

Optimal heavy-traffic queue length scaling in an incompletely saturated switch

Siva Theja Maguluri¹  · Sai Kiran Burle² · R. Srikant²

Received: 24 October 2016 / Revised: 2 October 2017 / Published online: 28 November 2017
© Springer Science+Business Media, LLC, part of Springer Nature 2017

Abstract We consider an input-queued switch operating under the MaxWeight scheduling algorithm. This system is interesting to study because it is a model for Internet routers and data center networks. Recently, it was shown that the MaxWeight algorithm has optimal heavy-traffic queue length scaling when all ports are uniformly saturated. Here we consider the case when an arbitrary number of ports are saturated (which we call the incompletely saturated case), and each port is allowed to saturate at a different rate. We use a recently developed drift technique to show that the heavy-traffic queue length under the MaxWeight scheduling algorithm has optimal scaling with respect to the switch size even in these cases.

Keywords $n \times n$ switch · Heavy-traffic optimality · Drift method · Performance analysis

Mathematics Subject Classification 60K25 · 90B15

✉ Siva Theja Maguluri
siva.theja@gmail.com

Sai Kiran Burle
burle2@illinois.edu

R. Srikant
rsrikant@illinois.edu

¹ H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA, USA

² Department of ECE and CSL, University of Illinois at Urbana-Champaign, Urbana, IL, USA

1 Introduction

The $n \times n$ switch is a model that has been widely used and studied to understand the behavior of Internet routers and data center switches. The importance of this model in the design of Internet routers is well-known. In the mid to late 1990s when the Internet was exploding, it served as an important model to study and design scheduling algorithms for the switch fabric of Internet routers. The model is now used to understand the design of data centers used for cloud-computing services. Today's data centers consist of a massive number of servers organized in racks, which are interconnected through a data center network. An ideal data center network is a huge input-queued switch with one port for each server. However, real switches are much smaller and they have to be interconnected appropriately, and routing and scheduling algorithms have to be designed, so that the overall network emulates an $n \times n$ switch. Designing such a network is a challenging and active area of research; see [1, 14] for example. Here, we do not explicitly consider a data center network, but only consider the $n \times n$ switch which is the underlying model (see [14], which argues why the model is appropriate even for a data center network) and study the behavior of the well-known MaxWeight algorithm [22] for this model.

As mentioned in [15], the $n \times n$ switch model also serves as a canonical theoretical example of a problem which exhibits the so-called multi-dimensional state space collapse, which makes it difficult to study using traditional heavy-traffic theory. Recently, it has been shown in [11] that the heavy-traffic behavior of the mean queue length in an $n \times n$ switch operating under the MaxWeight scheduling algorithm can be precisely characterized using a nontrivial extension of a drift technique introduced in [6]. In particular, one of the key contributions of [11] is to extend the drift technique to cover the case of multi-dimensional state space collapse. The result in [11] also resolved an open question on the scaling behavior of the heavy-traffic queue length in a switch operating under the MaxWeight algorithm. In particular, it showed that the total heavy-traffic-scaled queue length is $O(n)$ or the mean heavy-traffic-scaled delay experienced by a packet is $O(1)$. While there have been other results establishing $O(1)$ delay scaling, the significance of the result in [11] is that the result holds for the original MaxWeight algorithm introduced in [22], with no additional scheduling operations required.

The results in [11] were obtained under the assumption that every input and output port of the switch is saturated (i.e., close to capacity), and the arrival rates to each input port and each output port are close to capacity by the same amount. For the purposes of this paper, we call a switch where only some of ports are saturated an “incompletely saturated switch.” The main purpose of this paper is to show that the MaxWeight algorithm has order-optimal scaling (in the number of ports of the switch) for the case of an incompletely saturated switch and for the case where each port has a different rate of saturation. The results in [11] were obtained by setting the drift of a test function to zero in steady state. This function was carefully chosen based on the underlying symmetry when all ports are saturated. Since we do not have such a symmetry in the incompletely saturated switch, a new test function has to be used. The main reason for this is that the geometry of the state space collapse is different here than in the all-port saturated case considered in [11]. In this paper, we propose a novel test function to use

for the incompletely saturated switch. Another major contribution of the paper is that we express this function in a simple form, so that it can be generalized for use in other problems that exhibit state space collapse into other regions. The case where each port has a different rate of saturation is of theoretical importance since it corresponds to the situation where the drift vector is not the identity matrix in the diffusion limit [8]. In fact, the diffusion limit is no longer symmetric in the components of the limiting stochastic process (i.e., the diffusion limit is not symmetric across the ports), but here we show that the technique in [11] works even in this asymmetric situation to produce an exact formula for a certain linear combination of the queue lengths. This result can be further used to show optimal queue length scaling under some conditions on the saturation rates.

The rest of the paper is organized as follows. The model of an $n \times n$ switch and the MaxWeight algorithm are presented in Sect. 2. General results on queue lengths are presented in Sect. 3. In Sect. 4, we present various extensions and special cases in order to interpret the general result. Concluding remarks are provided in the last section.

We remark that a very preliminary version of this paper appeared in [10].

1.1 Related work

The MaxWeight algorithm for general stochastic networks, of which the $n \times n$ switch is a special case, was presented in [22], where it was shown that algorithm is throughput-optimal. The special case of the switch was considered in [12], where it was shown that simpler algorithms such as MaxSize and maximal matchings are not throughput-optimal. The case of nonstochastic arrivals was considered in [24], where in addition to the throughput optimality of MaxWeight-type algorithms, a lower bound on the throughput loss of simpler algorithms such as maximal matching was established.

Here we are interested in performance metrics beyond throughput optimality. In particular, we are interested in understanding whether the MaxWeight matching algorithm for switches achieves small queue lengths, at least under a heavy-traffic scaling regime. Using diffusion limits, the heavy-traffic optimality of the MaxWeight algorithm in a switch where only one port is saturated was established in [21], although the final step of interchanging the order of the heavy-traffic scaling limit and letting time go to infinity was not undertaken there. Motivated by this result, [6] studied the switch directly in steady state, established heavy-traffic optimality, and introduced a new drift method of studying stochastic networks in heavy-traffic. However, it should be emphasized that the results in [6,21] apply only to the case of a single saturated port since they both rely on state of the system collapsing to a single dimension in the heavy-traffic limit. In a recent development, the case where all ports are uniformly saturated (thus leading to the more difficult case of multi-dimensional state space collapse) was studied in [11], where an exact expression for the heavy-traffic scaled queue length under the MaxWeight algorithm is derived. Additionally, this expression shows that the algorithm has heavy-traffic optimal scaling in the size of the switch, resolving an open conjecture stated in [15]. The results in [11] use and significantly extend the drift technique presented in [6].

State collapse in the case where multiple ports are saturated has been established in [2, 18, 19] and, using the state space collapse result in [18], a diffusion limit was

established in [8]. However, properties of the diffusion limit (such as its steady state distribution or mean queue lengths) were difficult to establish. The result in [11] can thus be interpreted as a derivation of the sum of the first moments of the limiting vector stochastic process, but obtained without going through the usual fluid/diffusion limit scaling arguments. An entirely different technique to study heavy-traffic optimality was presented in [17], where the authors approximate the scheduling decisions made by a switch which can change its schedule infinitely often to simulate a queueing network with product-form steady state distribution as in [4]. The resulting algorithm is heavy-traffic optimal, but has a very high computational complexity. The optimal scaling of the queue length as a function of the switch size in the nonheavy-traffic limit appears to be still open. Alternatively, one can consider asymptotic regimes other than the heavy-traffic limit. The best known results in this regard are the ones in [13, 16], but these require algorithms that are more involved than the original MaxWeight algorithm.

2 System model and background

In this section, we present the model of an input-queued switch, the MaxWeight scheduling algorithm, and some lemmas that will later play a key role in the results.

Note on notation For ease of understanding, and to allow the reader to compare and contrast with the results in [11], here we use definitions and notation consistent with [11]. Since an $n \times n$ switch has n^2 queues, we often deal with the Euclidean space \mathbb{R}^{n^2} . To describe the elements of the space \mathbb{R}^{n^2} , we will use the terms vector and matrix interchangeably, because these elements may be viewed either as n^2 -dimensional vectors or as $n \times n$ matrices. Both these views will be used in this paper for ease of exposition. We primarily use the inner product and norm of the vector space \mathbb{R}^{n^2} , but we represent the vectors in \mathbb{R}^{n^2} as $n \times n$ matrices for convenience. Thus, \mathbf{x} is a matrix with the (i, j) th component denoted by x_{ij} , and for two vectors \mathbf{x} and \mathbf{y} in \mathbb{R}^{n^2} , their inner product $\langle \mathbf{x}, \mathbf{y} \rangle$ and Euclidean norm $\|\mathbf{x}\|$ are defined by

$$\langle \mathbf{x}, \mathbf{y} \rangle \triangleq \sum_{i=1}^n \sum_{j=1}^n x_{ij} y_{ij}, \quad \|\mathbf{x}\| \triangleq \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle} = \sqrt{\sum_{i=1}^n \sum_{j=1}^n x_{ij}^2}.$$

For two vectors \mathbf{x} and \mathbf{y} in \mathbb{R}^{n^2} , $\mathbf{x} \leq \mathbf{y}$ means $x_{ij} \leq y_{ij}$ for every (i, j) . We use $\mathbf{1}$ to denote the all ones vector. Let $\mathbf{e}^{(i)}$ denote the vector defined by $e_{ij}^{(i)} = 1$ for all j and $e_{i'j}^{(i)} = 0$ for all $i' \neq i$ and for all j . Thus, $\mathbf{e}^{(i)}$ is a matrix whose i th row is all ones, with zeros everywhere else. Similarly, let $\tilde{\mathbf{e}}^{(j)}$ denote the vector defined by $\tilde{e}_{ij}^{(j)} = 1$ for all i and $\tilde{e}_{i'j'}^{(j)} = 0$ for all $j' \neq j$ and for all i , i.e., it is a matrix whose j th column is all ones, with zeros everywhere else. We use $\sum_i(\cdot)$ without the limits on summation to denote $\sum_{i=1}^n(\cdot)$. For a random process $\mathbf{q}(t)$ and a function $V(\cdot)$, we will sometimes use $V(t)$ to denote $V(\mathbf{q}(t))$. We use $\text{Var}(\cdot)$ to denote variance of a random variable and $\text{Cov}(\cdot)$ to denote covariance between two random variables.

2.1 The switch model and the MaxWeight algorithm

For the purposes of queueing-theoretical analysis, an $n \times n$ switch can be treated as an $n \times n$ matrix of queues operating in a time-slotted discrete-time fashion. Let $a_{ij}(t)$ denote the number of packet arrivals to the (i, j) th queue, i.e., the queue in the i th row and j th column. We let $\mathbf{a} \in \mathbb{R}^{n^2}$ denote the vector $(a_{ij})_{ij}$. For every queue (i, j) , the arrival process $a_{ij}(t)$ is a stochastic process that is i.i.d. across time, with mean $E[a_{ij}(t)] = \lambda_{ij}$ and variance $\text{Var}(a_{ij}(t)) = \sigma_{ij}^2$ for any time t . We assume that the arrival processes are independent across queues (the processes $a_{ij}(t)$ and $a_{i'j'}(t)$ are independent for $(i, j) \neq (i', j')$) and are also independent of the queue lengths or schedules chosen in the switch. We further assume that for all i, j, t , $a_{ij}(t) \leq a_{\max}$ for some $a_{\max} \geq 1$. The arrival rate vector is denoted by $\boldsymbol{\lambda} = (\lambda_{ij})_{ij}$ and the variance vector $(\sigma_{ij}^2)_{ij}$ is denoted by $(\boldsymbol{\sigma})^2$ or $\boldsymbol{\sigma}^2$. We will use $\boldsymbol{\sigma}$ to denote $(\sigma_{ij})_{ij}$. We denote the queue length of packets at input port i to be delivered at output port j at time t by $q_{ij}(t)$. Let $\mathbf{q} \in \mathbb{R}^{n^2}$ denote the vector of all queue lengths.

The key scheduling constraints are that (i) at most one packet can be removed from each queue in each time slot and (ii) at most one queue can be served in each row and each column in each time slot. These constraints arise from technological constraints in a real switch, where each row represents an input port and each column represents an output port; see [20] for example. The scheduling constraints can be captured in graph-theoretical language as follows: Let G denote a complete $n \times n$ bipartite graph with n^2 edges. Each node on the left side of the bipartite graph can be thought of as representing a row in the matrix of switches and each node on the right side represents a column. The schedule in each time slot is a matching on this graph G . Let s_{ij} be the amount of service provided to queue (i, j) in a given time slot. Thus, $s_{ij} = 1$ if the link between input port i and output port j is matched or scheduled, and $s_{ij} = 0$ otherwise, and we denote $\mathbf{s} = (s_{ij})_{ij}$. Then, the set of feasible schedules, $\mathcal{S} \subset \{0, 1\}^{n^2}$, is the set of all vectors \mathbf{s} which satisfy

$$\sum_{i=1}^n s_{ij} \leq 1, \sum_{j=1}^n s_{ij} \leq 1 \forall i, j \in \{1, 2, \dots, n\}.$$

Let \mathcal{S}^* denote the set of maximal feasible schedules. Then, it is easy to see that \mathcal{S}^* is the set of all vectors \mathbf{s} which satisfy

$$\sum_{i=1}^n s_{ij} = 1, \sum_{j=1}^n s_{ij} = 1 \forall i, j \in \{1, 2, \dots, n\}. \tag{1}$$

A scheduling policy or algorithm picks a schedule $\mathbf{s}(t)$ in every time slot based on the current queue length vector, $\mathbf{q}(t)$. In each time slot, the order of events is as follows: Queue lengths at the beginning of time slot t are $\mathbf{q}(t)$. A schedule $\mathbf{s}(t)$ is then picked for that time slot based on the queue lengths. Then, arrivals for that time $\mathbf{a}(t)$ happen. Finally, the packets are served, and there is unused service if there are no packets in a scheduled queue. The queue lengths are then updated to give the queue lengths for the next time slot. The queue lengths therefore evolve as follows:

$$\begin{aligned}
 q_{ij}(t + 1) &= [q_{ij}(t) + a_{ij}(t) - s_{ij}(t)]^+ \\
 &= q_{ij}(t) + a_{ij}(t) - s_{ij}(t) + u_{ij}(t), \\
 \mathbf{q}(t + 1) &= \mathbf{q}(t) + \mathbf{a}(t) - \mathbf{s}(t) + \mathbf{u}(t),
 \end{aligned}
 \tag{2}$$

where $[x]^+ = \max(0, x)$ is the projection onto the positive real axis, and $u_{ij}(t)$ is the unused service on link (i, j) . Unused service is 1 only when link (i, j) is scheduled, but has zero queue length; and it is 0 in all other cases. Thus, when $u_{ij}(t) = 1$, we have $q_{ij}(t) = 0, a_{ij}(t) = 0, s_{ij}(t) = 1$ and $q_{ij}(t + 1) = 0$. Therefore, we have $u_{ij}(t)q_{ij}(t) = 0, u_{ij}(t)a_{ij}(t) = 0$ and $u_{ij}(t)q_{ij}(t + 1) = 0$. Also note that since $u_{ij}(t) \leq s_{ij}(t)$, we have that $\sum_{i=1}^n u_{ij} \in \{0, 1\}$ and $\sum_{j=1}^n u_{ij} \in \{0, 1\}$ for all i, j .

The MaxWeight algorithm is a popular scheduling algorithm for switches. In every time slot t , each link (i, j) is given a weight equal to its queue length $q_{ij}(t)$ and the schedule with the maximum weight among the feasible schedules \mathcal{S} is chosen at that time slot. In other words, using queue lengths as the weights, the permutation matrix with the maximum weight is picked in every time slot. This algorithm is presented in Algorithm 1.

