\epsilon)}$ *are defined as in Theorem* 2. *Furthermore, for any* $B \in \mathbb{R}_+$,

$$P\left[\lim_{\epsilon \downarrow 0} \epsilon \left\langle \boldsymbol{\alpha}, \overline{\boldsymbol{q}}^{(\epsilon)} \right\rangle \geq B\right] \leq \frac{\overline{f}(\boldsymbol{\alpha})}{B}. \tag{35}$$

Proof of Theorem 3. For ease of exposition, we omit the dependence on ϵ of the variables. Let

$$v_1 = \lim_{\epsilon \downarrow 0} \epsilon \mathbb{E}[\overline{q}_1],\ v_2 = \lim_{\epsilon \downarrow 0} \epsilon \mathbb{E}[\overline{q}_2],\ v_3 = \lim_{\epsilon \downarrow 0} \epsilon \mathbb{E}[\overline{q}_3],$$
$$w_1 = \lim_{\epsilon \downarrow 0} \mathbb{E}[\overline{q}_1^+ \overline{u}_2],\ w_2 = \lim_{\epsilon \downarrow 0} \mathbb{E}[\overline{q}_1^+ \overline{u}_3],$$
$$w_3 = \lim_{\epsilon \downarrow 0} \mathbb{E}[\overline{q}_2^+ \overline{u}_1],\ w_4 = \lim_{\epsilon \downarrow 0} \mathbb{E}[\overline{q}_2^+ \overline{u}_3],$$
$$w_5 = \lim_{\epsilon \downarrow 0} \mathbb{E}[\overline{q}_2^+ \overline{u}_4],\ w_6 = \lim_{\epsilon \downarrow 0} \mathbb{E}[\overline{q}_3^+ \overline{u}_1],$$
$$w_7 = \lim_{\epsilon \downarrow 0} \mathbb{E}[\overline{q}_3^+ \overline{u}_2],\ w_8 = \lim_{\epsilon \downarrow 0} \mathbb{E}[\overline{q}_3^+ \overline{u}_4].$$

Then, the proof of (34) follows from Theorem 2 because the set $\mathcal{P}$ represents the system of equations presented there together with nonnegativity constraints for all the variables. In particular, Inequalities (31)–(33) represent nonnegativity constraints associated to $\overline{q}_4$. These must be considered because, even though $\overline{q}_4$ does not appear in the system of equations explicitly, there are underlying constraints of the system related to $\overline{q}_4$ that affect its performance. Specifically, using (24) and the definition of the preceding variables, we obtain that the inequalities

$$\lim_{\epsilon \downarrow 0} \epsilon \mathbb{E}[\overline{q}_4] \geq 0, \quad \lim_{\epsilon \downarrow 0} \mathbb{E}[\overline{q}_4^+ \overline{u}_i] \geq 0 \ \ \forall i \in \{1, 2, 3\}$$

can be rewritten as (31)–(33) and $w_3 + w_6 \geq 0$, but the last inequality is implied by $w_3 \geq 0$ and $w_6 \geq 0$, so we do not write it in the definition of $\mathcal{P}$.

Also, from Markov's inequality, we know

$$P\left[\lim_{\epsilon \downarrow 0} \epsilon \left\langle \boldsymbol{\alpha}, \overline{\boldsymbol{q}}^{(\epsilon)} \right\rangle \geq B\right] \leq \frac{\lim_{\epsilon \downarrow 0} \epsilon \mathbb{E}\left[\left\langle \boldsymbol{\alpha}, \overline{\boldsymbol{q}}^{(\epsilon)} \right\rangle\right]}{B} \leq \frac{\overline{f}(\boldsymbol{\alpha})}{B},$$

where the last inequality holds by (34). □

Theorem 3 gives explicit bounds for all linear combinations of the expected scaled queue lengths. Similar linear programs can be written to obtain bounds on higher moments and, consequently, tighter tail probabilities.

In the rest of this section, we present numerical results to compare the bounds that we obtain from the linear program presented in Theorem 3 with the mean values that we obtain from simulation. We test four different

Table 1. Numerical results for LP with objective function $\lim_{\epsilon\downarrow 0}\epsilon\mathbb{E}[\overline{q}_2^{(\epsilon)}+\overline{q}_3^{(\epsilon)}]$.

ϵ	Solution to LP	Mean from simulation	Error, %
0.01	0.375	0.378	0.87
0.05	0.374	0.351	6.69
0.10	0.371	0.336	10.38

objective functions, viz. $\lim_{\epsilon\downarrow 0}\epsilon\mathbb{E}[\overline{q}_i]$ for $i\in\{1,2,3\}$ and $\lim_{\epsilon\downarrow 0}\epsilon\mathbb{E}[\overline{q}_2+\overline{q}_3]$. We use the last function because, in this case, the system of equations has a unique solution as shown in Corollary 8.

For simplicity, we assume that the arrivals to each queue are Bernoulli processes with mean $\lambda_i^{(\epsilon)}=(1-\epsilon)/2$ for all $i\in[4]$. We take $\epsilon\in\{0.01,0.05,0.1\}$ to evaluate the performance under different traffic intensities.

To allow the system to reach steady state, we ran the simulation for 10^9 time slots when $\epsilon\in\{0.05,0.1\}$ and for 10^{10} time slots in the case of $\epsilon=0.01$. The reason is that, for smaller ϵ, the system takes more time to reach steady state. In both cases, we compute the mean value of the variables considering the last 2×10^6 time slots.[1] We present our results in Tables 1 and 2. We ran three replicas of each experiment, and we obtained similar results. The results we present in Tables 1 and 2 are computed as an average of the three replicas.

In Table 1, we present the right-hand side of the expression proved in Corollary 8, the mean value of $\epsilon\left(\overline{q}_2^{(\epsilon)}+\overline{q}_3^{(\epsilon)}\right)$ obtained from the simulation, and the percentage error of the solution of the system of equations with respect to the simulation.

Observe that, as ϵ decreases, the solution to the LP becomes a better approximation for the simulated result. In fact, when $\epsilon=0.01$, the error is less than 1%. Even in the case of $\epsilon=0.1$, which is not considered heavy traffic, the error is around 10%.

In Table 2, we compute lower and upper bounds to the mean individual queue lengths, and we compare these results with the mean value of $\epsilon\overline{q}_1$, $\epsilon\overline{q}_2$, $\epsilon\overline{q}_3$, and $\epsilon\overline{q}_4$ obtained from simulation. The reason to present only one optimal value for all the queue lengths is that solving the linear program presented in Theorem 3 with objective function $\epsilon\mathbb{E}[\overline{q}_i^{(\epsilon)}]$ gives the same optimal value for all $i=1,2,3$ because of the symmetric arrival pattern. We additionally present the average between the minimum and maximum values of the individual queue lengths.

Observe that, for all the cases presented in Table 2, the mean obtained by simulation is between the lower and upper bounds obtained solving the LP. The bounds are not necessarily tight, but the average of both gives a good approximation of the mean individual queue lengths. Additionally, the LP presented in Theorem 1 is simple, and hence, it can be solved in fractions of a second as opposed to the simulation that may take hours.

5.3. Generalization to Other Queueing Systems and Higher Moments

In this section, we focused on a 2×2 input-queued switch in heavy traffic. We chose this system because it is one of the simplest queueing systems in which the CRP condition is not satisfied. However, the same approach can be applied to any queueing system in which the CRP condition is not met, which is what we discuss in this section. Specifically, we focus on a generalized switch with n queues, in which SSC occurs into a d-dimensional subspace.

