

HEAVY TRAFFIC QUEUE LENGTH BEHAVIOR IN A SWITCH UNDER THE MAXWEIGHT ALGORITHM*

BY SIVA THEJA MAGULURI[†] AND R. SRIKANT[‡]

*IBM T. J. Watson Research Center[†]
University of Illinois at Urbana-Champaign[‡]*

We consider a switch operating under the MaxWeight scheduling algorithm, under any traffic pattern such that all the ports are loaded. This system is interesting to study since the queue lengths exhibit a multi-dimensional state-space collapse in the heavy-traffic regime. We use a Lyapunov-type drift technique to characterize the heavy-traffic behavior of the expectation of the sum queue lengths in steady-state, under the assumption that all ports are saturated and all queues receive non-zero traffic. Under these conditions, we show that the heavy-traffic scaled queue length is given by $(1 - \frac{1}{2n})\|\sigma\|^2$, where σ is the vector of the standard deviations of arrivals to each port in the heavy-traffic limit. In the special case of uniform Bernoulli arrivals, the corresponding formula is given by $(n - \frac{3}{2} + \frac{1}{2n})$. The result shows that the heavy-traffic scaled queue length has optimal scaling with respect to n , thus settling one version of an open conjecture; in fact, it is shown that the heavy-traffic queue length is at most within a factor of two from the optimal. We then consider certain asymptotic regimes where the load of the system scales simultaneously with the number of ports. We show that the MaxWeight algorithm has optimal queue length scaling behavior provided that the arrival rate approaches capacity sufficiently fast.

1. Introduction. Consider a collection of queues arranged in the form of an $n \times n$ matrix. The queues are assumed to operate in discrete-time and jobs arriving to the queues will be called packets. The following constraints are imposed on the service process of the queueing system: (a) at most one queue can be served in each time slot in each row of the matrix, (b) at most one queue can be served in each time slot in each column of the matrix, and (c) when a queue is served, at most one packet can be removed from the queue. Such a queueing system is called a *switch*.

A scheduling algorithm for the switch is a rule which selects the queues to be served in each time slot. A well-known algorithm called the MaxWeight

Received July 2015.

*The work presented here was supported in part by NSF Grant ECCS-1202065.

MSC 2010 subject classifications: 60K25, 90B15.

Keywords and phrases: Switch, scheduling, MaxWeight, state space collapse, heavy traffic.

algorithm is known to optimize the throughput in a switch. The algorithm was derived in a more general context in [1] and for the special context of the switch considered in here in [2], where it was also shown that other seemingly good policies are not throughput-optimal. An important open question that is not fully understood is whether the MaxWeight algorithm is also queue length or delay optimal in any sense. In [3], it was shown that the MaxWeight algorithm minimizes the sum of the squares of the queue lengths in heavy-traffic under a condition called *Complete Resource Pooling* (CRP). For the switch, the CRP condition means that the arriving traffic saturates at most one column or one row of the switch. The result relies on the fact that, under CRP and in the heavy-traffic regime, there is a one-dimensional state-space collapse, i.e., the state of the system collapses to a line. When the CRP condition is not met, the state-space collapses to a lower-dimension, but is not one-dimensional. State-space collapse without the CRP condition was established in [4] when the arrivals are deterministic. For stochastic arrivals, state-space collapse for the fluid limit was studied in [5], and a diffusion limit has been established in [6]. However, a characterization of the steady-state behavior of the diffusion limit was still open.

In this paper, we use the Lyapunov-type drift technique introduced in [7]. The basic idea is to set the drift of an appropriately chosen function equal to zero in steady-state to obtain both upper and lower bounds on quantities of interest, such as the moments of the queue lengths. Setting the drift of a function to zero in steady state is analogous to the basic adjoint relation (BAR) of diffusion limit systems such as the ones studied in [8]. To obtain upper bounds one has to establish state-space collapse in a sense that is somewhat different than the one in [3]: the main difference being that the state-space collapse is expressed in terms of the moments of the queue lengths in steady-state. This form of state-space collapse can then be readily used in the drift condition to obtain the upper bound. However, in [7], the usefulness of the drift technique was only established under the CRP condition. In this paper, we consider the switch when all the ports are saturated, i.e., in the heavy-traffic regime, the traffic in all rows and columns approach capacity, and the CRP condition is violated. The main contribution of the paper is to characterize the expected steady-state queue lengths in heavy-traffic even though the CRP condition is violated. As mentioned earlier, when the CRP condition is violated, the state does not typically collapse to a single dimension. The main challenge in our proof is due to the difficulty in characterizing the behavior of the queue length process under such a multi-dimensional state-space collapse. Characterizing the behavior

of the queue lengths under multi-dimensional state-space collapse has been difficult, in general, except in rare cases; see [9, 10] for two such examples in other contexts.