Algorithm 1 MaxWeight scheduling algorithm for an input-queued switch

Consider the complete bipartite graph described earlier. Let the queue length $q_{ij}(t)$ be the weight of the edge between input port i and output port j . A maximum weight matching in this graph is chosen as the schedule in every time slot, i.e.,

$$\mathbf{s}(t) = \arg \max_{\mathbf{s} \in \mathcal{S}} \sum_{ij} q_{ij}(t)s_{ij} = \arg \max_{\mathbf{s} \in \mathcal{S}} \langle \mathbf{q}(t), \mathbf{s} \rangle.
 \tag{3}$$

Ties are broken uniformly at random.

Note that there is always a maximum weight schedule that is maximal. If the MaxWeight schedule chosen at time t , \mathbf{s} , is not maximal, there exists a maximal schedule $\mathbf{s}^* \in \mathcal{S}^*$ such that $\mathbf{s} \leq \mathbf{s}^*$. For any link (i, j) such that $s_{ij} = 0$ and $s_{ij}^* = 1, q_{ij}(t) = 0$. If not, \mathbf{s} would not have been a maximum weight schedule. Therefore, we can pretend that the actual schedule chosen is \mathbf{s}^* and the links (i, j) that are in \mathbf{s}^* but not in \mathbf{s} have an unused service of 1. Note that this does not change the scheduling algorithm, but it is just a notational convenience. Therefore, without loss of generality, we assume that the schedule chosen in each time slot is a maximal schedule, i.e.,

$$\mathbf{s}(t) \in \mathcal{S}^* \text{ for all time } t.$$

Hence, the MaxWeight schedule picks one of the $n!$ possible permutations from the set \mathcal{S}^* in each time slot.

Under i.i.d. arrivals, the queue lengths process $\mathbf{q}(t)$ is a Markov chain. The switch is said to be stable under a scheduling policy if the sum of all the queue lengths is finite in an appropriate stochastic sense (see [20] for example). The capacity region of the switch is the set of arrival rates λ for which the switch is stable under some scheduling policy. A policy that stabilizes the switch under any arrival rate in the capacity region

is said to be throughput-optimal. It is well-known [12,22] that the capacity region \mathcal{C} of the switch is convex hull of all feasible schedules:

$$\begin{aligned} \mathcal{C} &= \text{Conv}(\mathcal{S}) \\ &= \left\{ \boldsymbol{\lambda} \in \mathbb{R}_+^{n^2} : \sum_{i=1}^n \lambda_{ij} \leq 1, \sum_{j=1}^n \lambda_{ij} \leq 1 \forall i, j \in \{1, \dots, n\} \right\} \\ &= \left\{ \boldsymbol{\lambda} \in \mathbb{R}_+^{n^2} : \langle \boldsymbol{\lambda}, \mathbf{e}^{(i)} \rangle \leq 1, \langle \boldsymbol{\lambda}, \tilde{\mathbf{e}}^{(j)} \rangle \leq 1 \forall i, j \in \{1, \dots, n\} \right\}. \end{aligned} \tag{4}$$

For any arrival rate vector $\boldsymbol{\lambda}$, $\rho \triangleq \max_{ij} \{ \sum_i \lambda_{ij}, \sum_j \lambda_{ij} \}$ is called the load. It is also known that the queue lengths process is positive recurrent under the MaxWeight algorithm whenever the arrival rate is in the capacity region \mathcal{C} (equivalently, load $\rho < 1$) and therefore is throughput-optimal [12,22], [20, Chap 4].

For any arrival rate in the capacity region \mathcal{C} , due to positive recurrence of $\mathbf{q}(t)$, we have that a steady state distribution exists under the MaxWeight policy. Let $\bar{\mathbf{q}}$ denote the steady state random vector. In this paper, we focus on the weighted average queue length under the steady state distribution, i.e., $\mathbb{E}[\sum_{i,j} \alpha_{ij} \bar{q}_{ij}]$, for some weights α_{ij} , which can be shown to exist as in [11]. We consider a set of switch systems indexed by a parameter ϵ , with arrival rate $\boldsymbol{\lambda}^\epsilon$ so that the arrival rate approaches the vector $\boldsymbol{\nu}$ on the boundary of the capacity region \mathcal{C} in the limit as $\epsilon \rightarrow 0$. This is called the heavy-traffic limit. We are interested in the weighted average queue length in the heavy-traffic limit, i.e., $\lim_{\epsilon \rightarrow 0} \mathbb{E}[\sum_{i,j} \alpha_{ij} \bar{q}_{ij}]$. In particular, in this paper, we will consider cases where the sum of the arrival rates at some rows and some columns approach 1, and they may approach 1 at different rates at each column and row.

2.2 Kingman bound for a discrete-time queue

To establish our results, we later show that the total queue length along each row and each column is lower-bounded. For this purpose, we use a bound on the steady state queue length in a simple discrete-time queue [6]. While the well-known Kingman bound is for continuous-time G/G/1 queues, due to the similarity in establishing the result, the bound for the discrete-time case is also called the Kingman bound in [6] and we use the same terminology here. We state the version of the result for the special case where a queue can serve only one packet per slot here since this is what is used in this paper. In this special case, instead of a bound, one has an exact expression for the mean queue length which we present below.

Lemma 1 *Consider a single server operating in discrete time. In each time slot, packets arrive according to an i.i.d. arrival process $\alpha(t)$ with mean λ and variance σ^2 . Let q denote the queue length. Each packet needs exactly one time slot of service. The server operates according to any nonidling policy, serving one packet in every time slot whenever the queue is nonempty. Then, the queue is positive recurrent as long as $\lambda < 1$, and the steady state mean queue length is given by*

$$\mathbb{E}[q] = \frac{\sigma^2}{2(1 - \lambda)} - \frac{\lambda}{2}.$$

We note that the first term on the right-hand side of the above equation is what is referred to as the Kingman bound in [6].

2.3 Moment bounds from Lyapunov drift conditions

In later sections of the paper, we establish state space collapse results by obtaining moment bounds on certain quantities related to the queue length vector based on drift of a Lyapunov function. A key ingredient in this approach is to obtain moment bounds from drift conditions. A lemma from [7] was used in [6] to obtain these bounds, and a different result from [3] was used in [11] to obtain tighter bounds. Here we state [3, Theorem 1] in the form it was stated in [11].

Lemma 2 *For an irreducible and aperiodic Markov chain $\{X(t)\}_{t \geq 0}$ over a countable state space \mathcal{X} , suppose $Z : \mathcal{X} \rightarrow \mathbb{R}_+$ is a nonnegative-valued Lyapunov function. We define the drift of Z at X as*

$$\Delta Z(X) \triangleq [Z(X(t + 1)) - Z(X(t))] \mathcal{I}(X(t) = X),$$

where $\mathcal{I}(\cdot)$ is the indicator function. Thus, $\Delta Z(X)$ is a random variable that measures the amount of change in the value of Z in one step, starting from state X . This drift is assumed to satisfy the following conditions:

C1 *There exists an $\eta > 0$, and a $\zeta < \infty$ such that, for any $t = 1, 2, \dots$ and for all $X \in \mathcal{X}$ with $Z(X) \geq \zeta$,*

$$\mathbb{E}[\Delta Z(X) | X(t) = X] \leq -\eta.$$

C2 *There exists a $D < \infty$ such that, for all $X \in \mathcal{X}$,*

$$\mathbb{P}(|\Delta Z(X)| \leq D) = 1.$$

Further assume that the Markov chain $\{X(t)\}_t$ converges in distribution to a random variable \bar{X} . Then, for any $r = 1, 2, \dots$,

$$\mathbb{E}[Z(\bar{X})^r] \leq (2\zeta)^r + (4D)^r \left(\frac{D+\eta}{\eta}\right)^r r!.$$

3 Incompletely saturated switch

In this section, we will study the switch system when an arbitrary number of ports are saturated. We consider the switch where $n_1 \leq n$ input ports (rows) and $n_2 \leq n$ output ports (columns) are saturated. Without loss of generality, we assume that input ports (rows) $1, 2, \dots, n_1$ and output ports (columns) $1, 2, \dots, n_2$ are saturated. Thus, we

consider a point \mathbf{v} on the boundary of the capacity region that lies in $\text{Relint}(\mathcal{F}_{n_1 n_2})$, the relative interior¹ of the face $\mathcal{F}_{n_1 n_2}$ defined by

$$\begin{aligned} \mathcal{F}_{n_1 n_2} &\triangleq \left\{ \mathbf{v} \in \mathcal{C} : \sum_{j=1}^n v_{ij} = 1 \ \forall i \in \{1, \dots, n_1\}, \sum_{i=1}^n v_{ij} = 1 \ \forall j \in \{1, \dots, n_2\} \right\} \\ &= \left\{ \mathbf{v} \in \mathcal{C} : \langle \mathbf{v}, \mathbf{e}^{(i)} \rangle = 1 \ \forall i \in \{1, \dots, n_1\}, \right. \\ &\quad \left. \langle \boldsymbol{\lambda}, \tilde{\mathbf{e}}^{(j)} \rangle = 1 \ \forall j \in \{1, \dots, n_2\} \right\}. \end{aligned} \tag{5}$$

In other words, if we let $\delta_i = 1 - \sum_j v_{ij} = 1 - \langle \mathbf{v}, \mathbf{e}^{(i)} \rangle$ and $\tilde{\delta}_j = 1 - \sum_i v_{ij} = \langle \mathbf{v}, \tilde{\mathbf{e}}^{(j)} \rangle$, we have that $\delta_i = 0$ for $i = 1, \dots, n_1$, $\tilde{\delta}_j = 0$ for $j = 1, \dots, n_2$ and $\delta_i > 0$ for $i > n_1$, $\tilde{\delta}_j > 0$ for $j > n_2$.

We consider a sequence of systems indexed by ϵ . In this section, we consider an i.i.d. arrival process $\mathbf{a}^{(\epsilon)}(t)$ with mean and variance given by

$$\begin{aligned} \mathbb{E}[\mathbf{a}^{(\epsilon)}(t)] &= \boldsymbol{\lambda}^{(\epsilon)} = \mathbf{v} - \epsilon \mathbf{k}, \\ \text{Var}[\mathbf{a}^{(\epsilon)}(t)] &= (\boldsymbol{\sigma}^{(\epsilon)})^2, \end{aligned}$$

such that as $\epsilon \rightarrow 0$, $(\boldsymbol{\sigma}^{(\epsilon)})^2 \rightarrow \boldsymbol{\sigma}^2$. Here $\mathbf{k} \in \mathbb{R}_+^{n_2}$ is vector that represents the rates of saturation of different ports. Define

$$\begin{aligned} \kappa_i(\mathbf{k}) &\triangleq \langle \mathbf{k}, \mathbf{e}^{(i)} \rangle = \sum_j k_{ij} \text{ and} \\ \tilde{\kappa}_j(\mathbf{k}) &\triangleq \langle \mathbf{k}, \tilde{\mathbf{e}}^{(j)} \rangle = \sum_i k_{ij}. \end{aligned} \tag{6}$$

For simplicity of notation, we will suppress the dependence on \mathbf{k} . Note that $\sum_i \kappa_i = \sum_j \tilde{\kappa}_j$. Let $\kappa_{avg} \triangleq \sum_i \kappa_i / n = \sum_j \tilde{\kappa}_j / n$, $\kappa_{\min} = \min_i \kappa_i$, $\tilde{\kappa}_{\min} = \min_j \tilde{\kappa}_j$ and similarly for κ_{\max} , $\tilde{\kappa}_{\max}$. Note that in [11], the setting when $\mathbf{k} = \mathbf{v}$ is studied, in which case $\kappa_i = 1$ and $\tilde{\kappa}_j = 1$ for all i, j . In order to make sure that the heavy-traffic parameter ϵ is comparable to this case, we assume without loss of generality that $\kappa_{avg} = 1$. In other words, we normalize the vector \mathbf{k} by assuming that $\langle \mathbf{k}, \mathbf{1} \rangle = n$. Such a normalized \mathbf{k} is called the saturation rate vector. We will study the switch in the heavy-traffic limit as $\epsilon \downarrow 0$. Define

$$\begin{aligned} \gamma_i^{(\epsilon)} &\triangleq 1 - \sum_j \lambda_{ij}^{(\epsilon)} = \delta_i + \epsilon \kappa_i, \\ \tilde{\gamma}_j^{(\epsilon)} &\triangleq 1 - \sum_i \lambda_{ij}^{(\epsilon)} = \tilde{\delta}_j + \epsilon \tilde{\kappa}_j. \end{aligned}$$

¹ The relative interior of a set is defined as its interior relative to its affine hull [5, Section 2.1.3].

Note that $\gamma_i^{(\epsilon)} = \epsilon\kappa_i$ for $i \leq n_1$, and $\tilde{\gamma}_j^{(\epsilon)} = \epsilon\tilde{\kappa}_j$ for $j \leq n_2$. For the unsaturated ports, $\lim_{\epsilon \downarrow 0} \gamma_i^{(\epsilon)} = \delta_i > 0$ for $i > n_1$, and $\lim_{\epsilon \downarrow 0} \tilde{\gamma}_j^{(\epsilon)} = \tilde{\delta}_j > 0$ for $j > n_2$.

3.1 Universal lower bound

We now present lower bounds on the steady state queue lengths that are satisfied by any scheduling algorithm.

Proposition 1 Consider a set of switch systems with the arrival processes $\mathbf{a}^{(\epsilon)}(t)$ described above, parameterized by $0 < \epsilon < 1$, such that the mean arrival rate vector is $\lambda^\epsilon = \mathbf{v} - \epsilon\boldsymbol{\kappa}$ for some $\mathbf{v} \in \text{Relint}(\mathcal{F}_{n_1 n_2})$, and the variance is $(\boldsymbol{\sigma}^{(\epsilon)})^2$. For $1 \leq i, j \leq n$, $\gamma_i, \tilde{\gamma}_j$ are defined as above. Fix a scheduling policy under which the switch system is stable for any $0 < \epsilon < 1$. Let $\mathbf{q}^{(\epsilon)}(t)$ denote the queue lengths process under this policy for each system. Suppose that this process converges in distribution to a steady state random vector $\bar{\mathbf{q}}^{(\epsilon)}$. Then, for each of these systems, the steady state mean queue lengths can be lower-bounded as follows:

$$\mathbb{E} \left[\sum_j \bar{q}_{ij}^{(\epsilon)} \right] \geq \frac{\sum_j (\sigma_{ij}^{(\epsilon)})^2}{2\gamma_i} - \frac{1 - \gamma_i}{2} \text{ for all } 1 \leq i \leq n, \tag{7}$$

$$\mathbb{E} \left[\sum_i \bar{q}_{ij}^{(\epsilon)} \right] \geq \frac{\sum_i (\sigma_{ij}^{(\epsilon)})^2}{2\tilde{\gamma}_j} - \frac{1 - \tilde{\gamma}_j}{2} \text{ for all } 1 \leq j \leq n. \tag{8}$$

Therefore, in the heavy-traffic limit as $\epsilon \downarrow 0$, if $(\boldsymbol{\sigma}^{(\epsilon)})^2 \rightarrow \boldsymbol{\sigma}^2$, for the saturated ports we have

$$\liminf_{\epsilon \downarrow 0} \epsilon \mathbb{E} \left[\sum_j \bar{q}_{ij}^{(\epsilon)} \right] \geq \frac{\sum_j \sigma_{ij}^2}{2\kappa_i} \text{ for all } 1 \leq i \leq n_1,$$

$$\liminf_{\epsilon \downarrow 0} \epsilon \mathbb{E} \left[\sum_i \bar{q}_{ij}^{(\epsilon)} \right] \geq \frac{\sum_i \sigma_{ij}^2}{2\tilde{\kappa}_j} \text{ for all } 1 \leq i \leq n_2,$$

and for the unsaturated ports we have

$$\liminf_{\epsilon \downarrow 0} \epsilon \mathbb{E} \left[\sum_j \bar{q}_{ij}^{(\epsilon)} \right] \geq 0 \text{ for all } i > n_1,$$

$$\liminf_{\epsilon \downarrow 0} \epsilon \mathbb{E} \left[\sum_i \bar{q}_{ij}^{(\epsilon)} \right] \geq 0 \text{ for all } i > n_2.$$

Proof The queue lengths at each port can be lower-bounded by a single-server queue as follows: Consider $\sum_j q_{ij}^{(\epsilon)}(t)$, the total queue length at input port (row) i . It can

be lower-bounded sample path wise by a coupled single-server queue with arrival process $\sum_j a_{ij}^{(\epsilon)}(t)$ [11, Proposition 1]. The mean and variance of the arrival process for this single-server queue are then $(1 - \gamma_i)$ and $(\sigma_{ij}^{(\epsilon)})^2$, respectively, because of the independence of the arrival processes across the queues in the matrix. Then, using the Kingman bound for a single-server queue in Lemma 1, we get (7). Similarly lower bounding the total queue length for output port (column) j , $\sum_i q_{ij}^{(\epsilon)}(t)$, by a single-server queue, we get (8). Taking the heavy-traffic limits using the facts that $\gamma_i^{(\epsilon)} = \epsilon \kappa_i$, $\tilde{\gamma}_j^{(\epsilon)} = \epsilon \tilde{\kappa}_j$ for saturated ports and $\lim_{\epsilon \downarrow 0} \gamma_i^{(\epsilon)} > 0$, $\lim_{\epsilon \downarrow 0} \tilde{\gamma}_j^{(\epsilon)} > 0$ for unsaturated ports, gives the heavy-traffic limits. \square

3.2 State space collapse

In this subsection, we will show that under the MaxWeight algorithm the queue lengths vector concentrates close to a lower-dimensional cone. In order to make this more precise, we need to first present the following definitions.