Eryilmaz and Srikant [7] show how to compute the moments of $\|\boldsymbol{q}_{\|\mathcal{H}}\|$ using the drift method in queueing systems that satisfy the CRP condition. In this case, setting to zero the drift of $V(\boldsymbol{q})=\|\boldsymbol{q}_{\|\mathcal{H}}\|^{m+1}$ in the steady state and using SSC allows us to compute the m^{th} moment because of the following reason. When one sets to zero the drift of $V(\boldsymbol{q})$, terms of the form $q^+_{\|\mathcal{H}i}u_{\|\mathcal{H}i}$ arise, and because $\boldsymbol{q}^+_{\|\mathcal{H}}$ and $\boldsymbol{u}_{\|\mathcal{H}}$ belong to the same one-dimensional subspace, these terms can be approximated by $q_i^+u_i$, which is zero by definition of unused service.

On the other hand, if the CRP condition is not satisfied, then $\boldsymbol{q}$ lives in a d-dimensional subspace, where $d>1$. In this case, for each i, $q^+_{\|\mathcal{H}i}u_{\|\mathcal{H}i}$ cannot be approximated by $q_i^+u_i$ because of the following reason. In heavy traffic, we only have the approximation (with some abuse of notation) $q^+_{\|\mathcal{H}i}u_{\|\mathcal{H}i}\approx q_i^+(u_{k_1}+u_{k_2}+\ldots+u_{k_d})$, where $k_1,\ldots,k_d$

Table 2. Numerical results for individual queue lengths.

Value of ϵ	Minimum	Maximum	Average minimum and maximum	Simulation			
				Mean $\epsilon\overline{q}_1$	Mean $\epsilon\overline{q}_2$	Mean $\epsilon\overline{q}_3$	Mean $\epsilon\overline{q}_4$
0.01	0.062	0.312	0.187	0.187	0.187	0.192	0.191
0.05	0.062	0.312	0.187	0.174	0.176	0.175	0.176
0.10	0.062	0.309	0.186	0.168	0.168	0.168	0.169

represent the d dimensions that characterize SSC. In other words, cross-terms arise exactly as the qu terms in Theorems 2, 4, and 5 for the input-queued switch. In the following analysis, we present the number of equations and variables that appear in a general queueing system with d-dimensional SSC.

In order to obtain the m^{th} moment of the queue lengths, we should construct a system of equations that yields from setting to zero the drift of all the monomials of degree $m+1$. Because SSC occurs into a d-dimensional subspace, we need to consider all the possible monomials of degree $m+1$ in d variables. Setting to zero the drift of each monomial leads to an equation, so we have $\binom{m+d}{d-1}$ equations. Now we count the number of "new" variables with respect to the system of equations that arises after setting to zero the drift of monomials of degree k for all $k \leq m$. We say a variable is new for the system of equations that arises after setting to zero the monomials of degree $m+1$ if it does not appear in any system of equations of degree $k < m+1$. Observe that there are two types of new variables that do not vanish in the heavy-traffic limit. On one hand, we have the heavy-traffic limit of the expected value of products of the elements of $\boldsymbol{q}_{\|\mathcal{H}}$, and on the other hand, we have the heavy-traffic limit of the expected value of the product between the elements of $\boldsymbol{q}_{\|\mathcal{H}}$ and of the vector of unused service. We call them the q and qu variables, respectively. Specifically, the q variables are all monomials of degree m in d variables, so there are $\binom{m+d-1}{d-1}$ q variables. The qu variables that do not vanish in heavy traffic are of degree m in q and degree one in u. Also, the element corresponding to the unused service vector has to be different from the elements of the vector of queue lengths because the product between the queue length and the unused service of the same queue is zero by definition of unused service. Therefore, for each element of $\boldsymbol{u}_{\|\mathcal{H}}$, we need to consider all possible combinations of q's, that is, all monomials of degree m in $d-1$ variables. Therefore, there are $d\binom{m+d-2}{d-2}$ qu variables. Thus, in total, we have $\binom{m+d-1}{d-1}+d\binom{m+d-2}{d-2}$ variables, and this number is larger than the number of equations.

Summarizing, if we use the method introduced in this section to compute the m^{th} moment of the queue lengths of a queueing system that experiences d-dimensional SSC, we obtain a system of equations of $\binom{m+d}{d-1}$ equations and $\binom{m+d-1}{d-1}+d\binom{m+d-2}{d-2}$ variables. Therefore, it is underdetermined. In other words, we need extra equations to find a unique solution to this system of equations. This analysis shows that the issues illustrated in Theorem 2 arise in any queueing system with multidimensional SSC.

6. Proof of Theorems 1 and 2

In this section, we present the proofs of the main theorems of this paper.

6.1. Proof of Theorem 1

In this section, we present the proof of the main theorem. We use the notation $\mathbb{E}_m[\cdot] = \mathbb{E}[\cdot \,|\, \overline{M} = m]$, and we omit the dependence on ϵ of the variables for simplicity of exposition. Before presenting the proof of the theorem, we present two lemmas that formalize some intuition about the random variables $\overline{B}_\ell$ and are essential in the proof of Theorem 1.

Recall that $\overline{B}_\ell \triangleq b^{(\overline{M},\ell)}$ and that, for each $m \in \mathcal{M}$, $b^{(m,\ell)}$ is the maximum $\boldsymbol{c}^{(\ell)}$-weighted service rate in $\mathcal{S}^{(m)}$. Similarly, $b^{(\ell)}$ can be interpreted as the maximum $\boldsymbol{c}^{(\ell)}$-weighted service rate in $\mathcal{C}$, and $\boldsymbol{c}^{(\ell)}$ and $b^{(\ell)}$ define a facet of $\mathcal{C}$. Hence, because the values of $\overline{B}_\ell$ occur according to the probability mass function of the channel state and the capacity region $\mathcal{C}$ can be interpreted as the "expected capacity region" according to (5), we should expect $\mathbb{E}[\overline{B}_\ell] = b^{(\ell)}$. Additionally, $\boldsymbol{c}^{(\ell)}$ and $b^{(m,\ell)}$ define a half-space that passes through the boundary of $\mathit{ConvexHull}(\mathcal{S}^{(m)})$, and hence, there must exist a vector $\boldsymbol{\nu}^{(m)}$ such that $b^{(m,\ell)} = \langle \boldsymbol{c}^{(\ell)}, \boldsymbol{\nu}^{(m)} \rangle$. We formalize these results in Lemma 2.

Lemma 2. *Let $\ell \in P$ and $m \in \mathcal{M}$. Then, there exists $\boldsymbol{\nu}^{(m)} \in \mathcal{S}^{(m)}$ such that $b^{(m,\ell)} = \langle \boldsymbol{c}^{(\ell)}, \boldsymbol{\nu}^{(m)} \rangle$. This implies that $b^{(\ell)} = \mathbb{E}[\overline{B}_\ell]$ for all $\ell \in P$.*

The proof of Lemma 2 follows immediately from the definition of the capacity region $\mathcal{C}$ in (5) and of the parameters $b^{(m,\ell)}$ in (8). We present the details in Online Appendix D.1.

As ϵ gets closer to zero, we know that $\boldsymbol{\lambda}^{(\epsilon)}$ gets closer to $\boldsymbol{\nu}$, and SSC implies that the vector of queue lengths can be approximated by its projection on $\mathcal{K}$. In other words, as $\epsilon \downarrow 0$, the vector of queue lengths can be well-approximated by a conic combination of the vectors $\boldsymbol{c}^{(\ell)}$ with $\ell \in P$. Therefore, because the scheduling problem is solved using the MaxWeight algorithm and given that the channel state is m, one should expect that $\langle \boldsymbol{c}^{(\ell)}, \overline{\boldsymbol{s}} \rangle = b^{(m,\ell)}$ with high probability. In the next lemma, we formalize this intuition.