The difficulty in understanding the steady-state queue length behavior of the MaxWeight algorithm has meant that it is unknown whether the MaxWeight algorithm minimizes the expected total queue length in steady-state. One way to pose the optimality question is to increase the number of queues in the system, or increase the arrival to a point close to the boundary of the capacity region (the heavy-traffic regime), or do both, and study whether the MaxWeight algorithm is queue-length-optimal in a scaling sense. A conjecture regarding the scaling behavior for any algorithm, both in heavy-traffic and under all traffic conditions, has been stated in [11]. The authors first heard about the non-heavy-traffic version of this conjecture from A. L. Stolyar in 2005. The conjecture seemed to be difficult to verify for the MaxWeight algorithm, and so a number of other algorithms have been developed to achieve either optimal or near-optimal scaling behavior; see [12, 13, 14]. The results in this paper establish the validity of one version of the conjecture (pertaining to the heavy-traffic regime) for the MaxWeight algorithm.

Note on Notation: The set of real numbers, and the set of non-negative real numbers are denoted by \mathbb{R} , and \mathbb{R}_+ respectively. We work in the n^2 – dimensional Euclidean space \mathbb{R}^{n^2} . We represent vectors in this space in bold font, by \mathbf{x} . We use two indices $1 \leq i \leq n$ and $1 \leq j \leq n$ for different components of \mathbf{x} . We represent the $(i, j)^{\text{th}}$ component by x_{ij} and thus, $\mathbf{x} = (x_{ij})_{ij}$. 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 with i^{th} row being all ones and zeros every where 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 with j^{th} column being all ones and zeros every where else. For a random process $\mathbf{q}(t)$ and a Lyapunov 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.

2. Preliminaries. In this section, we will present the model of an input queued switch, MaxWeight scheduling algorithm, some observations on the geometry of the capacity region and other preliminaries.

2.1. *System model and MaxWeight algorithm.* An input queued switch is a model for cross-bar switches that are widely used. An $n \times n$ switch has n input ports and n output ports. We consider a discrete time system. In each time slot t , packets arrive at any of the input ports to be delivered to any of the output ports. When scheduled, each packet needs one time slot to be transmitted across.

Each input port maintains n separate queues, one each for packets to be delivered to each of the n output ports. 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.

Let $a_{ij}(t)$ denote the number of packet arrivals at input port i at time t to be delivered to output port j , and we let $\mathbf{a} \in \mathbb{R}^{n^2}$ denote the vector $(a_{ij})_{ij}$. For every input-output pair (i, j) , the arrival process $a_{ij}(t)$ is a stochastic process that is i.i.d across time, with mean $\mathbb{E}[a_{ij}(t)] = \lambda_{ij}$ and variance $\text{Var}(a_{ij}(t)) = \sigma_{ij}^2$ for any time t . We assume that the arrival processes are independent across input-output pairs, (i.e, if $(i, j) \neq (i', j')$, the processes $a_{ij}(t)$ and $a_{i'j'}(t)$ are independent) 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 $a_{\max} \geq 1$ and $P(a_{ij}(t) = 0) > \epsilon_a$ for some $\epsilon_a > 0$. 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 $(\sigma_{ij})_{ij}$.