The heavy-traffic rate vector \mathbf{v} lies in the relative interior of the face $\mathcal{F}_{n_1 n_2}$ which is at the intersection of hyperplanes with the $n_1 + n_2$ normal vectors, $\{\mathbf{e}^{(i)}$ for $1 \leq i \leq n_1\} \cup \{\tilde{\mathbf{e}}^{(j)}$ for $1 \leq j \leq n_2\}$. Call the cone spanned by these normal vector $\mathcal{K}_{n_1 n_2}$, and the subspace spanned by these normal vectors $\mathcal{S}_{n_1 n_2}$, i.e.,

$$\begin{aligned}
 \mathcal{K}_{n_1 n_2} &\triangleq \left\{ \mathbf{x} \in \mathbb{R}^{n^2} : \mathbf{x} = \sum_{i=1}^{n_1} w_i \mathbf{e}^{(i)} + \sum_{j=1}^{n_2} \tilde{w}_j \tilde{\mathbf{e}}^{(j)} \text{ where} \right. \\
 &\quad \left. w_i \in \mathbb{R}_+ \text{ for } 1 \leq i \leq n_1, \tilde{w}_j \in \mathbb{R}_+ \text{ for } 1 \leq j \leq n_2 \right\} \\
 &= \left\{ \mathbf{x} \in \mathbb{R}^{n^2} : \mathbf{x} = \sum_{i=1}^n w_i \mathbf{e}^{(i)} + \sum_{j=1}^n \tilde{w}_j \tilde{\mathbf{e}}^{(j)} \text{ where} \right. \\
 &\quad \left. w_i \in \mathbb{R}_+ \text{ for } 1 \leq i \leq n_1 \text{ and } w_i = 0 \text{ for } i > n_1, \right. \\
 &\quad \left. \tilde{w}_j \in \mathbb{R}_+ \text{ for } 1 \leq j \leq n_2 \text{ and } \tilde{w}_j = 0 \text{ for } j > n_2 \right\}. \\
 \\
 \mathcal{S}_{n_1 n_2} &\triangleq \left\{ \mathbf{x} \in \mathbb{R}^{n^2} : \mathbf{x} = \sum_{i=1}^{n_1} w_i \mathbf{e}^{(i)} + \sum_{j=1}^{n_2} \tilde{w}_j \tilde{\mathbf{e}}^{(j)} \text{ where} \right. \\
 &\quad \left. w_i \in \mathbb{R} \text{ for } 1 \leq i \leq n_1, \tilde{w}_j \in \mathbb{R} \text{ for } 1 \leq j \leq n_2 \right\} \tag{9} \\
 &= \left\{ \mathbf{x} \in \mathbb{R}^{n^2} : \mathbf{x} = \sum_{i=1}^n w_i \mathbf{e}^{(i)} + \sum_{j=1}^n \tilde{w}_j \tilde{\mathbf{e}}^{(j)} \text{ where} \right. \\
 &\quad \left. w_i \in \mathbb{R} \text{ for } 1 \leq i \leq n_1 \text{ and } w_i = 0 \text{ for } i > n_1, \right. \\
 &\quad \left. \tilde{w}_j \in \mathbb{R} \text{ for } 1 \leq j \leq n_2 \text{ and } \tilde{w}_j = 0 \text{ for } j > n_2 \right\}.
 \end{aligned}$$

The components of any vector \mathbf{x} in the subspace $\mathcal{S}_{n_1 n_2}$ can be written in the form $x_{ij} = w_i + \tilde{w}_j$, where $w_i \in \mathbb{R}$ for $1 \leq i \leq n_1$, $w_i = 0$ for $i > n_1$, $\tilde{w}_j \in \mathbb{R}$ for $1 \leq j \leq n_2$, and $\tilde{w}_j = 0$ for $j > n_1$. The same is true for any vector \mathbf{x} in the cone $\mathcal{K}_{n_1 n_2}$ with the further restriction that $w_i \geq 0$ for $1 \leq i \leq n_1$ and $\tilde{w}_j \geq 0$ for $1 \leq j \leq n_2$. This leads to the following lemma relating the structure of the cone $\mathcal{K}_{n_1 n_2}$ and the subspace $\mathcal{S}_{n_1 n_2}$.

Lemma 3 *Let $n_1 < n$ and $n_2 < n$. The cone $\mathcal{K}_{n_1 n_2}$ is the intersection of the space $\mathcal{S}_{n_1 n_2}$ and the positive orthant, i.e.,*

$$\mathcal{K}_{n_1 n_2} = \mathcal{S}_{n_1 n_2} \cap \mathbb{R}_+^{n^2}.$$

Proof From the definitions above, it is clear that $\mathcal{K}_{n_1 n_2} \subseteq \mathcal{S}_{n_1 n_2}$ and $\mathcal{K}_{n_1 n_2} \subseteq \mathbb{R}_+^{n^2}$. Therefore, we have $\mathcal{K}_{n_1 n_2} \subseteq \mathcal{S}_{n_1 n_2} \cap \mathbb{R}_+^{n^2}$. Now suppose that $\mathbf{x} \in \mathcal{S}_{n_1 n_2} \cap \mathbb{R}_+^{n^2}$. We have that $x_{ij} = w_i + \tilde{w}_j \geq 0$, where $w_i = 0$ for $i > n_1$ and $\tilde{w}_j = 0$ for $j > n_1$. Since $n_1 < n$, we have $x_{in} = w_i \geq 0$, and so we get that $w_i \geq 0$ for $1 \leq i \leq n_1$. Similarly, we get that $\tilde{w}_j \geq 0$ for $1 \leq j \leq n_2$, proving that $\mathbf{x} \in \mathcal{K}_{n_1 n_2}$ and so $\mathcal{S}_{n_1 n_2} \cap \mathbb{R}_+^{n^2} \subseteq \mathcal{K}_{n_1 n_2}$. \square

Let $\mathbf{x}_{\parallel \mathcal{K}}$ denote the projection of \mathbf{x} onto the convex cone $\mathcal{K}_{n_1 n_2}$, and let $\mathbf{x}_{\perp \mathcal{K}} \triangleq \mathbf{x} - \mathbf{x}_{\parallel \mathcal{K}}$ be the perpendicular component. Similarly, let $\mathbf{x}_{\parallel \mathcal{S}}$ denote the projection of \mathbf{x} onto the subspace $\mathcal{S}_{n_1 n_2}$, and let $\mathbf{x}_{\perp \mathcal{S}} \triangleq \mathbf{x} - \mathbf{x}_{\parallel \mathcal{S}}$ be the perpendicular component. For simplicity of notation, we suppress the dependence on n_1 and n_2 . In [11], \mathbf{x}_{\parallel} and \mathbf{x}_{\perp} were used to denote the projections we denote here by $\mathbf{x}_{\parallel \mathcal{K}}$ and $\mathbf{x}_{\perp \mathcal{K}}$, respectively. We will show that under the MaxWeight algorithm, all the moments of $\mathbf{q}_{\perp \mathcal{K}}$ are bounded in steady state independent of ϵ . Since the ℓ_1 norm of the queues length vector $\|\mathbf{q}\|_1$, is $\Omega(1/\epsilon)$, as shown in the previous subsection, this establishes that the perpendicular component $\mathbf{q}_{\perp \mathcal{K}}$ is a negligible part of the queue lengths vector \mathbf{q} for small ϵ . Thus, we establish state space collapse onto the cone $\mathcal{K}_{n_1 n_2}$.

For $\mathbf{v} \in \text{Relint}(\mathcal{F}_{n_1 n_2})$, the vector in the relative interior of the face $\mathcal{F}_{n_1 n_2}$, let $v_{\min} \triangleq \min_{i,j} v_{ij}$. We assume that $v_{\min} > 0$. Then, $v'_{\min} > 0$, where

$$v'_{\min} \triangleq \min \left\{ v_{\min}, \min_{i > n_1} \left\{ 1 - \langle \mathbf{v}, \mathbf{e}^{(i)} \rangle \right\}, \min_{j > n_2} \left\{ 1 - \langle \mathbf{v}, \tilde{\mathbf{e}}^{(j)} \rangle \right\} \right\}.$$

Proposition 2 *Consider a set of switch systems under the MaxWeight scheduling algorithm, with the arrival processes $\mathbf{a}^{(\epsilon)}(t)$, parameterized by $0 < \epsilon < 1$, and maximum possible arrivals in any queue a_{\max} . The mean arrival rate vector is $\lambda^\epsilon = \mathbf{v} - \epsilon \mathbf{k}$ for some $\mathbf{v} \in \text{Relint}(\mathcal{F}_{n_1 n_2})$ such that $v_{\min} \triangleq \min_{i,j} v_{ij} > 0$, and a normalized saturation rate vector $\mathbf{k} \in \mathbb{R}_+^{n^2}$ such that $\langle \mathbf{k}, \mathbf{1} \rangle = n$. Let the variance $(\boldsymbol{\sigma}^{(\epsilon)})^2$ of the arrival process be such that $\|\boldsymbol{\sigma}^{(\epsilon)}\|^2 \leq \tilde{\sigma}^2$ for some $\tilde{\sigma}^2$ not dependent on ϵ . Let $\mathbf{q}^{(\epsilon)}(t)$ denote the queue lengths process of each system, which is positive recurrent. Therefore, the process $\mathbf{q}^{(\epsilon)}(t)$ converges to a steady state random vector in distribution, which we denote by $\bar{\mathbf{q}}^{(\epsilon)}$. Then, for each system with $0 < \epsilon \leq v'_{\min}/2\|\mathbf{k}\|$, the steady state queue lengths vector satisfies*

$$\mathbb{E} \left[\|\bar{\mathbf{q}}_{\perp \mathcal{S}}^{(\epsilon)}\|^r \right] \leq \mathbb{E} \left[\|\bar{\mathbf{q}}_{\perp \mathcal{K}}^{(\epsilon)}\|^r \right] \leq (M_r)^r \quad \forall r \in \{1, 2, \dots\},$$

where v'_{\min} is defined as above and M_r is a function of $r, \tilde{\sigma}, \mathbf{v}, a_{\max}, v'_{\min}$ but independent of ϵ .

Proof We omit the superscript (ϵ) in this proof for simplicity of notation. For the Markov chain \mathbf{q} , consider the Lyapunov function $W_{\perp \mathcal{K}}(\mathbf{q}) \triangleq \|\mathbf{q}_{\perp \mathcal{K}}\|$. We will use Lemma 2 to obtain moment bounds from the drift of $W_{\perp \mathcal{K}}(\cdot)$. Similar to [11, Proposition 2], under the MaxWeight scheduling algorithm it can be shown that

$$\begin{aligned} \mathbb{E} \left[\Delta W_{\perp \mathcal{K}}(\mathbf{q}) \mid \mathbf{q}(t) = \mathbf{q} \right] &\leq \frac{1}{2\|\mathbf{q}_{\perp \mathcal{K}}\|} \left(\|\lambda\|^2 + \|\sigma\|^2 + n - 2\epsilon \langle \mathbf{q}_{\perp \mathcal{K}}, \mathbf{k} \rangle \right. \\ &\quad \left. + 2\mathbb{E} \left[\langle \mathbf{q}_{\parallel \mathcal{K}}, \mathbf{s}(t) - \mathbf{v} \rangle \mid \mathbf{q}(t) = \mathbf{q} \right] + 2 \min_{\mathbf{r} \in \mathcal{C}} \langle \mathbf{q}, \mathbf{v} - \mathbf{r} \rangle \right). \end{aligned} \tag{10}$$

Recall that since $\mathbf{s} \in \mathcal{S}^*$, $\langle \mathbf{e}^{(i)}, \mathbf{s}(t) \rangle = 1$ and $\langle \tilde{\mathbf{e}}^{(j)}, \mathbf{s}(t) \rangle = 1$ for all i, j, t . Similarly, since $\mathbf{v} \in \mathcal{F}_{n_1 n_2}$, for $1 \leq i \leq n_1$ and $1 \leq j \leq n_2$, we have $\langle \mathbf{e}^{(i)}, \mathbf{v} \rangle = 1$ and $\langle \tilde{\mathbf{e}}^{(j)}, \mathbf{v} \rangle = 1$. By the definition of the cone $\mathcal{K}_{n_1 n_2}$, the vector $\mathbf{q}_{\parallel \mathcal{K}}$ can be written as $\mathbf{q}_{\parallel \mathcal{K}} = \sum_{i=1}^{n_1} w_i \mathbf{e}^{(i)} + \sum_{j=1}^{n_2} \tilde{w}_j \tilde{\mathbf{e}}^{(j)}$. Putting all these together, we get $\langle \mathbf{q}_{\parallel \mathcal{K}}, \mathbf{s}(t) - \mathbf{v} \rangle = 0$. We now use the following claim to bound the last term in (10).

Claim 1 For any $\mathbf{q} \in \mathbb{R}^{n^2}$ and $\mathbf{v} \in \text{Relint}(\mathcal{F}_{n_1 n_2})$ such that $v_{\min} > 0$,

$$\mathbf{v} + \frac{v'_{\min}}{\|\mathbf{q}_{\perp \mathcal{K}}\|} \mathbf{q}_{\perp \mathcal{K}} \in \mathcal{C}.$$

Proof We will verify that $\mathbf{v} + \frac{v'_{\min}}{\|\mathbf{q}_{\perp \mathcal{K}}\|} \mathbf{q}_{\perp \mathcal{K}}$ satisfies all the conditions in the definition of \mathcal{C} in (4). Note that $\frac{\mathbf{q}_{\perp \mathcal{K}}}{\|\mathbf{q}_{\perp \mathcal{K}}\|}$ is a unit vector along some direction. Since $v_{ij} \geq v'_{\min}$, clearly $\mathbf{v} + \frac{v'_{\min}}{\|\mathbf{q}_{\perp \mathcal{K}}\|} \mathbf{q}_{\perp \mathcal{K}} \in \mathbb{R}_+^{n^2}$.