Lemma 3. *For each* $m \in \mathcal{M}$ *and* $\ell \in P$, *define* $\pi^{(m,\ell)} \triangleq P[\langle \boldsymbol{c}^{(\ell)}, \overline{\boldsymbol{s}} \rangle = b^{(m,\ell)} | \overline{M} = m]$. *Then,* $1 - \pi^{(m,\ell)}$ *is* $O(\epsilon)$.

The proof of Lemma 3 is a generalization of Eryilmaz and Srikant [7, claim 1], and we present it in Online Appendix D.2 for completeness.

Now we prove Theorem 1.

Proof of Theorem 1. First, observe that $\langle \overline{\boldsymbol{q}}, \boldsymbol{w} \rangle = \langle \overline{\boldsymbol{q}}_{\|\mathcal{H}}, \boldsymbol{\nu} \rangle$. To show this statement, define $\boldsymbol{w}_\perp \triangleq \boldsymbol{w} - \boldsymbol{\nu}$ for all $\boldsymbol{w} \in \cap_{\ell\in P}\mathcal{F}^{(\ell)}$ and observe that $\langle \boldsymbol{c}^{(\ell)}, \boldsymbol{w}_\perp \rangle = 0$ because both $\boldsymbol{\nu}, \boldsymbol{w} \in \mathcal{F}^{(\ell)}$ for all $\ell \in P$. Then,

$$\langle \overline{\boldsymbol{q}}_{\|\mathcal{H}}, \boldsymbol{\nu} \rangle = \langle \overline{\boldsymbol{q}}_{\|\mathcal{H}}, \boldsymbol{w} - \boldsymbol{w}_\perp \rangle = \langle \overline{\boldsymbol{q}}_{\|\mathcal{H}}, \boldsymbol{w} \rangle \overset{(*)}{=} \langle \overline{\boldsymbol{q}}, \boldsymbol{w} \rangle,$$

where $(*)$ holds because $\boldsymbol{w} \in \cap_{\ell\in P}\mathcal{F}^{(\ell)}$ and because $\overline{\boldsymbol{q}}_{\|\mathcal{H}} = \overline{\boldsymbol{q}} - \overline{\boldsymbol{q}}_{\perp\mathcal{H}}$. Hence, in the rest of the proof, we focus on computing bounds for $\mathbb{E}[\langle \overline{\boldsymbol{q}}_{\|\mathcal{H}}, \boldsymbol{\nu} \rangle]$.

We set to zero the drift of $V_{\|\mathcal{H}}(\boldsymbol{q}) = \|\boldsymbol{q}_{\|\mathcal{H}}\|^2$ and bound separately each of the terms that arise. Before setting the drift to zero, we need to make sure that $\mathbb{E}[V_{\|\mathcal{H}}(\overline{\boldsymbol{q}}_{\|\mathcal{H}})]$ is finite. This result can be proved using the Foster–Lyapunov theorem with Lyapunov function $Z(\boldsymbol{q}) = \|\boldsymbol{q}\|^2$. This proves that $\mathbb{E}[\|\overline{\boldsymbol{q}}\|^2]$ is finite. Then, because projection is nonexpansive, we have that $\mathbb{E}[\|\overline{\boldsymbol{q}}_{\|\mathcal{H}}\|^2]$ is also finite. The proof is simple, so we omit the details for ease of exposition. Now, setting to zero the drift of $V_{\|\mathcal{H}}(\boldsymbol{q})$, we obtain

$$\begin{aligned}
0 &= \mathbb{E}\left[\|\overline{\boldsymbol{q}}^+_{\|\mathcal{H}}\|^2 - \|\overline{\boldsymbol{q}}_{\|\mathcal{H}}\|^2\right] \\
&\overset{(a)}{=} \mathbb{E}\left[\|\overline{\boldsymbol{a}}_{\|\mathcal{H}} - \overline{\boldsymbol{s}}_{\|\mathcal{H}}\|^2 + 2\langle \overline{\boldsymbol{q}}_{\|\mathcal{H}}, \overline{\boldsymbol{a}}_{\|\mathcal{H}} - \overline{\boldsymbol{s}}_{\|\mathcal{H}} \rangle - \|\overline{\boldsymbol{u}}_{\|\mathcal{H}}\|^2 + 2\langle \overline{\boldsymbol{q}}^+_{\|\mathcal{H}}, \overline{\boldsymbol{u}}_{\|\mathcal{H}} \rangle\right],
\end{aligned} \tag{36}$$

where (a) holds by the dynamics of the queues presented in (2) and reorganizing terms. Let

$$\mathcal{T}_1 \triangleq 2\mathbb{E}\left[\langle \overline{\boldsymbol{q}}_{\|\mathcal{H}}, \overline{\boldsymbol{s}}_{\|\mathcal{H}} - \overline{\boldsymbol{a}}_{\|\mathcal{H}} \rangle\right], \quad \mathcal{T}_2 \triangleq \mathbb{E}\left[\|\overline{\boldsymbol{a}}_{\|\mathcal{H}} - \overline{\boldsymbol{s}}_{\|\mathcal{H}}\|^2\right],$$
$$\mathcal{T}_3 \triangleq \mathbb{E}\left[\|\overline{\boldsymbol{u}}_{\|\mathcal{H}}\|^2\right] \quad \text{and} \quad \mathcal{T}_4 \triangleq 2\mathbb{E}\left[\langle \overline{\boldsymbol{q}}^+_{\|\mathcal{H}}, \overline{\boldsymbol{u}}_{\|\mathcal{H}} \rangle\right].$$

Then, reorganizing the terms in (36), we obtain $\mathcal{T}_1 = \mathcal{T}_2 - \mathcal{T}_3 + \mathcal{T}_4$. We compute each term separately. We start with $\mathcal{T}_1$.

$$\begin{aligned}
\mathcal{T}_1 &\overset{(a)}{=} 2\mathbb{E}\left[\langle \overline{\boldsymbol{q}}_{\|\mathcal{H}}, \overline{\boldsymbol{s}} - \overline{\boldsymbol{a}} \rangle\right] \\
&\overset{(b)}{=} 2\epsilon\mathbb{E}\left[\langle \overline{\boldsymbol{q}}_{\|\mathcal{H}}, \boldsymbol{\nu} \rangle\right] + \mathbb{E}\left[\langle \overline{\boldsymbol{q}}_{\|\mathcal{H}}, \overline{\boldsymbol{s}} - \boldsymbol{\nu} \rangle\right] \\
&\overset{(c)}{=} 2\epsilon\mathbb{E}\left[\left\langle \overline{\boldsymbol{q}}_{\|\mathcal{H}}, \boldsymbol{\nu} \right\rangle\right] + O(\sqrt{\epsilon}),
\end{aligned} \tag{37}$$

where (a) holds by the orthogonality principle, (b) holds because $\mathbb{E}[\overline{\boldsymbol{a}}] = (1-\epsilon)\boldsymbol{\nu}$ and because $\overline{\boldsymbol{a}}$ is independent of the vector of queue lengths, and (c) holds by Claim 1.