In each time slot, each input port can be matched to only one output port and similarly, each output port can be mapped to only one input port. These constraints can be captured in a graph. Let G denote a complete $n \times n$ bipartite graph with n^2 edges between the set of input ports and the set of output ports. The schedule in each time slot is a matching on this graph G . We let $s_{ij} = 1$ 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, $\mathcal{S} \subset \mathbb{R}^{n^2}$ is defined as follows.

$$\mathcal{S} = \left\{ \mathbf{s} \in \{0, 1\}^{n^2} : \sum_{i=1}^n s_{ij} \leq 1, \sum_{j=1}^n s_{ij} \leq 1 \forall i, j \in \{1, 2, \dots, n\} \right\}.$$

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

$$\mathcal{S}^* = \left\{ \mathbf{s} \in \{0, 1\}^{n^2} : \sum_{i=1}^n s_{ij} = 1, \sum_{j=1}^n s_{ij} = 1 \forall i, j \in \{1, 2, \dots, n\} \right\}.$$

Each element in this set corresponds to a perfect matching on the graph G . Each of these maximal feasible schedules is also a permutation π on the set $1, 2, \dots, n$ with $\pi(i) = j$ if $s_{ij} = 1$.

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.

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

where $[x]^+ = \max(0, x)$ is the projection onto positive real axis, $u_{ij}(t)$ 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, we have that 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 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, i.e.,

$$\limsup_{C \rightarrow \infty} \limsup_{t \rightarrow \infty} \mathbb{P} \left(\sum_{ij} q_{ij}(t) \geq C \right) = 0.$$

If the queue lengths process $\mathbf{q}(t)$ is positive recurrent under a scheduling policy, then we have stability. 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. The MaxWeight Algorithm is a popular scheduling algorithm for the 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. This algorithm is presented in Algorithm 1. It is possible to show that the Markov chain $\mathbf{q}(t)$ is irreducible and aperiodic under the MaxWeight algorithm for an appropriately defined state space [15, Exercise 4.2]. It is well known [1, 2] that the capacity region \mathcal{C} of the switch is the convex hull of all feasible schedules,

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

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.

Algorithm 1 MaxWeight scheduling algorithm for an input-queued switch

Consider the complete bipartite graph between the input ports and output ports. Let the queue length $q_{ij}(t)$ be the weight of the edge between input port i and output port j at time t . In time slot t , the schedule chosen is given by the maximum weight matching, i.e.,

$$(1) \quad \mathbf{s}(t) = \arg \max_{\mathbf{s} \in \mathcal{S}} \sum_{ij} q_{ij}(t) s_{ij} = \arg \max_{\mathbf{s} \in \mathcal{S}} \langle \mathbf{q}(t), \mathbf{s} \rangle$$

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} and \mathbf{s}^* have an unused service of 1. Note that this does not change the scheduling algorithm, but it is just a convenience in the proof. Therefore, without loss of generality, we assume that the schedule chosen in each time slot is a maximal schedule, i.e.,

$$(2) \quad \mathbf{s}(t) \in \mathcal{S}^* \text{ for all time } t.$$

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

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 MaxWeight policy. Let $\bar{\mathbf{q}}$ denote the steady state random vector. In this paper, we focus on the average queue length under the steady state distribution, i.e., $\mathbb{E}[\sum_{i,j} \bar{q}_{ij}]$. We consider a set of systems indexed by ϵ with arrival rate $\boldsymbol{\lambda}^\epsilon = (1 - \epsilon)\boldsymbol{\nu}$, where $\boldsymbol{\nu}$ is an arrival rate on the boundary of the capacity region \mathcal{C} such that all the input and output ports are saturated and $\nu_{ij} > 0$ for all i, j . The load of each system is then $(1 - \epsilon)$. We will study the switch when $\epsilon \downarrow 0$. This is called the heavy traffic limit. We first show a universal lower bound on the average queue length in heavy traffic limit, i.e., on $\lim_{\epsilon \rightarrow 0} \mathbb{E}[\sum_{i,j} \bar{q}_{ij}]$. We then show that under MaxWeight policy, the limiting average queue length is within a factor of less than 2 of the universal lower bound and thus MaxWeight has optimal average queue length scaling. We will show these bounds using Lyapunov drift conditions. We will use several different quadratic Lyapunov functions through out the paper.