It is well-known that for any $\mathbf{x} \in \mathcal{K}_{n_1 n_2}$, $\langle \mathbf{q}_{\perp \mathcal{K}}, \mathbf{x} \rangle \leq 0$. Since $\mathbf{e}^{(i)} \in \mathcal{K}_{n_1 n_2}$ for $1 \leq i \leq n_1$, we have $\langle \mathbf{q}_{\perp \mathcal{K}}, \mathbf{e}^{(i)} \rangle \leq 0$. Then, using the fact that $\mathbf{v} \in \mathcal{F}_{n_1 n_2}$, we have, for $1 \leq i \leq n_1$,

$$\left\langle \mathbf{v} + \frac{v'_{\min}}{\|\mathbf{q}_{\perp \mathcal{K}}\|} \mathbf{q}_{\perp \mathcal{K}}, \mathbf{e}^{(i)} \right\rangle \leq 1.$$

For $i > n_1$,

$$\begin{aligned} \left\langle \mathbf{v} + \frac{v'_{\min}}{\|\mathbf{q}_{\perp \mathcal{K}}\|} \mathbf{q}_{\perp \mathcal{K}}, \mathbf{e}^{(i)} \right\rangle &= \langle \mathbf{v}, \mathbf{e}^{(i)} \rangle + v'_{\min} \left\langle \frac{\mathbf{q}_{\perp \mathcal{K}}}{\|\mathbf{q}_{\perp \mathcal{K}}\|}, \mathbf{e}^{(i)} \right\rangle \\ &\stackrel{(a)}{\leq} \langle \mathbf{v}, \mathbf{e}^{(i)} \rangle + v'_{\min} \leq 1, \end{aligned}$$

where (a) follows from the Cauchy–Schwartz inequality and the last inequality follows from the definition of v'_{\min} . It can similarly be shown that $\left\langle \mathbf{v} + \frac{v'_{\min}}{\|\mathbf{q}_{\perp\mathcal{K}}\|} \mathbf{q}_{\perp\mathcal{K}}, \tilde{\mathbf{e}}^{(j)} \right\rangle \leq 1$ for $1 \leq j \leq n_2$ as well as $j > n_2$, proving the claim. \square

Using the claim, the last term in (10) can be bounded as

$$\begin{aligned} 2 \min_{\mathbf{r} \in \mathcal{C}} \langle \mathbf{q}, \mathbf{v} - \mathbf{r} \rangle &\leq 2 \left\langle \mathbf{q}, \mathbf{v} - \left(\mathbf{v} + \frac{v'_{\min}}{\|\mathbf{q}_{\perp\mathcal{K}}\|} \mathbf{q}_{\perp\mathcal{K}} \right) \right\rangle \\ &= -2 \left\langle \mathbf{q}, \frac{v'_{\min}}{\|\mathbf{q}_{\perp\mathcal{K}}\|} \mathbf{q}_{\perp\mathcal{K}} \right\rangle \\ &= -2v'_{\min} \|\mathbf{q}_{\perp\mathcal{K}}\|. \end{aligned}$$

Using this in (10) and bounding the $-2\epsilon \langle \mathbf{q}_{\perp\mathcal{K}}, \mathbf{k} \rangle$ term using the Cauchy–Schwartz inequality, we get

$$\begin{aligned} \mathbb{E} \left[\Delta W_{\perp\mathcal{K}}(\mathbf{q}) \mid \mathbf{q}(t) = \mathbf{q} \right] &\leq \frac{\|\lambda\|^2 + \|\sigma\|^2 + n}{2\|\mathbf{q}_{\perp\mathcal{K}}\|} - v'_{\min} + \epsilon \|\mathbf{k}\| \\ &\leq \frac{\|\lambda\|^2 + \|\sigma\|^2 + n}{2\|\mathbf{q}_{\perp\mathcal{K}}\|} - \frac{v'_{\min}}{2}, \text{ whenever } \epsilon \leq \frac{v'_{\min}}{2\|\mathbf{k}\|} \\ &\leq -\frac{v'_{\min}}{4}, \text{ for all } \mathbf{q} \text{ such that } W_{\perp\mathcal{K}}(\mathbf{q}) \geq \frac{2(\|\lambda\|^2 + \|\sigma\|^2 + n)}{v'_{\min}}. \end{aligned}$$

Thus, the conditions of Lemma 2 is valid with $\zeta = \frac{2(\|\lambda\|^2 + \|\sigma\|^2 + n)}{v'_{\min}}$ and $\eta = \frac{v'_{\min}}{4}$.

Moreover, ζ can be upper-bounded by $\zeta \leq \frac{2(\|\mathbf{v}\|^2 + \|\sigma\|^2 + n)}{v'_{\min}}$, an expression that doesn't contain ϵ . The conditions of Lemma 2 can be verified using nonexpansivity of projection and the fact that the maximum arrivals at every time are a_{\max} [11]. Then, from Lemma 2, we get the bound on $\mathbb{E} \left[\|\bar{\mathbf{q}}_{\perp\mathcal{K}}^{(\epsilon)}\|^r \right]$ in the proposition. Since $\mathcal{K}_{n_1 n_2} \subseteq \mathcal{S}_{n_1 n_2}$, we have $\|\bar{\mathbf{q}}_{\perp\mathcal{S}}^{(\epsilon)}\|^r \leq \|\bar{\mathbf{q}}_{\perp\mathcal{K}}^{(\epsilon)}\|^r$, completing the proof. \square

3.3 Asymptotically tight upper and lower bounds under the MaxWeight policy

In this subsection, we will use the state space collapse result from the previous section to obtain lower and upper bounds on weighted sums of queue lengths under the MaxWeight algorithm that are tight in heavy-traffic limit. It turns out that the queue length behavior under the MaxWeight algorithm when there is at least one unsaturated port is qualitatively different from the case when all ports are saturated. The reason for this is discussed in Corollary 5 in Sect. 4. So, in this subsection, we focus on the case when at least one port is not saturated. If all the input ports are saturated, from the definition of the capacity region, it follows that all the output ports are also saturated, i.e., whenever $n_1 = n$, we also have $n_2 = n$. Similarly, if all the output ports are saturated, it again follows that all the input ports are saturated. Since we are interested in an incompletely saturated switch, in this section we assume $n_1 < n$ and $n_2 < n$.

The queue length bounds are obtained by setting the drift of the following function to zero in steady state:

$$V(\mathbf{q}) = \|\mathbf{q}\|_S^2.$$

Its drift is defined as

$$\Delta V(\mathbf{q}) \triangleq [V(\mathbf{q}(t + 1)) - V(\mathbf{q}(t))] \mathcal{I}(\mathbf{q}(t) = \mathbf{q}).$$

We now state the main result of the paper in a general form. In Sect. 4, we will interpret this result as well as present various special cases.

Theorem 1 Consider a set of switch systems under the MaxWeight scheduling algorithm, with the arrival processes $\mathbf{a}^{(\epsilon)}(t)$, parameterized by $0 < \epsilon < 1$, and maximum possible arrivals in any queue a_{\max} . The mean arrival rate vector is $\lambda^\epsilon = \mathbf{v} - \epsilon \mathbf{k}$ for some $\mathbf{v} \in \text{Relint}(\mathcal{F}_{n_1 n_2})$ such that $v_{\min} \triangleq \min_{ij} v_{ij} > 0$, and a normalized saturation rate vector $\mathbf{k} \in \mathbb{R}_+^{n^2}$ such that $\langle \mathbf{k}, \mathbf{1} \rangle = n$. Let the variance $(\sigma^{(\epsilon)})^2$ of the arrival process be such that $\|\sigma^{(\epsilon)}\|^2 \leq \tilde{\sigma}^2$ for some $\tilde{\sigma}^2$ not dependent on ϵ and assume that $n_1 < n$ and $n_2 < n$. Let $\mathbf{q}^{(\epsilon)}(t)$ denote the queue lengths process of each system, which is positive recurrent. Therefore, the process $\mathbf{q}^{(\epsilon)}(t)$ converges to a steady state random vector in distribution, which we denote by $\bar{\mathbf{q}}^{(\epsilon)}$. Then, for each system with $0 < \epsilon \leq v'_{\min}/2\|\mathbf{k}\|$, the steady state queue lengths vector satisfies

$$\frac{1}{2\epsilon} \left\langle (\sigma^{(\epsilon)})^2, \boldsymbol{\zeta} \right\rangle - B_1(\epsilon) \leq \mathbb{E} \left[\left\langle \bar{\mathbf{q}}^{(\epsilon)}, \boldsymbol{\alpha} \right\rangle \right] \leq \frac{1}{2\epsilon} \left\langle (\sigma^{(\epsilon)})^2, \boldsymbol{\zeta} \right\rangle + B_2(\epsilon)$$

for any fixed weight vector $\boldsymbol{\alpha} \in \mathbb{R}^{n^2}$ such that $\langle \boldsymbol{\alpha}, \mathbf{e}^{(i)} \rangle = n\kappa_i$ for $i \leq n_1$ and $\langle \boldsymbol{\alpha}, \tilde{\mathbf{e}}^{(j)} \rangle = n\tilde{\kappa}_j$ for $j \leq n_2$, where $B_1(\epsilon)$ as well as $B_2(\epsilon)$ are $o(\frac{1}{\epsilon})$, i.e., $\lim_{\epsilon \rightarrow 0} \epsilon B_1(\epsilon) = 0$ and $\lim_{\epsilon \rightarrow 0} \epsilon B_2(\epsilon) = 0$. The vector $\boldsymbol{\zeta}$ is defined by

$$\zeta_{ij} \triangleq \begin{cases} 2 - \frac{2n-n_1-n_2}{n^2-n_1n_2} & \text{if } i \leq n_1 \text{ and } j \leq n_2 \\ 1 + \frac{n_2}{n^2-n_1n_2} & \text{if } i \leq n_1 \text{ and } j > n_2 \\ 1 + \frac{n_1}{n^2-n_1n_2} & \text{if } i > n_1 \text{ and } j \leq n_2 \\ 0 & \text{if } i > n_1 \text{ and } j > n_2. \end{cases} \tag{11}$$

Thus, in the heavy-traffic limit as $\epsilon \downarrow 0$, we have

$$\lim_{\epsilon \rightarrow 0} \epsilon \mathbb{E} \left[\left\langle \bar{\mathbf{q}}^{(\epsilon)}, \boldsymbol{\alpha} \right\rangle \right] = \frac{1}{2} \langle \sigma^2, \boldsymbol{\zeta} \rangle.$$

Moreover, for any $i > n_1$ and $j > n_2$,

$$\lim_{\epsilon \rightarrow 0} \epsilon \mathbb{E} \left[\bar{q}_{ij}^{(\epsilon)} \right] = 0.$$

Note that, in general, the weights α_{ij} are allowed to be negative. We now present the proof of the theorem.

Proof We consider the switch for a fixed $0 < \epsilon \leq v'_{\min}/2\|\mathbf{k}\|$. For simplicity of notation, we again omit the superscript (ϵ) in this proof. Similar to the notation in [11], we use $\bar{\mathbf{q}}$ to denote the steady state queue length vector and $\bar{\mathbf{a}}$ to denote the steady state arrival vector, which is identically distributed to the vector $\mathbf{a}(t)$ at any time t . We use $\mathbf{s}(\bar{\mathbf{q}})$ and $\mathbf{u}(\bar{\mathbf{q}})$ for the schedule and unused service to explicitly show their dependence on the queue lengths. If the queue length at time t is $\bar{\mathbf{q}}$, then the queue length at time $t + 1$, $\bar{\mathbf{q}} + \bar{\mathbf{a}} - \mathbf{s}(\bar{\mathbf{q}}) + \mathbf{u}(\bar{\mathbf{q}})$, is denoted by $\bar{\mathbf{q}}^+$. Since $\bar{\mathbf{q}}$ is the steady state queue length, it has the same distribution as $\bar{\mathbf{q}}^+$.

It can be easily shown using Lemma 2 that, in steady state, $\mathbb{E}[\|\bar{\mathbf{q}}\|^2]$ is finite, and consequently we have

$$\mathbb{E}[V(\bar{\mathbf{q}})] < \infty \quad \text{and} \quad \mathbb{E}[\|\bar{\mathbf{q}}\|_1] = \mathbb{E}\left[\sum_{ij} \bar{q}_{ij}\right] < \infty, \tag{12}$$

where $\|\cdot\|_1$ denotes the ℓ_1 norm. See Lemma 5 in [11] for details. Setting the drift of $V(\bar{\mathbf{q}})$ to zero in steady state, we get

$$\begin{aligned} 0 &= \mathbb{E}[\Delta V(\bar{\mathbf{q}})] \\ &= \mathbb{E}[V(\bar{\mathbf{q}} + \bar{\mathbf{a}} - \mathbf{s}(\bar{\mathbf{q}}) + \mathbf{u}(\bar{\mathbf{q}})) - V(\bar{\mathbf{q}})] \\ &= \mathbb{E}\left[\|(\bar{\mathbf{q}} + \bar{\mathbf{a}} - \mathbf{s}(\bar{\mathbf{q}}) + \mathbf{u}(\bar{\mathbf{q}}))_{\parallel\mathcal{S}}\|^2 - \|\bar{\mathbf{q}}_{\parallel\mathcal{S}}\|^2\right] \\ &\stackrel{(a)}{=} \mathbb{E}\left[\|\bar{\mathbf{q}}_{\parallel\mathcal{S}} + (\bar{\mathbf{a}} - \mathbf{s}(\bar{\mathbf{q}}))_{\parallel\mathcal{S}} + \mathbf{u}_{\parallel\mathcal{S}}(\bar{\mathbf{q}})\|^2 - \|\bar{\mathbf{q}}_{\parallel\mathcal{S}}\|^2\right] \\ &= \mathbb{E}\left[\|\bar{\mathbf{q}}_{\parallel\mathcal{S}} + (\bar{\mathbf{a}} - \mathbf{s}(\bar{\mathbf{q}}))_{\parallel\mathcal{S}}\|^2 + 2\langle \bar{\mathbf{q}}_{\parallel\mathcal{S}} + (\bar{\mathbf{a}} - \mathbf{s}(\bar{\mathbf{q}}))_{\parallel\mathcal{S}}, \mathbf{u}_{\parallel\mathcal{S}}(\bar{\mathbf{q}}) \rangle\right] \\ &\quad + \mathbb{E}\left[\|\mathbf{u}_{\parallel\mathcal{S}}(\bar{\mathbf{q}})\|^2 - \|\bar{\mathbf{q}}_{\parallel\mathcal{S}}\|^2\right] \\ &= \mathbb{E}\left[\|\bar{\mathbf{q}}_{\parallel\mathcal{S}} + (\bar{\mathbf{a}} - \mathbf{s}(\bar{\mathbf{q}}))_{\parallel\mathcal{S}}\|^2 + 2\langle \bar{\mathbf{q}}_{\parallel\mathcal{S}} + (\bar{\mathbf{a}} - \mathbf{s}(\bar{\mathbf{q}}))_{\parallel\mathcal{S}}, \mathbf{u}_{\parallel\mathcal{S}}(\bar{\mathbf{q}}) \rangle\right] \\ &\quad - \mathbb{E}\left[\|\mathbf{u}_{\parallel\mathcal{S}}(\bar{\mathbf{q}})\|^2 + \|\bar{\mathbf{q}}_{\parallel\mathcal{S}}\|^2\right] \\ &\stackrel{(b)}{=} \mathbb{E}\left[\|\bar{\mathbf{q}}_{\parallel\mathcal{S}} + (\bar{\mathbf{a}} - \mathbf{s}(\bar{\mathbf{q}}))_{\parallel\mathcal{S}}\|^2 + 2\langle \bar{\mathbf{q}}^+_{\parallel\mathcal{S}}, \mathbf{u}_{\parallel\mathcal{S}}(\bar{\mathbf{q}}) \rangle - \|\mathbf{u}_{\parallel\mathcal{S}}(\bar{\mathbf{q}})\|^2 - \|\bar{\mathbf{q}}_{\parallel\mathcal{S}}\|^2\right] \\ &= \mathbb{E}\left[\|(\bar{\mathbf{a}} - \mathbf{s}(\bar{\mathbf{q}}))_{\parallel\mathcal{S}}\|^2 + 2\langle \bar{\mathbf{q}}_{\parallel\mathcal{S}}, (\bar{\mathbf{a}} - \mathbf{s}(\bar{\mathbf{q}}))_{\parallel\mathcal{S}} \rangle - \|\mathbf{u}_{\parallel\mathcal{S}}(\bar{\mathbf{q}})\|^2 + 2\langle \bar{\mathbf{q}}^+_{\parallel\mathcal{S}}, \mathbf{u}_{\parallel\mathcal{S}}(\bar{\mathbf{q}}) \rangle\right], \end{aligned}$$

where (a) and (b) follow from the fact that projection onto a subspace is linear. Therefore, we have

$$2\mathbb{E}[\langle \bar{\mathbf{q}}_{\parallel\mathcal{S}}, (\mathbf{s}(\bar{\mathbf{q}}) - \bar{\mathbf{a}})_{\parallel\mathcal{S}} \rangle] = \mathbb{E}\left[\|(\bar{\mathbf{a}} - \mathbf{s}(\bar{\mathbf{q}}))_{\parallel\mathcal{S}}\|^2\right] - \mathbb{E}\left[\|\mathbf{u}_{\parallel\mathcal{S}}(\bar{\mathbf{q}})\|^2\right] \tag{13}$$

$$+ 2\mathbb{E}\left[\langle \bar{\mathbf{q}}^+_{\parallel\mathcal{S}}, \mathbf{u}_{\parallel\mathcal{S}}(\bar{\mathbf{q}}) \rangle\right]. \tag{14}$$