Claim 1. *Consider a set of generalized switches as described in Theorem 1. Then,*

$$\left|\mathbb{E}\left[\left\langle \overline{\boldsymbol{q}}_{\|\mathcal{H}}, \overline{\boldsymbol{s}} - \boldsymbol{\nu} \right\rangle\right]\right| \text{ is } O(\sqrt{\epsilon}).$$

We present the proof of Claim 1 in Online Appendix D.3. Now, we compute $\mathcal{T}_2$. Expanding the product, we obtain

$$\mathcal{T}_2 = \mathbb{E}\left[\|\overline{\boldsymbol{a}}_{\|\mathcal{H}} - \overline{\boldsymbol{s}}_{\|\mathcal{H}}\|^2\right] = \mathbb{E}\left[\|\overline{\boldsymbol{a}}_{\|\mathcal{H}}\|^2\right] + \mathbb{E}\left[\|\overline{\boldsymbol{s}}_{\|\mathcal{H}}\|^2\right] - 2\mathbb{E}\left[\langle \overline{\boldsymbol{a}}_{\|\mathcal{H}}, \overline{\boldsymbol{s}}_{\|\mathcal{H}} \rangle\right]. \tag{38}$$

We compute each term in (38) separately. For the first two terms, we solve the least squares problem, and we use the projection matrix on the subspace $\mathcal{H}$, denoted as H. Let $h_{i,j}$ be its element (i, j) for each $i, j \in [n]$. For the first term, we have

$$\begin{aligned}
\mathbb{E}\left[\|\overline{\boldsymbol{a}}_{\|\mathcal{H}}\|^2\right] &= \mathbb{E}\left[\|H\overline{\boldsymbol{a}}\|^2\right] \\
&\overset{(a)}{=} \sum_{i=1}^n\sum_{j=1}^n h_{i,j}\mathrm{Cov}[\overline{a}_i, \overline{a}_j] + \sum_{i=1}^n\sum_{j=1}^n h_{i,j}\mathbb{E}[\overline{a}_i]\mathbb{E}[\overline{a}_j] \\
&\overset{(b)}{=} \mathbf{1}^T\left(H \circ \Sigma_a^{(\epsilon)}\right)\mathbf{1} + (1-\epsilon)^2\boldsymbol{\nu}^T H\boldsymbol{\nu},
\end{aligned} \tag{39}$$

where (a) holds solving the least squares problem by definition of the norm because H is a projection matrix (and, therefore, $H = H^T = H^2$) and by definition of covariance, and (b) holds by definition of the Hadamard's product and because $\mathbb{E}[\overline{a}_i] = \lambda_i^{(\epsilon)} = (1-\epsilon)\nu_i$ for each $i \in [n]$. For the second term in (38), we obtain

$$\begin{aligned}
\mathbb{E}\left[\|\overline{\boldsymbol{s}}_{\|\mathcal{H}}\|^2\right] &= \mathbb{E}\left[\|H\overline{\boldsymbol{s}}\|^2\right] \\
&\overset{(a)}{=} \mathbb{E}\left[\overline{\boldsymbol{s}}^T C(C^TC)^{-1}C^T\overline{\boldsymbol{s}}\right] \\
&\overset{(b)}{=} \sum_{\ell_1\in\tilde{P}}\sum_{\ell_2\in\tilde{P}}(C^TC)^{-1}_{\ell_1,\ell_2}\mathbb{E}\left[\langle \boldsymbol{c}^{(\ell_1)},\overline{\boldsymbol{s}}\rangle\langle \boldsymbol{c}^{(\ell_2)},\overline{\boldsymbol{s}}\rangle\right] \\
&\overset{(c)}{=} \sum_{\ell_1\in\tilde{P}}\sum_{\ell_2\in\tilde{P}}(C^TC)^{-1}_{\ell_1,\ell_2}\sum_{m\in\mathcal{M}}\psi_m\mathbb{E}_m\left[\langle \boldsymbol{c}^{(\ell_1)},\overline{\boldsymbol{s}}\rangle\langle \boldsymbol{c}^{(\ell_2)},\overline{\boldsymbol{s}}\rangle\right] \\
&\overset{(d)}{=} \boldsymbol{1}^T\left((C^TC)^{-1}\circ\Sigma_B\right)\boldsymbol{1} + \boldsymbol{\nu}^T H\boldsymbol{\nu} - O(\epsilon).
\end{aligned} \tag{40}$$

Here, $(C^TC)^{-1}_{\ell_1,\ell_2}$ is the element (ℓ_1,ℓ_2) of the matrix $(C^TC)^{-1}$ for each $\ell_1,\ell_2 \in \tilde{P}$. Here, ($a$) holds solving the least squares problems and because $H = C(C^TC)^{-1}C^T$ by definition of the projection matrix; (b) holds by definition of matrix multiplication and because $C^T\overline{\boldsymbol{s}}$ is a vector with elements $\langle \boldsymbol{c}^{(\ell)},\overline{\boldsymbol{s}}\rangle$ for $\ell\in\tilde{P}$; (c) holds by the law of total probability, conditioning on the channel state; and (d) holds using Lemma 3, the definition of covariance, and reorganizing the terms. Now, we compute the last term in (38). We obtain

$$-2\mathbb{E}\left[\langle\overline{\boldsymbol{a}}_{\|\mathcal{H}},\overline{\boldsymbol{s}}_{\|\mathcal{H}}\rangle\right] \overset{(a)}{=} -2\mathbb{E}\left[\overline{\boldsymbol{a}}^T H\overline{\boldsymbol{s}}\right] \overset{(b)}{=} -2(1-\epsilon)\boldsymbol{\nu}^T\mathbb{E}[H\overline{\boldsymbol{s}}] \overset{(c)}{=} -2(1-\epsilon)\boldsymbol{\nu}^T H\boldsymbol{\nu} + O(\epsilon), \tag{41}$$

where (a) holds because, for any vector $\boldsymbol{x}$, we have $\boldsymbol{x}_{\|\mathcal{H}} = H\boldsymbol{x}$ by the solution of the least squares problem and because H is a projection matrix; (b) holds because $\overline{\boldsymbol{a}}$ is independent of $\overline{\boldsymbol{s}}$ and $\mathbb{E}[\overline{\boldsymbol{a}}] = \boldsymbol{\lambda}^{(\epsilon)} = (1-\epsilon)\boldsymbol{\nu}$; and ($c$) holds because $H = C(C^TC)^{-1}C^T$ because $C^T\overline{\boldsymbol{s}}$ has elements $\langle \boldsymbol{c}^{(\ell)},\overline{\boldsymbol{s}}\rangle$ with $\ell\in\tilde{P}$ by Lemmas 2 and 3 and because $\boldsymbol{\nu}\in\mathcal{F}^{(\ell)}$. Therefore, using (39)–(41) in (38), we obtain

$$\left|\mathcal{T}_2 - \left(\boldsymbol{1}^T\left(H\circ\Sigma_a^{(\epsilon)}\right)\boldsymbol{1} + \boldsymbol{1}^T\left((C^TC)^{-1}\circ\Sigma_B\right)\boldsymbol{1} + \epsilon^2\boldsymbol{\nu}^T H\boldsymbol{\nu}\right)\right| \text{ is } O(\epsilon). \tag{42}$$

Now, we compute $\mathcal{T}_3$. We obtain

$$0 \le \mathcal{T}_3 = \mathbb{E}\left[\|\overline{\boldsymbol{u}}_{\|\mathcal{H}}\|^2\right] \overset{(a)}{\le} \sum_{\ell\in P}\mathbb{E}\left[\langle \boldsymbol{c}^{(\ell)},\overline{\boldsymbol{u}}\rangle^2\right] \overset{(b)}{\le} nS_{\max}C_{\max}\sum_{\ell\in P}\mathbb{E}\left[\langle \boldsymbol{c}^{(\ell)},\overline{\boldsymbol{u}}\rangle\right] \overset{(c)}{=} O(\epsilon),$$

where $C_{\max} = \max_{\ell\in P, i\in[n]}\{c_i^{(\ell)}\}$, and it is a finite constant. Here, (a) holds because the vectors $\boldsymbol{c}^{(\ell)}$ are not necessarily orthogonal for all $\ell\in P$; (b) holds because, by definition of the unused service, we know $\overline{\boldsymbol{u}} \le \overline{\boldsymbol{s}} \le S_{\max}\boldsymbol{1}$ with probability one; and (c) holds by Claim 2.