2.2. Geometry of the capacity region. The capacity region \mathcal{C} is a coordinate convex polytope in \mathbb{R}^{n^2} . Here, we review some basic definitions. For any set $P \in \mathbb{R}^m$, its dimension is defined by

$$\dim(P) \triangleq \min\{\dim(A) | P \subseteq A, A \text{ is an affine space}\}.$$

So the capacity region \mathcal{C} has dimension n^2 . A hyperplane H is said to be a supporting hyperplane of a polytope P if $P \cap H \neq \emptyset$, $P \cap H_+ \neq \emptyset$ and $P \cap H_- = \emptyset$ where H_+ and H_- are the open half-spaces determined by the hyperplane H . For any supporting hyperplane H of polytope P , $P \cap H$ is called a face [16]. A face of a polytope is also a polytope with lower dimension. A face F of polytope P with dimension $\dim(F) = \dim(P) - 1$ is called a *facet*. Heavy traffic optimality of MaxWeight algorithm for generalized switches is shown in [3, 7] when a single input or output port is saturated or in other words when approaching an arrival rate vector on a facet of the capacity region. However, in this paper, we are interested in the case when all the ports are saturated. The arrival rate vector $\boldsymbol{\nu}$ in this case does not lie on a facet and so, that result is not applicable here.

When $\boldsymbol{\nu}$ is the arrival rate vector on the boundary of the capacity region such that all the input ports and all the output ports are saturated, it lies on the face \mathcal{F} ,

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

It is easy to see that \mathcal{F} as defined here is indeed a face by observing that the hyperplane $\langle \boldsymbol{\lambda}, \mathbf{1} \rangle = n$ is a supporting hyperplane of the capacity region \mathcal{C} and it contains any rate vector $\boldsymbol{\nu}$ where all the ports are saturated. The face \mathcal{F} has dimension $(n-1)^2 = n^2 - (2n-1)$, and lies in the affine space formed by the intersection of the $2n$ constraints, $\{\sum_{i=1}^n \lambda_{ij} = 1 \text{ for all } j\}$, and $\{\sum_{j=1}^n \lambda_{ij} = 1 \text{ for all } i\}$. Of these $2n$ constraints, one is linearly dependent of the others and we have $2n-1$ linearly independent constraints. The face \mathcal{F} is actually the convex combination of the maximal feasible schedules \mathcal{S}^* , i.e., $\mathcal{F} = \text{Conv}(\mathcal{S}^*)$. These results follow from the fact that the face \mathcal{F} is the Birkhoff polytope B_n that contains all the $n \times n$ doubly stochastic matrices. It is known [17, page 20] that B_n lies in the $(n-1)^2$ dimensional affine space of the constraints and is the convex hull of the permutation matrices.

A facet of a polytope has a unique supporting hyperplane defining the facet. It was shown in [7] that when the arrival rate vector approaches a rate vector in the relative interior of a facet, in the limit, the queue length vector concentrates along the direction of the normal vector of the unique supporting hyperplane. However, a lower dimensional face can be defined by one of several hyperplanes, and so there is no unique normal vector. A lower dimensional face is always an intersection of two or more facets. We are interested in the case when the arrival rate vector approaches the vector $\boldsymbol{\nu}$ that lies on the face \mathcal{F} . The face \mathcal{F} is the intersection of the $2n$ facets, $\{\langle \mathbf{e}^{(i)}, \boldsymbol{\lambda} \rangle = 1\} \cap \mathcal{C}$ for all i , and $\{\langle \tilde{\mathbf{e}}^{(j)}, \boldsymbol{\lambda} \rangle = 1\} \cap \mathcal{C}$ for all j . We will show in section 4 that in the heavy traffic limit, the queue length vector concentrates within the cone spanned by the $2n$ normal vectors, $\{\mathbf{e}^{(i)} \text{ for all } i\} \cup \{\tilde{\mathbf{e}}^{(j)} \text{ for all } j\}$. We will call this cone \mathcal{K} . Here, we will present some definitions and other results related to this cone. More formally, the cone \mathcal{K} can be defined as follows.

$$\mathcal{K} = \left\{ \mathbf{x} \in \mathbb{R}^{n^2} : \mathbf{x} = \sum_i w_i \mathbf{e}^{(i)} + \sum_j \tilde{w}_j \tilde{\mathbf{e}}^{(j)} \text{ where } w_i \in \mathbb{R}_+, \tilde{w}_j \in \mathbb{R}_+ \forall i, j \right\}.$$

The following lemma presents several equivalent ways of characterizing the cone \mathcal{K} . The proof of the lemma is presented in Appendix A.