We will now study each of the terms in this equation. Consider the LHS term in (13). Since any vector of the form $\mathbf{x}_{\perp\mathcal{S}}$ is orthogonal to the space $\mathcal{S}_{n_1n_2}$, we get

$$2\mathbb{E}\left[\langle \bar{\mathbf{q}}_{\parallel\mathcal{S}}, (\mathbf{s}(\bar{\mathbf{q}}) - \bar{\mathbf{a}})_{\parallel\mathcal{S}} \rangle\right]$$

$$\begin{aligned}
 &= 2\mathbb{E} \left[\langle \bar{\mathbf{q}}_{\parallel \mathcal{S}}, (\mathbf{s}(\bar{\mathbf{q}}) - \bar{\mathbf{a}})_{\parallel \mathcal{S}} \rangle \right] + 2\mathbb{E} \left[\langle \bar{\mathbf{q}}_{\parallel \mathcal{S}}, (\mathbf{s}(\bar{\mathbf{q}}) - \bar{\mathbf{a}})_{\perp \mathcal{S}} \rangle \right] \\
 &= 2\mathbb{E} \left[\langle \bar{\mathbf{q}}_{\parallel \mathcal{S}}, \mathbf{s}(\bar{\mathbf{q}}) - \bar{\mathbf{a}} \rangle \right] \\
 &\stackrel{(a)}{=} 2\mathbb{E} \left[\langle \bar{\mathbf{q}}_{\parallel \mathcal{S}}, \mathbf{s}(\bar{\mathbf{q}}) - \boldsymbol{\lambda} \rangle \right] \\
 &= 2\mathbb{E} \left[\langle \bar{\mathbf{q}}_{\parallel \mathcal{S}}, \mathbf{s}(\bar{\mathbf{q}}) - (\mathbf{v} - \epsilon \mathbf{k}) \rangle \right] \\
 &= 2\epsilon \mathbb{E} \left[\langle \bar{\mathbf{q}}_{\parallel \mathcal{S}}, \mathbf{k} \rangle \right] + 2\mathbb{E} \left[\langle \bar{\mathbf{q}}_{\parallel \mathcal{S}}, \mathbf{s}(\bar{\mathbf{q}}) - \mathbf{v} \rangle \right] \\
 &\stackrel{(b)}{=} 2\epsilon \mathbb{E} \left[\left\langle \sum_{i=1}^{n_1} w_i \mathbf{e}^{(i)} + \sum_{j=1}^{n_2} \tilde{w}_j \tilde{\mathbf{e}}^{(j)}, \mathbf{k} \right\rangle \right] \\
 &\quad + 2\mathbb{E} \left[\left\langle \sum_{i=1}^{n_1} w_i \mathbf{e}^{(i)} + \sum_{j=1}^{n_2} \tilde{w}_j \tilde{\mathbf{e}}^{(j)}, \mathbf{s}(\bar{\mathbf{q}}) - \mathbf{v} \right\rangle \right] \\
 &\stackrel{(c)}{=} 2\epsilon \mathbb{E} \left[\sum_{i=1}^{n_1} w_i \langle \mathbf{e}^{(i)}, \mathbf{k} \rangle + \sum_{j=1}^{n_2} \tilde{w}_j \langle \tilde{\mathbf{e}}^{(j)}, \mathbf{k} \rangle \right] \\
 &\quad + 2\mathbb{E} \left[\sum_{i=1}^{n_1} w_i \langle \mathbf{e}^{(i)}, \mathbf{s}(\bar{\mathbf{q}}) - \mathbf{v} \rangle + \sum_{j=1}^{n_2} \tilde{w}_j \langle \tilde{\mathbf{e}}^{(j)}, \mathbf{s}(\bar{\mathbf{q}}) - \mathbf{v} \rangle \right] \\
 &\stackrel{(d)}{=} \frac{2\epsilon}{n} \mathbb{E} \left[\sum_{i=1}^{n_1} w_i \langle \mathbf{e}^{(i)}, \boldsymbol{\alpha} \rangle + \sum_{j=1}^{n_2} \tilde{w}_j \langle \tilde{\mathbf{e}}^{(j)}, \boldsymbol{\alpha} \rangle \right] \\
 &= \frac{2\epsilon}{n} \mathbb{E} \left[\left\langle \sum_{i=1}^{n_1} w_i \mathbf{e}^{(i)} + \sum_{j=1}^{n_2} \tilde{w}_j \tilde{\mathbf{e}}^{(j)}, \boldsymbol{\alpha} \right\rangle \right] \\
 &= \frac{2\epsilon}{n} \mathbb{E} \left[\langle \bar{\mathbf{q}}_{\parallel \mathcal{S}}, \boldsymbol{\alpha} \rangle \right] \\
 &= \frac{2\epsilon}{n} \mathbb{E} \left[\langle \bar{\mathbf{q}}, \boldsymbol{\alpha} \rangle \right] - \frac{2\epsilon}{n} \mathbb{E} \left[\langle \bar{\mathbf{q}}_{\perp \mathcal{S}}, \boldsymbol{\alpha} \rangle \right], \tag{15}
 \end{aligned}$$

where (a) follows from the fact that the arrivals are independent of queue lengths. From the definition of the space $\mathcal{S}_{n_1 n_2}$ in (9), we know that the vector $\bar{\mathbf{q}}_{\parallel \mathcal{S}}$ can be represented as $\sum_{i=1}^{n_1} w_i \mathbf{e}^{(i)} + \sum_{j=1}^{n_2} \tilde{w}_j \tilde{\mathbf{e}}^{(j)}$ for some $w_i \in \mathbb{R}, 1 \leq i \leq n_1$, and $\tilde{w}_j \in \mathbb{R}, 1 \leq j \leq n_2$, giving (b). Since the schedule is always assumed to be maximal, from (1), we have that $\langle \mathbf{e}^{(i)}, \mathbf{s} \rangle = 1$ for all i . Since the first n_1 rows are saturated, we have from (5) that $\langle \mathbf{e}^{(i)}, \mathbf{v} \rangle = 1$ for $1 \leq i \leq n_1$. Therefore, we get that $\langle \mathbf{e}^{(i)}, \mathbf{s}(\bar{\mathbf{q}}) - \mathbf{v} \rangle = 0$ for $1 \leq i \leq n_1$, and similarly we have $\langle \tilde{\mathbf{e}}^{(j)}, \mathbf{s}(\bar{\mathbf{q}}) - \mathbf{v} \rangle = 0$ for $1 \leq j \leq n_2$. Consequently, the second term in (c) vanishes. From the definition of κ_i in (6) and the assumption on $\boldsymbol{\alpha}$, we have that $\langle \mathbf{e}^{(i)}, \mathbf{k} \rangle = \kappa_i = \langle \mathbf{e}^{(i)}, \boldsymbol{\alpha} \rangle / n$ for $1 \leq i \leq n_1$. Similarly for $1 \leq j \leq n_2$, we have $\langle \tilde{\mathbf{e}}^{(j)}, \mathbf{k} \rangle = \kappa_j = \langle \tilde{\mathbf{e}}^{(j)}, \boldsymbol{\alpha} \rangle / n$, giving us (d). Using the Cauchy–Schwartz inequality, we can bound the last term in (15) as follows:

$$\begin{aligned}
 -\mathbb{E}[\|\bar{\mathbf{q}}_{\perp \mathcal{S}}\|] \|\boldsymbol{\alpha}\| &\leq \mathbb{E}[\langle \bar{\mathbf{q}}_{\perp \mathcal{S}}, \boldsymbol{\alpha} \rangle] \leq \mathbb{E}[\|\bar{\mathbf{q}}_{\perp \mathcal{S}}\|] \|\boldsymbol{\alpha}\|, \\
 -M_1 \|\boldsymbol{\alpha}\| &\leq \mathbb{E}[\langle \bar{\mathbf{q}}_{\perp \mathcal{S}}, \boldsymbol{\alpha} \rangle] \leq M_1 \|\boldsymbol{\alpha}\|,
 \end{aligned}
 \tag{16}$$

where the last set of inequalities follow from the state space collapse in Proposition 2. Putting this back in (15), we get

$$-\frac{2\epsilon}{n} M_1 \|\boldsymbol{\alpha}\| \leq 2\mathbb{E}[\langle \bar{\mathbf{q}}_{\parallel \mathcal{S}}, (\mathbf{s}(\bar{\mathbf{q}}) - \bar{\mathbf{a}})_{\parallel \mathcal{S}} \rangle] - \frac{2\epsilon}{n} \mathbb{E}[\langle \bar{\mathbf{q}}, \boldsymbol{\alpha} \rangle] \leq \frac{2\epsilon}{n} M_1 \|\boldsymbol{\alpha}\|. \tag{17}$$

We now consider the first term on the RHS of (13). Let $\mathbf{f}_1, \mathbf{f}_2, \dots, \mathbf{f}_L$ be an orthonormal basis of the space $\mathcal{S}_{n_1 n_2}$, where L is the dimension of the space $\mathcal{S}_{n_1 n_2}$. From the definition of the space $\mathcal{S}_{n_1 n_2}$, we know that each of these vectors \mathbf{f}_l can be written as $f_{lij} = v_{li} + \tilde{v}_{lj}$ for some $v_{li}, \tilde{v}_{lj} \in \mathbb{R}$ for all i, j with $v_{li} = 0$ for $i > n_1, \tilde{v}_{lj} = 0$ for $j > n_2$. Then the norm of the projection onto the subspace $\mathcal{S}_{n_1 n_2}$ can be written in terms of the projections onto the basis vectors as follows:

$$\begin{aligned}
 &\mathbb{E}[\|\bar{\mathbf{a}} - \mathbf{s}(\bar{\mathbf{q}})\|_{\mathcal{S}}^2] \\
 &= \mathbb{E}\left[\sum_l \langle \bar{\mathbf{a}} - \mathbf{s}(\bar{\mathbf{q}}), \mathbf{f}_l \rangle^2\right] \\
 &= \sum_l \mathbb{E}\left[\left(\sum_{ij} (\bar{a}_{ij} - s_{ij}(\bar{\mathbf{q}})) f_{lij}\right)^2\right] \\
 &= \sum_l \mathbb{E}\left[\left(\sum_{ij} (\bar{a}_{ij} - s_{ij}(\bar{\mathbf{q}}))(v_{li} + \tilde{v}_{lj})\right)^2\right] \\
 &= \sum_l \mathbb{E}\left[\left(\sum_i v_{li} \left(\sum_j (\bar{a}_{ij} - s_{ij}(\bar{\mathbf{q}}))\right) + \sum_j \tilde{v}_{lj} \left(\sum_i (\bar{a}_{ij} - s_{ij}(\bar{\mathbf{q}}))\right)\right)^2\right] \\
 &\stackrel{(a)}{=} \sum_l \mathbb{E}\left[\left(\sum_i v_{li} \left(\sum_j \bar{a}_{ij} - (1 - \kappa_i \epsilon)\right) + \sum_j \tilde{v}_{lj} \left(\sum_i \bar{a}_{ij} - (1 - \tilde{\kappa}_j \epsilon)\right) - \epsilon \left(\sum_i \kappa_i v_{li} + \sum_j \tilde{\kappa}_j \tilde{v}_{lj}\right)\right)^2\right] \\
 &\stackrel{(b)}{=} \sum_l \text{Var}\left(\sum_i v_{li} \sum_j \bar{a}_{ij} + \sum_j \tilde{v}_{lj} \sum_i \bar{a}_{ij}\right) + \epsilon^2 \sum_l \left(\sum_i \kappa_i v_{li} + \sum_j \tilde{\kappa}_j \tilde{v}_{lj}\right)^2 \\
 &\stackrel{(c)}{=} \sum_l \left[\text{Var}\left(\sum_i v_{li} \sum_j \bar{a}_{ij}\right) + \text{Var}\left(\sum_j \tilde{v}_{lj} \sum_i \bar{a}_{ij}\right) + 2\text{Cov}\left(\sum_i v_{li} \sum_{j'} \bar{a}_{ij'}, \sum_j \tilde{v}_{lj} \sum_{i'} \bar{a}_{i'j}\right)\right] + \epsilon^2 \sum_l \langle \mathbf{k}, \mathbf{f}_l \rangle^2
 \end{aligned}$$

$$\begin{aligned}
 &\stackrel{(d)}{=} \sum_l \left[\sum_i v_l^2 \text{Var} \left(\sum_j \bar{a}_{ij} \right) + \sum_j \tilde{v}_l^2 \text{Var} \left(\sum_i \bar{a}_{ij} \right) \right. \\
 &\quad \left. + 2 \sum_{ij} v_l \tilde{v}_l \text{Cov} \left(\sum_{j'} \bar{a}_{ij'}, \sum_{i'} \bar{a}_{i'j} \right) \right] + \epsilon^2 \|\mathbf{k}_{\parallel S}\|^2 \\
 &\stackrel{(e)}{=} \sum_l \left[\sum_i v_l^2 \sum_j \sigma_{ij}^2 + \sum_j \tilde{v}_l^2 \sum_i \sigma_{ij}^2 + 2 \sum_{ij} v_l \tilde{v}_l \sigma_{ij}^2 \right] + \epsilon^2 \|\mathbf{k}_{\parallel S}\|^2 \\
 &= \sum_l \left[\sum_{ij} (v_l + \tilde{v}_l)^2 \sigma_{ij}^2 \right] + \epsilon^2 \|\mathbf{k}_{\parallel S}\|^2 \\
 &= \sum_{ij} \sum_l f_{lij}^2 \sigma_{ij}^2 + \epsilon^2 \|\mathbf{k}_{\parallel S}\|^2 \\
 &= \sum_{ij} \sum_l \langle \chi^{(ij)}, \mathbf{f}_l \rangle \sigma_{ij}^2 + \epsilon^2 \|\mathbf{k}_{\parallel S}\|^2 \\
 &= \sum_{ij} \|\chi_{\parallel S}^{(ij)}\|^2 \sigma_{ij}^2 + \epsilon^2 \|\mathbf{k}_{\parallel S}\|^2 \\
 &= \frac{1}{n} \langle \sigma^2, \xi \rangle + \epsilon^2 \|\mathbf{k}_{\parallel S}\|^2, \tag{18}
 \end{aligned}$$

where $\chi^{(ij)}$ is the matrix with 1 in the (i, j) th position and 0 everywhere else. Since a maximal schedule is always picked, from (1), we get (a). Equation (b) follows from the fact that the total arrival rate for row i is $(1 - \kappa_i \epsilon)$ and that for column j it is $(1 - \tilde{\kappa}_j \epsilon)$. Expanding the variance of two random variables along with the definition of κ_i and $\tilde{\kappa}_j$ in (6) gives (c). Independence of the arrival processes across the ports gives (d) and (e). The following lemma, which is proved in Appendix 1, gives (18).