Claim 2. *Consider a set of generalized switches as described in Theorem 1. Then,*

$$\sum_{\ell\in P}\mathbb{E}\left[\langle \boldsymbol{c}^{(\ell)},\overline{\boldsymbol{u}}\rangle\right] = \epsilon\sum_{\ell\in P}b^{(\ell)} - O(\epsilon). \tag{43}$$

We present the proof of Claim 2 in Online Appendix D.4. Therefore,

$$\mathcal{T}_3 = O(\epsilon). \tag{44}$$

The last step is to prove that

$$|\mathcal{T}_4| = O(\sqrt{\epsilon}). \tag{45}$$

We provide the proof in Online Appendix D.5. Putting Equations (37), (42), (44), and (45) together, we obtain

$$\left|\mathbb{E}\left[\langle\overline{\boldsymbol{q}}^{(\epsilon)},\boldsymbol{w}\rangle\right] - \frac{1}{2\epsilon}\left(\boldsymbol{1}^T\left(H\circ\Sigma_a^{(\epsilon)}\right)\boldsymbol{1} + \boldsymbol{1}^T\left((C^TC)^{-1}\circ\Sigma_B\right)\boldsymbol{1}\right)\right| \le K(\epsilon).$$

This completes the proof. □

Clearly, this result and proof are much more general and more involved than the proof in the special case of an input-queued switch developed by Maguluri and Srikant [20] and Maguluri et al. [21]. The bound in Theorem 1 is expressed in terms of a general projection of the second moments of arrival and service processes onto the space $\mathcal{H}$. We point out a couple of conceptual differences from the proof in the case of the input-queued switch. First, in the proof of asymptotic upper bounds in an input-queued switch, the scheduling policy is not used. This means that, for an input-queued switch, any scheduling policy that exhibits SSC also has the same asymptotic upper bounds. In our proof here, we use the scheduling policy to upper bound the term $\mathcal{T}_1$ in Claim 1. Thus, we may not claim that any scheduling policy that exhibits SSC in Proposition 2 satisfies the bound in Theorem 1. Second, whereas SSC into the cone $\mathcal{K}$ is established by Maguluri and Srikant [20] and Maguluri et al. [21] in the case of an input-queued switch, only the weaker result about collapse into the space $\mathcal{H}$ is used to obtain heavy-traffic queue length bounds. In contrast, we use the collapse into the cone $\mathcal{K}$ in the proof of Theorem 1 to lower bound the term $\mathcal{T}_1$. Both these differences are because $\overline{s}_{\|\mathcal{H}}$ is constant for all maximal schedules $\overline{s}\in\mathcal{S}$ in the case of an input-queued switch, whereas in the case of the generalized switch, this is not necessarily true.

6.2. Proof of Theorem 2

For ease of exposition, in this proof, we use subscript $\|$ instead of $\|\mathcal{H}$ because we only use projection on the subspace $\mathcal{H}$ and not on the cone $\mathcal{K}$.

Proof of Theorem 2. We know that SSC occurs into a subspace of dimension $2N-1=3$. Therefore, three variables are necessary to compute the most general quadratic polynomial. In fact, we know $\overline{q}_{\|4}=\overline{q}_{\|2}+\overline{q}_{\|3}-\overline{q}_{\|1}$. Then, we only need to consider the variables $\overline{q}_{\|1}$, $\overline{q}_{\|2}$, and $\overline{q}_{\|3}$. The most general quadratic polynomial with these variables is

$$V(\boldsymbol{q})=\alpha_1\overline{q}_{\|1}^2+\alpha_2\overline{q}_{\|2}^2+\alpha_3\overline{q}_{\|3}^2+\alpha_4\overline{q}_{\|1}\overline{q}_{\|2}+\alpha_5\overline{q}_{\|1}\overline{q}_{\|3}+\alpha_6\overline{q}_{\|2}\overline{q}_{\|3},$$

where $\alpha_i\in\mathbb{R}$ for all $i\in[6]$.

Setting to zero the drift of $V(\boldsymbol{q})$ is equivalent to setting to zero the drift of each monomial separately. Then, we set to zero the drift of the following six test functions:

$$V_1(\boldsymbol{q})=\overline{q}_{\|1}^2,\ V_2(\boldsymbol{q})=\overline{q}_{\|2}^2,\ V_3(\boldsymbol{q})=\overline{q}_{\|3}^2,$$
$$V_4(\boldsymbol{q})=\overline{q}_{\|1}\overline{q}_{\|2},\ V_5(\boldsymbol{q})=\overline{q}_{\|1}\overline{q}_{\|3}\text{ and }V_6(\boldsymbol{q})=\overline{q}_{\|2}\overline{q}_{\|3}.$$

Before setting to zero the drift of $V_i(\boldsymbol{q})$ for $i\in[6]$ observe that, by definition of the cone $\mathcal{K}$ in (16), we have, for any vector $\boldsymbol{y}\in\mathbb{R}^4$,

$$y_{\|1}=\frac{y_1+y_2}{2}+\frac{y_1+y_3}{2}-\frac{y_1+y_2+y_3+y_4}{4}=\frac{3y_1+y_2+y_3-y_4}{4},\tag{46}$$
$$y_{\|2}=\frac{y_1+y_2}{2}+\frac{y_2+y_4}{2}-\frac{y_1+y_2+y_3+y_4}{4}=\frac{y_1+3y_2-y_3+y_4}{4},\tag{47}$$
$$y_{\|3}=\frac{y_3+y_4}{2}+\frac{y_1+y_3}{2}-\frac{y_1+y_2+y_3+y_4}{4}=\frac{y_1-y_2+3y_3+y_4}{4}.\tag{48}$$

Then, because the switch is completely saturated, we have

$$\mathbb{E}\left[\overline{a}_{\|i}\right]=\frac{1-\epsilon}{2}+\frac{1-\epsilon}{2}-\frac{2(1-\epsilon)}{4}=\frac{1-\epsilon}{2}\qquad\forall i\in[4],\tag{49}$$

and because $\overline{s}$ is a maximal schedule, we have

$$\overline{s}_{\|i}=\frac{1}{2}+\frac{1}{2}-\frac{2}{4}=\frac{1}{2}\qquad\forall i\in[4].\tag{50}$$

We first set to zero the drift of $V_1(\boldsymbol{q})$. We obtain

$$\begin{aligned}0&=\mathbb{E}\left[(\overline{q}_{\|1}^+)^2-\overline{q}_{\|1}^2\right]\\&=\mathbb{E}\left[(\overline{q}_{\|1}^+-\overline{u}_{\|1}+\overline{u}_{\|1})^2-\overline{q}_{\|1}^2\right]\\&=\mathbb{E}\left[(\overline{q}_{\|1}^+-\overline{u}_{\|1})^2+\overline{u}_{\|1}^2+2(\overline{q}_{\|1}^+-\overline{u}_{\|1})\overline{u}_{\|1}-\overline{q}_{\|1}^2\right]\\&\overset{(*)}{=}\mathbb{E}\left[(\overline{q}_{\|1}+\overline{a}_{\|1}-\overline{s}_{\|1})^2-\overline{u}_{\|1}^2+2\overline{q}_{\|1}^+\overline{u}_{\|1}-\overline{q}_{\|1}^2\right]\\&=\mathbb{E}\left[(\overline{a}_{\|1}-\overline{s}_{\|1})^2+2\overline{q}_{\|1}(\overline{a}_{\|1}-\overline{s}_{\|1})-\overline{u}_{\|1}^2+2\overline{q}_{\|1}^+\overline{u}_{\|1}\right],\end{aligned}\tag{51}$$

where $(*)$ holds by (2) and reorganizing the terms. We compute each term separately. For the first term, we have

$$\begin{aligned}\mathbb{E}\left[\left(\overline{a}_{\|1}-\overline{s}_{\|1}\right)^2\right] &\overset{(a)}{=} \mathbb{E}\left[\left(\overline{a}_{\|1}-\frac{1}{2}\right)^2\right]\\ &\overset{(b)}{=} \mathrm{Var}[\overline{a}_{\|1}]+(\mathbb{E}[\overline{a}_{\|1}])^2+\frac{1}{4}-\mathbb{E}[\overline{a}_{\|1}]\\ &= \mathrm{Var}[\overline{a}_{\|1}]+\left(\mathbb{E}[\overline{a}_{\|1}]-\frac{1}{2}\right)^2\\ &\overset{(c)}{=} \mathrm{Var}\left[\frac{3\overline{a}_1+\overline{a}_2+\overline{a}_3-\overline{a}_4}{4}\right]+\frac{\epsilon^2}{4}\\ &\overset{(d)}{=} \frac{9\left(\sigma_{a_1}^{(\epsilon)}\right)^2+\left(\sigma_{a_2}^{(\epsilon)}\right)^2+\left(\sigma_{a_3}^{(\epsilon)}\right)^2+\left(\sigma_{a_4}^{(\epsilon)}\right)^2}{16}+\frac{\epsilon^2}{4},\end{aligned}\tag{52}$$

where (a) holds by (50), (b) holds by definition of variance and reorganizing terms, (c) holds by definition of $\overline{a}_{\|1}$ as in (46) and by (49), and (d) holds because the arrival processes to different queues are independent. For the second term, we obtain

$$2\mathbb{E}\left[\overline{q}_{\|1}(\overline{a}_{\|1}-\overline{s}_{\|1})\right]\overset{(a)}{=}2\mathbb{E}\left[\overline{q}_{\|1}\left(\overline{a}_{\|1}-\frac{1}{2}\right)\right]\overset{(b)}{=}2\mathbb{E}\left[\overline{q}_{\|1}\right]\left(\mathbb{E}[\overline{a}_{\|1}]-\frac{1}{2}\right)\overset{(c)}{=}-\epsilon\mathbb{E}\left[\overline{q}_{\|1}\right],\tag{53}$$

where (a) holds by (50), (b) holds because the arrival processes are independent of the queue lengths, and (c) holds by (49). For the third term, observe

$$0\le\mathbb{E}\left[\overline{u}_{\|1}^2\right]\le\mathbb{E}\left[\|\overline{\boldsymbol{u}}_{\|}\|^2\right].$$

From the proof of Theorem 1, we know $\mathbb{E}[\|\overline{\boldsymbol{u}}_{\|}\|^2]$ is $O(\epsilon)$ (see (44)). Therefore,

$$\mathbb{E}\left[\overline{u}_{\|1}^2\right] \text{ is } O(\epsilon).\tag{54}$$

Now, we compute the last term. By definition of $\overline{\boldsymbol{q}}_{\|}$ and $\overline{\boldsymbol{q}}_{\perp}$, we have

$$2\mathbb{E}\left[\overline{q}_{\|1}^+\overline{u}_{\|1}\right]=2\mathbb{E}[\overline{q}_1^+\overline{u}_{\|1}]-2\mathbb{E}[\overline{q}_{\perp1}^+\overline{u}_{\|1}].$$

Claim 3. *Consider the queueing system described in Theorem 2. Then,*

$$\mathbb{E}[\overline{q}_{\perp1}^+\overline{u}_{\|1}] \text{ is } O(\sqrt{\epsilon}).$$

The proof of Claim 3 is presented in Online Appendix E.1. Then,

$$\begin{aligned}2\mathbb{E}\left[\overline{q}_{\|1}^+\overline{u}_{\|1}\right] &= 2\mathbb{E}[\overline{q}_1^+\overline{u}_{\|1}]+O(\sqrt{\epsilon})\\ &\overset{(a)}{=}\frac{1}{2}\mathbb{E}[\overline{q}_1^+(3\overline{u}_1+\overline{u}_2+\overline{u}_3-\overline{u}_4)]+O(\sqrt{\epsilon})\\ &\overset{(b)}{=}\frac{1}{2}\mathbb{E}[\overline{q}_1^+(\overline{u}_2+\overline{u}_3-\overline{u}_4)]+O(\sqrt{\epsilon}),\end{aligned}$$

where (a) holds by (46) and (b) holds by (3).

Claim 4. *Consider the queueing system described in Theorem 2. Then,*

$$\mathbb{E}[\overline{q}_1^+\overline{u}_4]=\mathbb{E}[\overline{q}_2^+\overline{u}_4]+\mathbb{E}[\overline{q}_3^+\overline{u}_4]+O(\sqrt{\epsilon}).$$

The proof of Claim 4 is presented in Online Appendix E.2. Therefore, we obtain

$$2\mathbb{E}\left[\overline{q}_{\|1}^+\overline{u}_{\|1}\right]=\frac{1}{2}\mathbb{E}[\overline{q}_1^+(\overline{u}_2+\overline{u}_3)]-\frac{1}{2}\mathbb{E}[\overline{q}_2^+\overline{u}_4]-\frac{1}{2}\mathbb{E}[\overline{q}_3^+\overline{u}_4]+O(\sqrt{\epsilon}).\tag{55}$$

Using (52)–(55) in (51) and reorganizing the terms, we obtain

$$\epsilon\mathbb{E}\left[\overline{q}_{\|1}\right]=\frac{9\left(\sigma_{a_1}^{(\epsilon)}\right)^2+\left(\sigma_{a_2}^{(\epsilon)}\right)^2+\left(\sigma_{a_3}^{(\epsilon)}\right)^2+\left(\sigma_{a_4}^{(\epsilon)}\right)^2}{16}+\frac{1}{2}\mathbb{E}\left[\overline{q}_1^+(\overline{u}_2+\overline{u}_3)\right]-\frac{1}{2}\mathbb{E}\left[\overline{q}_2^+\overline{u}_4\right]-\frac{1}{2}\mathbb{E}\left[\overline{q}_3^+\overline{u}_4\right]+\frac{\epsilon^2}{4}+O(\sqrt{\epsilon}).$$

Taking the limit as $\epsilon\downarrow 0$ on both sides, we obtain (17). The proof of (18) and of (19) hold similarly after setting to zero the drift of $V_2(\boldsymbol{q})$ and $V_3(\boldsymbol{q})$, respectively. We omit the details for brevity.