Lemma 4 For all $1 \leq i, j \leq n$,

$$\|\chi_{\parallel S}^{(ij)}\|^2 = \frac{\xi_{ij}}{n},$$

where ξ is defined in (11).

We will now focus on the second term on the RHS of (13). In order to do this, we will first obtain a bound on the unused service. We know from (12) that $\mathbb{E}[\sum_{ij} \bar{q}_{ij}]$ is finite and so $\mathbb{E}[\sum_j \bar{q}_{ij}]$ finite for each i . For $i \leq n_1$, setting the drift of $\sum_j \bar{q}_{ij}$ to zero in steady state, we get

$$\begin{aligned}
 \mathbb{E} \left[\sum_j \bar{q}_{ij} \right] &= \mathbb{E} \left[\sum_j \bar{q}_{ij}^+ \right] \\
 &= \mathbb{E} \left[\sum_j \bar{q}_{ij} + \bar{a}_{ij} - s_{ij}(\bar{\mathbf{q}}) + u_{ij}(\bar{\mathbf{q}}) \right],
 \end{aligned}$$

$$0 = \sum_j (v_{ij} - \epsilon k_{ij}) - 1 + \mathbb{E} \left[\sum_j u_{ij}(\bar{\mathbf{q}}) \right],$$

$$\mathbb{E} \left[\sum_j u_{ij}(\bar{\mathbf{q}}) \right] = \epsilon \kappa_i,$$

where the last equality follows from (5) since the row i is saturating. Similarly, for $j \leq n_2$, we have that $\mathbb{E} [\sum_i u_{ij}(\bar{\mathbf{q}})] = \epsilon \tilde{\kappa}_j$.

For any $\mathbf{x} \in \mathbb{R}^{n^2}$, let $\hat{\mathbf{x}} \in \mathbb{R}^{n^2}$ denote its projection on to the space spanned by the vectors $\chi^{(ij)}$ for $i \leq n_1$ or $j \leq n_2$. Call this space $\mathcal{X}_{n_1 n_2}$. Clearly, this space contains the space $\mathcal{S}_{n_1 n_2}$, i.e., $\mathcal{S}_{n_1 n_2} \subseteq \mathcal{X}_{n_1 n_2}$. In other words, the vector $\hat{\mathbf{x}} \in \mathbb{R}^{n^2}$ is obtained by replacing the $(n - n_1)(n - n_2)$ components with $i > n_1$ and $j > n_2$ with zeros, i.e.,

$$\hat{x}_{ij} \triangleq \begin{cases} x_{ij} & \text{if } i \leq n_1 \text{ or } j \leq n_2 \\ 0 & \text{if } i > n_1 \text{ and } j > n_2. \end{cases}$$

Moreover, for any $\mathbf{x} \in \mathbb{R}^{n^2}$, $\mathbf{x} - \hat{\mathbf{x}}$ is orthogonal to the space $\mathcal{X}_{n_1 n_2}$ and so is also orthogonal to the space $\mathcal{S}_{n_1 n_2}$.

Now the second term on the RHS of (13) can be upper-bounded as follows:

$$\begin{aligned} \mathbb{E} \left[\|\mathbf{u}_{\parallel \mathcal{S}}(\bar{\mathbf{q}})\|^2 \right] &= \mathbb{E} \left[\|(\hat{\mathbf{u}}(\bar{\mathbf{q}}) + \mathbf{u}(\bar{\mathbf{q}}) - \hat{\mathbf{u}}(\bar{\mathbf{q}}))_{\parallel \mathcal{S}}\|^2 \right] \\ &\stackrel{(a)}{=} \mathbb{E} \left[\|(\hat{\mathbf{u}}(\bar{\mathbf{q}}))_{\parallel \mathcal{S}} + (\mathbf{u}(\bar{\mathbf{q}}) - \hat{\mathbf{u}}(\bar{\mathbf{q}}))_{\parallel \mathcal{S}}\|^2 \right] \\ &\stackrel{(b)}{=} \mathbb{E} \left[\|(\hat{\mathbf{u}}(\bar{\mathbf{q}}))_{\parallel \mathcal{S}}\|^2 \right] \\ &\stackrel{(c)}{\leq} \mathbb{E} \left[\|(\hat{\mathbf{u}}(\bar{\mathbf{q}}))\|^2 \right] \\ &= \mathbb{E} \left[\sum_{ij} \hat{u}_{ij}^2(\bar{\mathbf{q}}) \right] \\ &\stackrel{(d)}{=} \mathbb{E} \left[\sum_{ij} \hat{u}_{ij}(\bar{\mathbf{q}}) \right] \\ &\leq \mathbb{E} \left[\sum_{i=1}^{n_1} \sum_j u_{ij}(\bar{\mathbf{q}}) \right] + \mathbb{E} \left[\sum_{j=1}^{n_2} \sum_i u_{ij}(\bar{\mathbf{q}}) \right] \\ &\stackrel{(e)}{=} \epsilon \sum_{i=1}^{n_1} \kappa_i + \epsilon \sum_{j=1}^{n_2} \tilde{\kappa}_j \\ &\leq 2\epsilon \langle \mathbf{k}, \mathbf{1} \rangle \\ &= 2\epsilon n, \end{aligned} \tag{19}$$

where (a) follows from linearity of projection onto a subspace. Since, for any vector \mathbf{x} , $\mathbf{x} - \widehat{\mathbf{x}}$ is orthogonal to the space $\mathcal{S}_{n_1 n_2}$, we get (b). Inequality (c) is true due to the nonexpansive property of projection. Since, $u_{ij} \in \{0, 1\}$, we have (d). Since the saturation rate vector \mathbf{k} is assumed to be normalized, we get (19). Using the trivial lower bound of zero, we have

$$0 \leq \mathbb{E} \left[\|\mathbf{u}_{\parallel \mathcal{S}}(\overline{\mathbf{q}})\|^2 \right] \leq 2\epsilon n. \tag{20}$$

We will now consider the final term, the one in (14).

$$\begin{aligned} 2\mathbb{E} \left[\langle \overline{\mathbf{q}}_{\parallel \mathcal{S}}^+, \mathbf{u}_{\parallel \mathcal{S}}(\overline{\mathbf{q}}) \rangle \right] &= 2\mathbb{E} \left[\langle \overline{\mathbf{q}}_{\parallel \mathcal{S}}^+, \mathbf{u}(\overline{\mathbf{q}}) \rangle \right] \\ &\stackrel{(a)}{=} 2\mathbb{E} \left[\langle \overline{\mathbf{q}}_{\parallel \mathcal{S}}^+, \widehat{\mathbf{u}}(\overline{\mathbf{q}}) \rangle \right] \\ &\stackrel{(b)}{=} 2\mathbb{E} \left[\langle \overline{\mathbf{q}}^+, \widehat{\mathbf{u}}(\overline{\mathbf{q}}) \rangle \right] - 2\mathbb{E} \left[\langle \overline{\mathbf{q}}_{\perp \mathcal{S}}^+, \widehat{\mathbf{u}}(\overline{\mathbf{q}}) \rangle \right] \\ &= -2\mathbb{E} \left[\langle \overline{\mathbf{q}}_{\perp \mathcal{S}}^+, \widehat{\mathbf{u}}(\overline{\mathbf{q}}) \rangle \right], \end{aligned} \tag{21}$$

where (a) follows from the fact that $\mathbf{x} - \widehat{\mathbf{x}}$ is orthogonal to the space $\mathcal{S}_{n_1 n_2}$. Since, by the definition of unused service in (2), we have that $\overline{q}_{ij}^+ = 0$ whenever $u_{ij}(\overline{\mathbf{q}}) > 0$, the first term in (b) vanishes. Therefore, using the Cauchy–Schwartz inequality, we get

$$\begin{aligned} -2\sqrt{\mathbb{E} \left[\|\overline{\mathbf{q}}_{\perp \mathcal{S}}^+\|^2 \right] \mathbb{E} \left[\|\widehat{\mathbf{u}}(\overline{\mathbf{q}})\|^2 \right]} &\leq 2\mathbb{E} \left[\langle \overline{\mathbf{q}}_{\perp \mathcal{S}}^+, \mathbf{u}_{\parallel \mathcal{S}}(\overline{\mathbf{q}}) \rangle \right] \leq 2\sqrt{\mathbb{E} \left[\|\overline{\mathbf{q}}_{\perp \mathcal{S}}^+\|^2 \right] \mathbb{E} \left[\|\widehat{\mathbf{u}}(\overline{\mathbf{q}})\|^2 \right]}, \\ -2M_2\sqrt{\mathbb{E} \left[\|\widehat{\mathbf{u}}(\overline{\mathbf{q}})\|^2 \right]} &\leq 2\mathbb{E} \left[\langle \overline{\mathbf{q}}_{\parallel \mathcal{S}}^+, \mathbf{u}_{\parallel \mathcal{S}}(\overline{\mathbf{q}}) \rangle \right] \leq 2M_2\sqrt{\mathbb{E} \left[\|\widehat{\mathbf{u}}(\overline{\mathbf{q}})\|^2 \right]}, \end{aligned} \tag{22}$$

$$-2M_2\sqrt{2\epsilon n} \leq 2\mathbb{E} \left[\langle \overline{\mathbf{q}}_{\parallel \mathcal{S}}^+, \mathbf{u}_{\parallel \mathcal{S}}(\overline{\mathbf{q}}) \rangle \right] \leq 2M_2\sqrt{2\epsilon n}, \tag{23}$$

where (22) is obtained by using the fact that, in steady state, $\mathbb{E} \left[\|\overline{\mathbf{q}}_{\perp \mathcal{S}}^+\|^2 \right] = \mathbb{E} \left[\|\overline{\mathbf{q}}_{\perp \mathcal{S}}\|^2 \right]$, which is bounded by M_2 from state space collapse in Proposition 2. We get (23) from the bound in (19).

Substituting (17), (18), (20) and (23) in (13) and (14), we get the theorem with

$$B_1(\epsilon) = M_1 \|\boldsymbol{\alpha}\| - \frac{n\epsilon}{2} \|\mathbf{k}_{\parallel \mathcal{S}}\|^2 + n^2 + 2M_2 \frac{n\sqrt{n}}{\sqrt{2\epsilon}},$$

$$B_2(\epsilon) = M_1 \|\boldsymbol{\alpha}\| + \frac{n\epsilon}{2} \|\mathbf{k}_{\parallel \mathcal{S}}\|^2 + 2M_2 \frac{n\sqrt{n}}{\sqrt{2\epsilon}}.$$

□

The main idea of the proof is to set the drift of a carefully chosen test function $V(\cdot)$ to zero. The choice of this function is crucial to obtain tight heavy-traffic bounds. We will now briefly motivate our choice of the function $V(\cdot)$. For a discrete-time single-server (G/G/1) queue, with $q(t)$ that evolves according to $q(t+1) = q(t) + a(t) - s(t) + u(t)$, the right test function to obtain tight queue length bounds is q^2 . Such a bound is known

as the Kingman bound [20, Section 10.1]. Next, consider a load balancing system under the ‘Join the shortest queue’ (JSQ) policy operating in discrete time. There are a finite number of servers, each with a separate queue, similar to supermarket checkout lanes. Whenever a user arrives into the system, (s)he joins the queue with the shortest length, breaking ties uniformly at random. Tight heavy-traffic queue length bounds are obtained for this system in [6] by first showing that the queue lengths collapse to a single dimension where they are all equal. Then, tight bounds are obtained by setting the drift of the quadratic test function $\mathbb{E}[(\sum_i \bar{q}_i)^2]$ to zero in steady state. This function is the same as $\|\mathbf{q}_{\parallel}\|^2$ (up to a factor of n , which is not important), where \mathbf{q}_{\parallel} denotes the projection of the queue length vector \mathbf{q} onto the region of state space collapse, which is the line along the vector that has ones in all components.

These examples motivate us to choose the norm square of projection of the queue lengths vector onto the region of state space collapse as the test function in general. In the case of the switch, this would be $\|\mathbf{q}_{\parallel\mathcal{K}}\|^2$. However, a projection operator onto a convex cone is difficult to study because it is not linear. Moreover, a closed form expression for projection onto a general cone is not known. Therefore, we relax this function and use $\|\mathbf{q}_{\parallel\mathcal{S}}\|^2$ as the test function. Since we use the relaxed function, it is sufficient to use state space collapse into the space $\mathcal{S}_{n_1n_2}$, as opposed to the stronger form of state space collapse into the cone $\mathcal{K}_{n_1n_2}$. Note that, in the proof of Theorem 1, we use state space collapse only at two instances, viz., (16) and (22), and both these use the weaker result, $\mathbb{E}[\|\bar{\mathbf{q}}_{\perp\mathcal{S}}^{(\epsilon)}\|^r]$. In other words, we only use the fact that the state collapses to the space $\mathcal{S}_{n_1n_2}$ and not the cone $\mathcal{K}_{n_1n_2}$. Such a relaxation works because of the property of the cone $\mathcal{K}_{n_1n_2}$ that it is just the intersection of the space it spans, $\mathcal{S}_{n_1n_2}$, with the positive orthant $\mathbb{R}_+^{n_2}$, as proved in Lemma 3. Therefore, any positive queue length vector that is in the space $\mathcal{S}_{n_1n_2}$ is also in the cone $\mathcal{K}_{n_1n_2}$. Since queue length vectors are nonnegative, if we know that a queue length vector collapses onto the space $\mathcal{S}_{n_1n_2}$, we know that it should collapse onto the cone $\mathcal{K}_{n_1n_2}$.

Even though we only need the weaker version of state space collapse, viz., the bound on $\mathbb{E}[\|\bar{\mathbf{q}}_{\perp\mathcal{S}}^{(\epsilon)}\|^r]$, we proved a stronger version, viz., the bound on $\mathbb{E}[\|\bar{\mathbf{q}}_{\perp\mathcal{K}}^{(\epsilon)}\|^r]$ in Proposition 2 for completeness. The proof of weaker state space collapse can be much simpler because projection onto a subspace is a linear operator, while projection onto a convex cone is not. Therefore, we can write down the complete proof of Theorem 1 (including the weaker version of Proposition 2) without even referring to the cone $\mathcal{K}_{n_1n_2}$.

The intuition behind our choice of the test function as discussed above may be applicable to other problem settings beyond just the switch system. For any problem that exhibits state space collapse into a cone that satisfies a property similar to Lemma 3, we may be able to use $\|\mathbf{q}_{\parallel\mathcal{S}}\|^2$ as the test function with appropriately defined inner product and norm. Moreover, as illustrated in the proof of Theorem 1, we don’t need to know the explicit closed form of $\mathbf{q}_{\parallel\mathcal{S}}$ to use this test function. Our presentation of the test function in the form $\|\mathbf{q}_{\parallel\mathcal{S}}\|^2$ that is easily generalizable to other problem settings is a major contribution of this paper. Such a test function with a novel inner product has been successfully used in [9, 23]. The reason for the choice of $\|\mathbf{q}_{\parallel\mathcal{S}}\|^2$ as a test function is the following: When a quadratic test function is chosen, and the

drift is set to zero, we get several terms as in (14). Among these terms, the cross term between $\bar{\mathbf{q}}^+$ and $\mathbf{u}(\bar{\mathbf{q}})$, as in (14), is the crucial term to bound in heavy-traffic. Under state space collapse, this term should be small and can be bounded using an argument similar to that in the set of equations ending in (21).

We now make a short remark about the scheduling policy. In the proof of Theorem 1, we do not use any details about the scheduling policy, except for the fact that it is throughput-optimal and exhibits state space collapse as in Proposition 2. Therefore, the queue length bounds in Theorem 1 hold true under any throughput-optimal algorithm that exhibits state space collapse as in Proposition 2. Moreover, in the proof of state space collapse in Proposition 2, we use the details of the MaxWeight scheduling policy only in the Claim 1 in Sect. 3.2.

4 Discussion

In this section, we present various corollaries and extensions of Theorem 1, and interpret the results. The following corollary gives a bound on the sum of the queue lengths.

Corollary 1 *Consider the set of switch systems operating under the MaxWeight algorithm as described in Theorem 1, with $0 \leq n_1, n_2 < n$. Then, in the heavy-traffic limit, we have*

$$\frac{1}{2\kappa_{\max}} \langle \sigma^2, \zeta \rangle \leq \lim_{\epsilon \rightarrow 0} \epsilon \mathbb{E} \left[\sum_{ij} \bar{q}_{ij}^{(\epsilon)} \right] \leq \frac{1}{2\kappa_{\min}} \langle \sigma^2, \zeta \rangle,$$

where $\kappa_{\max} = \max_{i \leq n_1, j \leq n_2} \{\kappa_i, \tilde{\kappa}_j\}$ and $\kappa_{\min} = \min_{i \leq n_1, j \leq n_2} \{\kappa_i, \tilde{\kappa}_j\}$. Moreover, under any stable algorithm the sum of all the queue lengths is lower-bounded by

$$\lim_{\epsilon \rightarrow 0} \epsilon \mathbb{E} \left[\sum_{ij} \bar{q}_{ij}^{(\epsilon)} \right] \geq \frac{1}{2} \max \left\{ \langle \sigma^2, \zeta' \rangle, \langle \sigma^2, \zeta'' \rangle \right\},$$

where $\zeta'_{ij} = 1/\kappa_i$ and $\zeta''_{ij} = 1/\tilde{\kappa}_j$.

Proof It is easy to see that there exists a weight vector $\alpha \in \mathbb{R}_+^{n^2}$ such that $\kappa_{\min} \leq \alpha_{ij}$ for all i, j , $\langle \alpha, \mathbf{e}^{(i)} \rangle = n\kappa_i$ for $i \leq n_1$ and $\langle \alpha, \tilde{\mathbf{e}}^{(j)} \rangle = n\tilde{\kappa}_j$ for $j \leq n_2$. Using such a weight vector in Theorem 1, we get the upper bound. Similarly, by picking an α such that $\kappa_{\max} \geq \alpha_{ij}$ for all i, j we get the lower bound. \square

If the saturation rate vector \mathbf{k} is such that $\kappa_{\max}/\kappa_{\min}$ is a constant that does not scale with n , it is easy to see that the sum of all the queue lengths under MaxWeight are within a constant factor of the universal lower bound after noting that $\zeta_{ij} \leq 2$ for all i, j .

For some values of \mathbf{k} , we get an exact expression for the sum of queue lengths in heavy-traffic.

Corollary 2 *Suppose that $\mathbf{k} = \mathbf{v}$ in the set of switch systems described in Theorem 1, so that $\lambda^\epsilon = \mathbf{v}(1 - \epsilon)$ for some $\mathbf{v} \in \text{Relint}(\mathcal{F}_{n_1 n_2})$ such that $v_{\min} \triangleq \min_{ij} v_{ij} > 0$ and $0 \leq n_1, n_2 < n$. Then, in heavy-traffic we have*

$$\lim_{\epsilon \rightarrow 0} \epsilon \mathbb{E} \left[\sum_{ij} \bar{q}_{ij}^{(\epsilon)} \right] = \frac{1}{2} \langle \boldsymbol{\sigma}, \boldsymbol{\zeta} \rangle \text{ and}$$

$$\lim_{\epsilon \rightarrow 0} \epsilon \mathbb{E} \left[\bar{q}_{ij}^{(\epsilon)} \right] = 0 \quad \forall i > n_1, j > n_2.$$

Proof Since $\kappa_i = 1$ and $\tilde{\kappa}_j = 1$ for $i \leq n_1$ and $j \leq n_2$, we pick $\alpha_{ij} = 1$ for all i, j . Since the weight vector $\boldsymbol{\alpha}$ satisfies the condition in Theorem 1, the corollary follows. □

The following corollary considers the case when exactly one port of the switch is saturated. Under this condition, the switch is said to satisfy the complete resource pooling condition and was studied in [21]. In this case, the state collapses onto a line. Without loss of generality, the corollary is stated when an input port is saturated.

Corollary 3 *Consider the set of switch systems operating under the MaxWeight algorithm as described in Theorem 1, with $n_1 = 1, n_2 = 0$ and $\mathbf{k} = \mathbf{v}$ so that $\lambda^\epsilon = \mathbf{v}(1 - \epsilon)$ for some $\mathbf{v} \in \text{Relint}(\mathcal{F}_{10})$ such that $v_{\min} \triangleq \min_{ij} v_{ij} > 0$. Then, in heavy-traffic we have*

$$\lim_{\epsilon \rightarrow 0} \epsilon \mathbb{E} \left[\sum_j \bar{q}_{1j}^{(\epsilon)} \right] = \frac{\sum_{1j} \sigma_{1j}^2}{2} \text{ and}$$

$$\lim_{\epsilon \rightarrow 0} \epsilon \mathbb{E} \left[\bar{q}_{ij}^{(\epsilon)} \right] = 0 \quad \forall i > 1, j.$$

Moreover, MaxWeight is heavy-traffic optimal, i.e., $\lim_{\epsilon \rightarrow 0} \epsilon \mathbb{E} \left[\sum_{ij} \bar{q}_{ij}^{(\epsilon)} \right]$ is minimized under the MaxWeight algorithm.

The proof follows directly from Theorem 1 and the universal lower bound in Proposition 1. In order to clearly see the scaling of queue lengths in terms of n , we state the following corollary under Bernoulli arrivals, which again follows from Corollary 2 and Proposition 1.

Corollary 4 *Consider the set of switch systems operating under the MaxWeight algorithm, as described in Theorem 1. Suppose that the arrival process for each queue is Bernoulli with the arrival rate vector $\lambda^{(\epsilon)}$, where $\lambda_{ij}^{(\epsilon)} = (1 - \epsilon)/2n$ for $i > n_1, j > n_2$ and $\lambda_{ij}^{(\epsilon)} = (1 - \epsilon)/n$ otherwise, with $0 \leq n_1, n_2 < n$ so that n_1 inputs and n_2 outputs are saturating. Then, in the heavy-traffic limit, we have*

$$\lim_{\epsilon \rightarrow 0} \epsilon \mathbb{E} \left[\sum_{ij} \bar{q}_{ij}^{(\epsilon)} \right] = \frac{n_1 + n_2}{2} \left(1 - \frac{1}{n} \right).$$

Moreover, under any stable algorithm the sum of all the queue lengths is lower-bounded by

$$\lim_{\epsilon \rightarrow 0} \epsilon \mathbb{E} \left[\sum_{ij} \bar{q}_{ij}^{(\epsilon)} \right] \geq \frac{\max\{n_1, n_2\}}{2} \left(1 - \frac{1}{n} \right).$$

So, we know that MaxWeight is within less than a factor of two away from heavy-traffic optimality under incomplete saturation. A similar observation was made under the completely saturated case in [11]. It is not clear if the gap is because the lower bound is loose or because MaxWeight is a constant factor away from heavy-traffic optimality. Under the MaxWeight algorithm, the sum of all queue lengths is as if we have $(n_1 + n_2)$ separate queues, each with variance $(1 - \frac{1}{n})$, which is the total variance of the arrivals in each row or column. This is because there are $(n_1 + n_2)$ independent constraints on the capacity region that are tight in the limit. This is the same reason why the state collapses to the $(n_1 + n_2)$ -dimensional space $\mathcal{S}_{n_1 n_2}$. So, in general the number of tight constraints in the limit is important.

Theorem 1 is valid only for the incompletely saturated switch, $n_1 < n$ and $n_2 < n$. However, a similar result can be proved for the completely saturated case $n_1 = n_2 = n$ as follows.

Corollary 5 Consider the set of switch systems operating under the MaxWeight algorithm, parameterized by $0 < \epsilon < 1$, as described in Theorem 1, with the only difference being that, $n_1 = n_2 = n$. Then, as long as $0 < \epsilon \leq v'_{\min}/2\|\mathbf{k}\|$, the steady state queue lengths vector satisfies

$$\left(1 - \frac{1}{2n} \right) \frac{\|\boldsymbol{\sigma}^{(\epsilon)}\|^2}{\epsilon} - B_3(\epsilon, n) \leq \mathbb{E} \left[\langle \bar{\mathbf{q}}^{(\epsilon)}, \boldsymbol{\alpha} \rangle \right] \leq \left(1 - \frac{1}{2n} \right) \|\boldsymbol{\sigma}^{(\epsilon)}\|^2 + B_4(\epsilon, n)$$

for any fixed weight vector $\boldsymbol{\alpha} \in \mathbb{R}^{n^2}$ such that $\langle \boldsymbol{\alpha}, \mathbf{e}^{(i)} \rangle = n\kappa_i$ and $\langle \boldsymbol{\alpha}, \tilde{\mathbf{e}}^{(j)} \rangle = n\tilde{\kappa}_j$ for all $1 \leq i, j \leq n$, where $B_3(\epsilon)$ and $B_4(\epsilon)$ are $o(\frac{1}{\epsilon})$. Thus, in the heavy-traffic limit as $\epsilon \downarrow 0$, we have

$$\lim_{\epsilon \rightarrow 0} \epsilon \mathbb{E} \left[\langle \bar{\mathbf{q}}^{(\epsilon)}, \boldsymbol{\alpha} \rangle \right] = \left(1 - \frac{1}{2n} \right) \|\boldsymbol{\sigma}\|^2.$$

Proof Note that Proposition 2 is valid in the case when $n_1 = n_2 = n$. Most of the proof of Theorem 1 also holds true, except for Lemma 4. The norm of the projections of unit vectors $\boldsymbol{\chi}^{(ij)}$ onto the cone \mathcal{S}_{nn} is different from ζ_{ij} . This fundamental difference in behavior for the case $n_1 = n_2 = n$ is for the following reason: For $n_1, n_2 < n$, the cone $\mathcal{S}_{n_1 n_2}$ is spanned by the vectors $\mathbf{e}^{(i)}, \tilde{\mathbf{e}}^{(j)}$ for $i \leq n_1, j \leq n_2$, which are linearly independent, and so the cone $\mathcal{S}_{n_1 n_2}$ has dimension $(n_1 + n_2)$. When $n_1 = n_2 = n$, the vectors $\mathbf{e}^{(i)}, \tilde{\mathbf{e}}^{(j)}$ for $i \leq n, j \leq n$ are not linearly independent because, clearly, $\sum_i \mathbf{e}^{(i)} = \sum_j \tilde{\mathbf{e}}^{(j)}$, and so the dimension of the cone \mathcal{S}_{nn} is smaller than $2n$. It can be shown that the cone \mathcal{S}_{nn} has dimension $(2n - 1)$ [25, page 20]. The proof will

be complete once we calculate $\|\chi_{\|S}^{(ij)}\|^2$ for all i, j . By symmetry, we have that these norms have the same value for all i, j , i.e., $\|\chi_{\|S}^{(ij)}\|^2 = \xi$ for some ξ , which can be calculated as follows: Suppose $\mathbf{f}_1, \mathbf{f}_2, \dots, \mathbf{f}_{2n-1}$ is an orthonormal basis of S_{nn} . We have

$$\begin{aligned} n^2\xi &= \sum_{ij} \|\chi_{\|S}^{(ij)}\|^2 \\ &= \sum_{ij} \sum_{l=1}^{2n-1} \langle \chi^{(ij)}, \mathbf{f}_l \rangle^2 \\ &= \sum_{l=1}^{2n-1} \sum_{ij} \langle \chi^{(ij)}, \mathbf{f}_l \rangle^2 \\ &\stackrel{(a)}{=} \sum_l^{2n-1} \|\mathbf{f}_l\|_{\|S}^2 \\ &= 2n - 1, \end{aligned}$$

where (a) follows from the fact that $\{\chi^{(ij)}\}_{ij}$ is an orthonormal basis of $\mathbb{R}^{(n^2)}$. Replacing ζ_{ij} by $\xi = (2n - 1)/n^2$, we get the corollary. \square

Similar to Corollary 1, we can get lower and upper bounds on the sum of queue lengths in heavy-traffic using κ_{\max} and κ_{\min} . We now state the following corollary to illustrate the use of the weight vectors α .

Corollary 6 *Consider the completely saturated switch system in Corollary 5 with $\mathbf{k} = \mathbf{v}$. Then, in the heavy-traffic limit the queue lengths satisfy the following relations:*

$$\begin{aligned} \lim_{\epsilon \rightarrow 0} \epsilon \mathbb{E} \left[\sum_{ij} \bar{q}_{ij}^{(\epsilon)} \right] &= \left(1 - \frac{1}{2n} \right) \|\sigma\|^2, \\ \lim_{\epsilon \rightarrow 0} \epsilon \mathbb{E} \left[\sum_{ij} \bar{q}_{ij} v_{ij}^{(\epsilon)} \right] &= \left(\frac{2n - 1}{2n^2} \right) \|\sigma\|^2, \\ \lim_{\epsilon \rightarrow 0} \epsilon \mathbb{E} \left[\sum_{ij} \bar{q}_{i\pi(i)}^{(\epsilon)} \right] &= \left(\frac{2n - 1}{2n^2} \right) \|\sigma\|^2 \text{ for any permutation } \pi. \end{aligned}$$

Proof The proof follows directly from Corollary 5 by choosing $\alpha = \mathbf{1}$, $\alpha = n\mathbf{v}$ and $\alpha = nP_\pi$, respectively, where P_π is the permutation matrix corresponding to the permutation π . \square

Note that the first result above is the main result of [11]. Even though the proof of Theorem 1 in [11] is written in a different style, it is equivalent to the proof presented

here, because the function used there is equivalent to the norm $\|\mathbf{q}\|_{\mathcal{S}}\|^2$. It can be shown (using the orthogonality principle) that, for any vector $\mathbf{q} \in \mathbb{R}^{n^2}$, its projection onto the space \mathcal{S}_{nn} is given by

$$q_{\|\mathcal{S}ij} = \frac{\sum_i q_{ij}}{n} + \frac{\sum_j q_{ij}}{n} - \frac{\sum_{ij} q_{ij}}{n^2}.$$

Taking the norm, we get

$$\|\mathbf{q}\|_{\mathcal{S}}\|^2 = \frac{1}{n} \left(\sum_i \left(\sum_j q_{ij} \right)^2 + \sum_j \left(\sum_i q_{ij} \right)^2 - \frac{1}{n} \left(\sum_{ij} q_{ij} \right)^2 \right),$$

which is the function used in [11] scaled by n . We now present a further special case, which was also studied in [11], to contrast it with the result in Corollary 4.