To obtain (20), we set to zero the drift of $V_4(\boldsymbol{q})$. After similar manipulation as earlier, we obtain

$$\begin{aligned}0&=\mathbb{E}\left[\overline{q}_{\|1}^+\overline{q}_{\|2}^+-\overline{q}_{\|1}\overline{q}_{\|2}\right]\\&=\mathbb{E}\left[\overline{q}_{\|1}(\overline{a}_{\|2}-\overline{s}_{\|2})\right]+\mathbb{E}\left[\overline{q}_{\|2}(\overline{a}_{\|1}-\overline{s}_{\|1})\right]+\mathbb{E}\left[(\overline{a}_{\|1}-\overline{s}_{\|1})(\overline{a}_{\|2}-\overline{s}_{\|2})\right]\\&\quad+\mathbb{E}\left[\overline{q}_{\|1}^+\overline{u}_{\|2}\right]+\mathbb{E}\left[\overline{q}_{\|2}^+\overline{u}_{\|1}\right]-\mathbb{E}\left[\overline{u}_{\|1}\overline{u}_{\|2}\right].\end{aligned}\tag{56}$$

We compute term by term. For the first term, we have

$$\mathbb{E}\left[\overline{q}_{\|1}(\overline{a}_{\|2}-\overline{s}_{\|2})\right]=-\frac{\epsilon}{2}\mathbb{E}\left[\overline{q}_{\|1}\right],\tag{57}$$

where we use that $\overline{s}_{\|2}=1/2$ and independence of the arrivals and queue lengths processes. Similarly, for the second term, we obtain

$$\mathbb{E}\left[\overline{q}_{\|2}(\overline{a}_{\|1}-\overline{s}_{\|1})\right]=-\frac{\epsilon}{2}\mathbb{E}\left[\overline{q}_{\|2}\right].\tag{58}$$

For the third term, we have

$$\begin{aligned}\mathbb{E}\left[(\overline{a}_{\|1}-\overline{s}_{\|1})(\overline{a}_{\|2}-\overline{s}_{\|2})\right]&\overset{(a)}{=}\mathbb{E}\left[\left(\overline{a}_{\|1}-\frac{1}{2}\right)\left(\overline{a}_{\|2}-\frac{1}{2}\right)\right]\\&\overset{(b)}{=}\mathrm{Cov}\left[\overline{a}_{\|1},\overline{a}_{\|2}\right]+\mathbb{E}\left[\overline{a}_{\|1}\right]\mathbb{E}\left[\overline{a}_{\|2}\right]-\frac{1}{2}\mathbb{E}\left[\overline{a}_{\|1}\right]-\frac{1}{2}\mathbb{E}\left[\overline{a}_{\|2}\right]+\frac{1}{4}\\&\overset{(c)}{=}\mathrm{Cov}\left[\frac{3\overline{a}_1+\overline{a}_2+\overline{a}_3-\overline{a}_4}{4},\frac{\overline{a}_1+3\overline{a}_2-\overline{a}_3+\overline{a}_4}{4}\right]+\frac{\epsilon^2}{4}\\&\overset{(d)}{=}\frac{3\left(\sigma_{a_1}^{(\epsilon)}\right)^2+3\left(\sigma_{a_2}^{(\epsilon)}\right)^2-\left(\sigma_{a_3}^{(\epsilon)}\right)^2-\left(\sigma_{a_4}^{(\epsilon)}\right)^2}{16}+\frac{\epsilon^2}{4},\end{aligned}\tag{59}$$

where (a) holds by (50); (b) holds by definition of covariance and reorganizing terms; (c) holds by (46), (47), and (49); and (d) holds because the arrival processes to different queues are independent. For the fourth term, we have

$$\begin{aligned}\mathbb{E}\left[\overline{q}_{\|1}^+\overline{u}_{\|2}\right]&\overset{(a)}{=}\mathbb{E}\left[\overline{q}_1^+\overline{u}_{\|2}\right]-\mathbb{E}\left[\overline{q}_{\perp 1}^+\overline{u}_{\|2}\right]\\&\overset{(b)}{=}\mathbb{E}\left[\overline{q}_1^+\overline{u}_{\|2}\right]+O(\sqrt{\epsilon})\\&\overset{(c)}{=}\frac{1}{4}\mathbb{E}\left[\overline{q}_1^+(\overline{u}_1+3\overline{u}_2-\overline{u}_3+\overline{u}_4)\right]+O(\sqrt{\epsilon})\\&\overset{(d)}{=}\frac{1}{4}\mathbb{E}\left[\overline{q}_1^+(3\overline{u}_2-\overline{u}_3+\overline{u}_4)\right]+O(\sqrt{\epsilon})\\&\overset{(e)}{=}\frac{1}{4}\mathbb{E}\left[\overline{q}_1^+(3\overline{u}_2-\overline{u}_3)\right]+\frac{1}{4}\mathbb{E}\left[\overline{q}_2^+\overline{u}_4\right]+\frac{1}{4}\mathbb{E}\left[\overline{q}_3^+\overline{u}_4\right]+O(\sqrt{\epsilon}),\end{aligned}\tag{60}$$

where (a) holds by definition of $\overline{\boldsymbol{q}}_{\|}$ and $\overline{\boldsymbol{q}}_{\perp}$, (b) holds similarly to Claim 3, (c) holds by (47), (d) holds by (3), and (e) holds by Claim 4. Similarly, for the fifth term, we have

$$\mathbb{E}\left[\overline{q}_{\|2}^+\overline{u}_{\|1}\right]=\frac{1}{4}\mathbb{E}\left[\overline{q}_2^+(3\overline{u}_1+\overline{u}_3-\overline{u}_4)\right]+O(\sqrt{\epsilon}).\tag{61}$$

For the sixth term, we have

$$\begin{aligned} 0 \le \mathbb{E}\left[\overline{u}_{\|1}\overline{u}_{\|2}\right] &\overset{(*)}{\le} \sqrt{\mathbb{E}\left[\overline{u}_{\|1}^2\right]\mathbb{E}\left[\overline{u}_{\|2}^2\right]} \\ &\le \sqrt{\mathbb{E}\left[\|\overline{\boldsymbol{u}}_{\|}\|^2\right]\mathbb{E}\left[\|\overline{\boldsymbol{u}}_{\|}\|^2\right]} \\ &= \mathbb{E}\left[\|\overline{\boldsymbol{u}}_{\|}\|^2\right], \end{aligned}$$

where $(*)$ holds by the Cauchy–Schwarz inequality. Also, because $\mathbb{E}[\|\overline{\boldsymbol{u}}_{\|}\|^2]$ is $O(\epsilon)$, we obtain

$$\mathbb{E}\left[\overline{u}_{\|1}\overline{u}_{\|2}\right] \text{ is } O(\epsilon). \tag{62}$$

Using (57)–(62) in (56), and reorganizing terms, we obtain

$$\begin{aligned} &\epsilon\mathbb{E}\left[\overline{q}_{\|1}\right] + \epsilon\mathbb{E}\left[\overline{q}_{\|2}\right] \\ &= \frac{3\left(\sigma_{a_1}^{(\epsilon)}\right)^2 + 3\left(\sigma_{a_2}^{(\epsilon)}\right)^2 - \left(\sigma_{a_3}^{(\epsilon)}\right)^2 - \left(\sigma_{a_4}^{(\epsilon)}\right)^2}{8} + \frac{\epsilon^2}{2} + O(\sqrt{\epsilon}) \\ &\quad + \frac{1}{2}\mathbb{E}\left[\overline{q}_1^+(3\overline{u}_2 - \overline{u}_3)\right] + \frac{1}{2}\mathbb{E}\left[\overline{q}_3^+\overline{u}_4\right] + \frac{1}{2}\mathbb{E}\left[\overline{q}_2^+(3\overline{u}_1 + \overline{u}_3)\right]. \end{aligned}$$

Taking the limit as $\epsilon \downarrow 0$ on both sides, we obtain (20). The proof of (21) and (22) hold similarly after setting to zero the drift of $V_5(\boldsymbol{q})$ and $V_6(\boldsymbol{q})$, respectively. This completes the proof of Theorem 2. □

7. Conclusion

In this paper, we study one of the most general single-hop SPNs with control in service: the generalized switch. This model subsumes several queueing systems, such as the input-queued switch, parallel-server systems, ad hoc wireless networks, etc. Our result is widely applicable because we do not assume the CRP condition, neither independence of the arrival processes.