Corollary 7 *Suppose that the traffic of the completely saturated systems in Corollary 5 is uniform Bernoulli with uniform saturation rate, i.e., $\lambda^{(\epsilon)} = (1 - \epsilon)/n\mathbf{1}$. In heavy-traffic we have*

$$\lim_{\epsilon \rightarrow 0} \epsilon \mathbb{E} \left[\sum_{ij} \bar{q}_{ij}^{(\epsilon)} \right] = \frac{2n - 1}{2} \left(1 - \frac{1}{n} \right).$$

Notice that the behavior of the queue length here is similar to $(2n - 1)$ separate queues. This is consistent with the discussion after Corollary 4 because in the heavy-traffic limit, among the $2n$ constraints in the capacity region, we have only $(2n - 1)$ linearly independent ones.

5 Conclusion

We consider the heavy-traffic queue length behavior in an input-queued switch operating under the MaxWeight algorithm. It was recently shown in [11] that, in the heavy-traffic regime, the queue length scales optimally with the size of the switch when all the ports in the switch saturate at the same rate. In this paper, we considered the case when an arbitrary set of ports saturate, and each port is allowed to saturate at different rate. We obtained an exact heavy-traffic characterization of a linear combination of queue lengths and showed that the MaxWeight algorithm achieves optimal scaling of the sum of all the queue lengths in heavy-traffic.

Acknowledgements We thank Prof. Alexander Stolyar and Prof. Michael Harrison for suggesting us the question about heavy-traffic queue length behavior in a switch when different ports saturate at different rates. This research was supported by NSF Grants CMMI-1562276, CIF-1409106, and ARO Grant W911NF-16-1-0259.

Appendix A: Proof of Lemma 4

Proof In order to calculate the norm of the projections of the unit vectors $\chi^{(ij)}$, we will first consider the projections of an arbitrary vector \mathbf{x} on to the space $\mathcal{S}_{n_1 n_2}$. From the definition of the cone $\mathcal{S}_{n_1 n_2}$, we know that the projection $\mathbf{x}_{\parallel \mathcal{S}}$ of \mathbf{x} can be decomposed as

$$\mathbf{x}_{\parallel \mathcal{S}} = \sum_{i=1}^{n_1} w_i \mathbf{e}^{(i)} + \sum_{j=1}^{n_2} \tilde{w}_j \tilde{\mathbf{e}}^{(j)}.$$

From the Orthogonality Principle, we have that

$$\begin{aligned} \langle \mathbf{x} - \mathbf{x}_{\parallel \mathcal{S}}, \mathbf{e}^{(i)} \rangle &= 0 \quad \text{for } i \leq n_1, \\ \langle \mathbf{x} - \mathbf{x}_{\parallel \mathcal{S}}, \tilde{\mathbf{e}}^{(j)} \rangle &= 0 \quad \text{for } j \leq n_2. \end{aligned}$$

Note that

$$\begin{aligned} \langle \mathbf{e}^{(i)}, \mathbf{e}^{(k)} \rangle &= \begin{cases} 0 & i \neq k \\ n & i = k \end{cases}, & \langle \tilde{\mathbf{e}}^{(j)}, \tilde{\mathbf{e}}^{(l)} \rangle &= \begin{cases} 0 & j \neq l \\ n & j = l \end{cases}, \\ \langle \mathbf{e}^{(i)}, \tilde{\mathbf{e}}^{(j)} \rangle &= 1 \quad \forall i, j. \end{aligned}$$

Therefore, we get

$$nw_i + \sum_{j=1}^{n_2} \tilde{w}_j = \sum_j x_{ij} \text{ for } i \leq n_1 \quad \text{and} \quad n\tilde{w}_j + \sum_{i=1}^{n_1} w_i = \sum_i x_{ij} \text{ for } i \leq n_2.$$

Defining $W = \sum_{i=1}^{n_1} w_i$ and $\tilde{W} = \sum_{j=1}^{n_2} \tilde{w}_j$, and summing each set of equations above, we get

$$nW + n_1 \tilde{W} = \sum_{i=1}^{n_1} \sum_j x_{ij} \quad \text{and} \quad n\tilde{W} + n_2 W = \sum_{j=1}^{n_2} \sum_i x_{ij}.$$

Solving for W and \tilde{W} we get

$$\begin{aligned} W &= \frac{n \sum_{i=1}^{n_1} \sum_j x_{ij} - n_1 \sum_{j=1}^{n_2} \sum_i x_{ij}}{n^2 - n_1 n_2} \quad \text{and} \\ \tilde{W} &= \frac{n \sum_{j=1}^{n_2} \sum_i x_{ij} - n_2 \sum_{i=1}^{n_1} \sum_j x_{ij}}{n^2 - n_1 n_2}. \end{aligned} \tag{24}$$

Therefore, w_i and \tilde{w}_j can be written as

$$w_i = \frac{\sum_j x_{ij} - \tilde{W}}{n} \text{ for } i \leq n_1 \text{ and } \tilde{w}_j = \frac{\sum_i x_{ij} - W}{n} \text{ for } i \leq n_2. \tag{25}$$

The norm of the projection $\mathbf{x}_{\parallel S}$ is then given by

$$\begin{aligned} \|\mathbf{x}_{\parallel S}\|^2 &= \langle \mathbf{x}_{\parallel S}, \mathbf{x}_{\parallel S} \rangle \\ &= \left\langle \sum_{i=1}^{n_1} w_i \mathbf{e}^{(i)} + \sum_{j=1}^{n_2} \tilde{w}_j \tilde{\mathbf{e}}^{(j)}, \sum_{i=1}^{n_1} w_i \mathbf{e}^{(i)} + \sum_{j=1}^{n_2} \tilde{w}_j \tilde{\mathbf{e}}^{(j)} \right\rangle \\ &= n \sum_{i=1}^{n_1} w_i^2 + \sum_{i=1}^{n_1} w_i \sum_{j=1}^{n_2} \tilde{w}_j + \sum_{j=1}^{n_2} \tilde{w}_j \sum_{i=1}^{n_1} w_i + n \sum_{j=1}^{n_2} \tilde{w}_j^2 \\ &= n \sum_{i=1}^{n_1} w_i^2 + n \sum_{j=1}^{n_2} \tilde{w}_j^2 + 2W\tilde{W}. \end{aligned} \tag{26}$$

We will now consider the four cases in the definition of ξ in (11) and calculate $\|\chi_{\parallel S}^{(ij)}\|^2$.

Case 1: When $1 \leq i \leq n_1$ and $1 \leq j \leq n_2$

Using (24) and (25), it is easy to see that, when $\mathbf{x} = \chi^{(ij)}$,

$$\begin{aligned} W &= \frac{n - n_1}{n^2 - n_1 n_2}, & \tilde{W} &= \frac{n - n_2}{n^2 - n_1 n_2}, \\ w_k &= \begin{cases} \frac{1 - \tilde{W}}{n} & k = i \\ -\frac{\tilde{W}}{n} & k \neq i \end{cases}, & \tilde{w}_l &= \begin{cases} \frac{1 - W}{n} & l = j \\ -\frac{W}{n} & l \neq j \end{cases}, \end{aligned}$$

$$\begin{aligned} \|\chi_{\parallel S}^{(ij)}\|^2 &= n \sum_{k=1}^{n_1} w_k^2 + n \sum_{l=1}^{n_2} \tilde{w}_l^2 + 2W\tilde{W} \\ &= n \left[(n_1 - 1) \left(\frac{-\tilde{W}}{n} \right)^2 + \left(\frac{1 - \tilde{W}}{n} \right)^2 + (n_2 - 1) \left(\frac{-W}{n} \right)^2 + \left(\frac{1 - W}{n} \right)^2 \right] \\ &\quad + 2 \frac{(n - n_1)(n - n_2)}{(n^2 - n_1 n_2)^2} \\ &= \frac{1}{n} \left[\frac{n_1(n - n_2)^2}{(n^2 - n_1 n_2)^2} - \frac{2(n - n_2)}{(n^2 - n_1 n_2)} + \frac{n_2(n - n_1)^2}{(n^2 - n_1 n_2)^2} - \frac{2(n - n_1)}{(n^2 - n_1 n_2)} + 2 \right] \\ &\quad + 2 \frac{(n - n_1)(n - n_2)}{(n^2 - n_1 n_2)^2} \\ &= \frac{2}{n} - \frac{2n - n_1 - n_2}{n(n^2 - n_1 n_2)}, \end{aligned}$$

where the last equality is obtained after direct algebraic simplification.

Case 2: When $1 \leq i \leq n_1$ and $j > n_2$

Again, using (24) and (25), it is easy to see that, when $\mathbf{x} = \chi^{(ij)}$,

$$W = \frac{n}{n^2 - n_1 n_2}, \quad \tilde{W} = \frac{-n_2}{n^2 - n_1 n_2},$$

$$w_k = \begin{cases} \frac{1-\tilde{W}}{n} & k = i \\ -\frac{\tilde{W}}{n} & k \neq i \end{cases}, \quad \tilde{w}_l = \frac{-W}{n},$$

$$\begin{aligned} \|\chi_{\|S}^{(ij)}\|^2 &= n \sum_{k=1}^{n_1} w_k^2 + n \sum_{l=1}^{n_2} \tilde{w}_l^2 + 2W\tilde{W} \\ &= n \left[(n_1 - 1) \left(\frac{-\tilde{W}}{n}\right)^2 + \left(\frac{1-\tilde{W}}{n}\right)^2 + n_2 \left(\frac{-W}{n}\right)^2 \right] + 2 \frac{(n)(-n_2)}{(n^2 - n_1 n_2)^2} \\ &= \frac{1}{n} \left[\frac{n_1(-n_2)^2}{(n^2 - n_1 n_2)^2} - \frac{2(-n_2)}{(n^2 - n_1 n_2)} + \frac{n_2(n)^2}{(n^2 - n_1 n_2)^2} + 1 \right] - \frac{2nn_2}{(n^2 - n_1 n_2)^2} \\ &= \frac{1}{n} + \frac{n_2}{n(n^2 - n_1 n_2)}. \end{aligned}$$

Case 3: When $i \geq n_1$ and $1 \leq j \leq n_2$

Using the same argument as in Case 2, by symmetry, we get

$$\|\chi_{\|S}^{(ij)}\|^2 = \frac{1}{n} + \frac{n_1}{n(n^2 - n_1 n_2)}.$$

Case 4: When $i \geq n_1$ and $j \geq n_2$

When $\mathbf{x} = \chi^{(ij)}$, with $i \geq n_1$ and $j \geq n_2$, it is orthogonal to the space $\mathcal{S}_{n_1 n_2}$, and so we get $\|\chi_{\|S}^{(ij)}\|^2 = 0$. □

References

1. Alizadeh, M., Yang, S., Sharif, M., Katti, S., McKeown, N., Prabhakar, B., Shenker, S.: pFabric: Minimal near-optimal datacenter transport. In: Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM, SIGCOMM '13, pp. 435–446 (2013)
2. Andrews, M., Jung, K., Stolyar, A.: Stability of the max-weight routing and scheduling protocol in dynamic networks and at critical loads. In: Proceedings of the Thirty-ninth Annual ACM Symposium on Theory of Computing, STOC '07, pp. 145–154 (2007)
3. Bertsimas, D., Gamarnik, D., Tsitsiklis, J.N.: Performance of multiclass Markovian queueing networks via piecewise linear Lyapunov functions. *Ann. Appl. Probab.* **11**(4), 1384–1428 (2001)
4. Bonald, T., Proutiere, A.: Insensitivity in processor-sharing networks. *Perform. Eval.* **49**(1), 193–209 (2002)
5. Boyd, S., Vandenberghe, L.: *Convex Optimization*. Cambridge University Press, New York, NY (2004)
6. Eryilmaz, A., Srikant, R.: Asymptotically tight steady-state queue length bounds implied by drift conditions. *Queueing Syst.* **72**(3–4), 311–359 (2012)
7. Hajek, B.: Hitting-time and occupation-time bounds implied by drift analysis with applications. *Adv. Appl. Probab.* **14**(3), 502–525 (1982)

8. Kang, W.N., Williams, R.J.: Diffusion approximation for an input-queued switch operating under a maximum weight matching policy. *Stoch. Syst.* **2**(2), 277–321 (2012). <https://doi.org/10.1214/12-SSY061>
9. Lu, Y., Maguluri, S.T., Squillante, M., Suk, T.: On optimal weighted-delay scheduling in input-queued switches (2017). Arxiv preprint [arXiv:1704.02302](https://arxiv.org/abs/1704.02302)
10. Maguluri, S.T., Burle, S.K., Srikant, R.: Optimal heavy-traffic queue length scaling in an incompletely saturated switch. In: Proceedings of the 2016 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Science, pp. 13–24. ACM (2016)
11. Maguluri, S.T., Srikant, R.: Heavy traffic queue length behavior in a switch under the MaxWeight algorithm. *Stoch. Syst.* **6**(1), 211–250 (2016). <https://doi.org/10.1214/15-SSY193>
12. McKeown, N., Anantharam, V., Walrand, J.: Achieving 100% throughput in an input queued switch. In: Proceedings of IEEE INFOCOM, pp. 296–302 (1996)
13. Neely, M.J., Modiano, E., Cheng, Y.S.: Logarithmic delay for $n \times n$ packet switches under the crossbar constraint. *IEEE/ACM Trans. Netw.* **15**(3), 657–668 (2007)
14. Perry, J., Ousterhout, A., Balakrishnan, H., Shah, D., Fugal, H.: Fastpass: a centralized “zero-queue” datacenter network. In: Proceedings of the 2014 ACM Conference on SIGCOMM, SIGCOMM '14, pp. 307–318. ACM, New York, NY (2014). <https://doi.org/10.1145/2619239.2626309>
15. Shah, D., Tsitsiklis, J., Zhong, Y.: Optimal scaling of average queue sizes in an input-queued switch: an open problem. *Queueing Syst.* **68**(3–4), 375–384 (2011)
16. Shah, D., Tsitsiklis, J.N., Zhong, Y.: On queue-size scaling for input-queued switches (2015). Arxiv preprint [arXiv:1405.4764](https://arxiv.org/abs/1405.4764)
17. Shah, D., Walton, N.S., Zhong, Y.: Optimal queue-size scaling in switched networks. *Ann. Appl. Probab.* **24**(6), 2207–2245 (2014)
18. Shah, D., Wischik, D.: Switched networks with maximum weight policies: fluid approximation and multiplicative state space collapse. *Ann. Appl. Probab.* **22**(1), 70–127 (2012)
19. Singh, R., Stolyar, A.: Maxweight scheduling: asymptotic behavior of unscaled queue-differentials in heavy traffic. arXiv preprint [arXiv:1502.03793](https://arxiv.org/abs/1502.03793) (2015)
20. Srikant, R., Ying, L.: *Communication Networks: An Optimization, Control and Stochastic Networks Perspective*. Cambridge University Press, Cambridge (2014)
21. Stolyar, A.L.: Maxweight scheduling in a generalized switch: state space collapse and workload minimization in heavy traffic. *Ann. Appl. Probab.* **14**(1), 1–53 (2004)
22. Tassiulas, L., Ephremides, A.: Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Trans. Autom. Control* **37**(12), 1936–1948 (1992)
23. Wang, W., Maguluri, S.T., Srikant, R., Ying, L.: Heavy-traffic delay insensitivity in connection-level models of data transfer with proportionally fair bandwidth sharing. In: IFIP WG 7.3 Performance 2017 (2017)
24. Weller, T., Hajek, B.: Scheduling nonuniform traffic in a packet-switching system with small propagation delay. *IEEE/ACM Trans. Netw.* **5**(6), 813–823 (1997)
25. Ziegler, G.: *Lectures on Polytopes*. Graduate Texts in Mathematics. Springer, New York (1995)