We showcase the generality of our result with three particular SPNs: the input-queued switch, parallel-server systems, and an ad hoc wireless network. Each of these results are interesting by themselves because they are studied separately in the literature, and we can easily compute them as applications of Theorem 1.

Additionally, we prove that, if the heavy-traffic limit is to a vertex of the capacity region, then SSC does not result in a reduction on the dimension of the state space. In other words, in this case, we observe full-dimensional SSC. Under this condition, regardless of the correlation among arrival processes, the mean of the linear combinations of the queue lengths that we obtain behave as if the queues were independent in heavy traffic.

Our result is widely applicable to several SPNs, but it only allows us to compute certain linear combinations of the queue lengths. In the case of an input-queued switch, this linear combination turns out to be the total queue length, and in parallel-server systems, the weights of the linear combination are the mean service rates.

We also show that obtaining other linear combinations is a nontrivial problem because using the drift method with polynomial test functions is equivalent to solving an underdetermined system of linear equations. The results we obtain in this paper can be also obtained by taking specific linear combinations of these equations, such that some unknowns cancel out. An immediate line of future work is to extend the method so that all the linear combinations can be computed. This would allow us to also obtain higher moments and, eventually, the joint distribution of the queue lengths.

Acknowledgment

The authors acknowledge professor Mohit Singh for the recommendation of using the least squares problem to prove Theorem 1.

Endnote

[1] The code is publicly available at https://github.com/dhurtadolange/2x2-switch-simulation.

References

[1] Bell SL, Williams RJ (2001) Dynamic scheduling of a system with two parallel servers in heavy traffic with resource pooling: Asymptotic optimality of a threshold policy. *Ann. Appl. Probab.* 11(3):608–649.

[2] Benson T, Akella A, Maltz DA (2010) Network traffic characteristics of data centers in the wild. *Proc. 10th ACM SIGCOMM Conf. Internet Measurement* (Association of Computer Machinery, New York), 267–280.
[3] Bertsimas D, Tsitsiklis JN (1997) *Introduction to Linear Optimization*, vol. 6 (Athena Scientific, Belmont, MA).
[4] Bertsimas D, Paschalidis IC, Tsitsiklis JN (1994) Optimization of multiclass queueing networks: Polyhedral and nonlinear characterizations of achievable performance. *Ann. Appl. Probab.* 4(1):43–75.
[5] Dai J, Lin W (2008) Asymptotic optimality of maximum pressure policies in stochastic processing networks. *Ann. Appl. Probab.* 18(6): 2239–2299.
[6] Dimakis A, Walrand J (2006) Sufficient conditions for stability of longest queue first scheduling. *Adv. Appl. Prob.* 38(2):505–521.
[7] Eryilmaz A, Srikant R (2012) Asymptotically tight steady-state queue length bounds implied by drift conditions. *Queueing Systems* 72(3–4): 311–359.
[8] Garnett O, Mandelbaum A (2000) An introduction to skills-based routing and its operational complexities. Teaching notes.
[9] Ghamami S, Ward AR (2013) Dynamic scheduling of a two-server parallel server system with complete resource pooling and reneging in heavy traffic: Asymptotic optimality of a two-threshold policy. *Math. Oper. Res.* 38(4):761–824.
[10] Gupta G, Shroff N (2010) Delay analysis for wireless networks with single hop traffic and general interference constraints. *IEEE/ACM Trans. Networking* 18(2):393–405.
[11] Hajek B (1982) Hitting-time and occupation-time bounds implied by drift analysis with applications. *Adv. Appl. Probab.* 14(3):502–525.
[12] Harrison J (1998) Heavy traffic analysis of a system with parallel servers: Asymptotic optimality of discrete review policies. *Ann. App. Probab.* 8(3):822–848.
[13] Harrison J (2013) *Brownian Models of Performance and Control* (Cambridge University Press, Cambridge, UK).
[14] Harrison J, López M (1999) Heavy traffic resource pooling in parallel-server systems. *Queueing Systems* 33:339–368.
[15] Kandula S, Sengupta S, Greenberg A, Patel P, Chaiken R (2009) The nature of data center traffic: Measurements & analysis. *Proc. Ninth ACM SIGCOMM Conf. Internet Measurement* (Association of Computer Machinery, New York), 202–208.
[16] Kang W, Williams R (2012) Diffusion approximation for an input-queued switch operating under a maximum weight matching policy. *Stochastic Systems* 2(2):277–321.
[17] Kang X, Wang W, Jaramillo JJ, Ying L (2014) On the performance of largest-deficit-first for scheduling real-time traffic in wireless networks. *IEEE/ACM Trans. Networking* 24(1):72–84.
[18] Kingman J (1962) Some inequalities for the queue GI/G/1. *Biometrika* 49(3/4):315–324.
[19] Kumar S, Kumar PR (1994) Performance bounds for queueing networks and scheduling policies. *IEEE Trans. Automatic Control* 39(8): 1600–1611.
[20] Maguluri ST, Srikant R (2016) Heavy traffic queue length behavior in a switch under the MaxWeight algorithm. *Stochastic Systems* 6(1): 211–250.
[21] Maguluri ST, Burle S, Srikant R (2018) Optimal heavy-traffic queue length scaling in an incompletely saturated switch. *Queueing Systems* 88(3–4):279–309.
[22] Maguluri ST, Hajek B, Srikant R (2014) The stability of longest-queue-first scheduling with variable packet sizes. *IEEE Transactions on Automatic Control* 59(8):2295–2300.
[23] Meyn S (2009) Stability and asymptotic optimality of generalized maxweight policies. *SIAM J. Control Optim.* 47(6):3259–3294.
[24] Shah D, Tsitsiklis J, Zhong Y (2011) Optimal scaling of average queue sizes in an input-queued switch: An open problem. *Queueing Systems* 68(3–4):375–384.
[25] Shi C, Wei Y, Zhong Y (2019) Process flexibility for multiperiod production systems. *Oper. Res.* 67(5):1300–1320.
[26] Stolyar A (2004) MaxWeight scheduling in a generalized switch: State space collapse and workload minimization in heavy traffic. *Ann. Appl. Probab.* 14(1):1–53.
[27] Tassiulas L, Ephremides A (1992) Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Trans. Automatic Control* 37(12):1936–1948.
[28] Wang W, Maguluri S, Srikant R, Ying L (2022) Heavy-traffic insensitive bounds for weighted proportionally fair bandwidth sharing policies. *Math. Oper. Res.*, ePub ahead of print February 7, https://doi.org/10.1287/moor.2021.1225.
[29] Williams R (2000) On dynamic scheduling of a parallel server system with complete resource pooling. McDonald DR, Turner SRE, eds. *Analysis of Communication Networks: Call Centres, Traffic and Performance*, Fields Institute Communications, vol. 28 (American Mathematical Society), 49–71.
[30] Williams R (2016) Stochastic processing networks. *Annual. Rev. Statist. Appl.* 3:323–345.