LEMMA 1. *For a vector $\mathbf{x} \in \mathbb{R}^{n^2}$, the following are equivalent.*

- (i) $\mathbf{x} \in \mathcal{K}$.
- (ii) *There are $w_i \in \mathbb{R}_+$ and $\tilde{w}_j \in \mathbb{R}_+$ for all $i, j \in \{1, 2, \dots, n\}$ such that $x_{ij} = w_i + \tilde{w}_j$.*
- (iii) $\mathbf{x} \in \mathbb{R}_+^{n^2}$ and when \mathbf{x} is used as the weight, all maximal schedules (permutation matrices) have same weight.

(iv) $\mathbf{x} \in \mathbb{R}_+^{n^2}$ and it satisfies

$$(3) \quad x_{ij} = \frac{1}{n} \sum_{j'=1}^n x_{ij'} + \frac{1}{n} \sum_{i'=1}^n x_{i'j} - \frac{1}{n^2} \sum_{i'=1}^n \sum_{j'=1}^n x_{i'j'} \quad \forall i, j.$$

Note that the representation of the vector \mathbf{x} in terms of w_i 's and \tilde{w}_j 's in (ii) above need not be unique. For example, suppose that $w_i \geq 1$ for all i , then setting $w'_i = w_i - 1$ for each i and $\tilde{w}'_j = \tilde{w}_j + 1$ for each j , we again have that $w'_i \in \mathbb{R}_+$, $\tilde{w}'_j \in \mathbb{R}_+$ and $x_{ij} = w'_i + \tilde{w}'_j$ for all i, j . Equation 3 means that any element in the matrix \mathbf{x} , is equal to the average of all the elements in its row plus the average of all the elements in its column minus the average of all the elements in the whole matrix. Suppose the queue lengths $\mathbf{q} \in \mathcal{K}$, then any queue length from an input port to an output port is equal to the average queue lengths at that input port plus the average queue lengths at that output port minus the average queue length of all the queues in the switch.

We now present some more properties of this cone. The cone \mathcal{K} lies in the $2n - 1$ dimensional subspace spanned by the $2n - 1$ independent vectors among the $2n$ vectors, $\{\mathbf{e}^{(i)} \text{ for all } i\} \cup \{\tilde{\mathbf{e}}^{(j)} \text{ for all } j\}$. Call this space $\mathcal{V}_{\mathcal{K}}$. For any two vectors $\mathbf{x}, \mathbf{y} \in \mathcal{F}$, $\mathbf{x} - \mathbf{y}$ is orthogonal to the subspace $\mathcal{V}_{\mathcal{K}}$, i.e.,

$$(4) \quad \mathbf{x} - \mathbf{y} \perp \mathcal{V}_{\mathcal{K}}.$$

This is easy to see since $\langle \mathbf{x}, \mathbf{e}^{(i)} \rangle = \langle \mathbf{y}, \mathbf{e}^{(i)} \rangle = 1$ for all i and $\langle \mathbf{x}, \tilde{\mathbf{e}}^{(j)} \rangle = \langle \mathbf{y}, \tilde{\mathbf{e}}^{(j)} \rangle = 1$ for all j . If $\mathcal{V}_{\mathcal{F}}$ denotes the subspace obtained by translating the affine space spanned by \mathcal{F} , it follows that the spaces $\mathcal{V}_{\mathcal{K}}$ and $\mathcal{V}_{\mathcal{F}}$ are orthogonal because translation means subtraction by a vector. Moreover, they span the entire space \mathbb{R}^{n^2} since their dimensions sum to n^2 .

As noted earlier, state space collapse to the cone \mathcal{K} has been shown in [4, 5], although our notion of state space collapse is a bit different as will be explained later. Equation 3 in Lemma 1 is essentially what is called a lifting map in [5]. However, we do not explicitly use the notion of workload or lifting map in the sense used in [5].

2.2.1. *Projection onto the cone \mathcal{K}* . The cone \mathcal{K} is closed and convex. For any $\mathbf{x} \in \mathbb{R}^{n^2}$, the closest point in the cone \mathcal{K} to \mathbf{x} is called the projection of \mathbf{x} on to the cone \mathcal{K} and we will denote it by \mathbf{x}_{\parallel} . More formally,

$$\mathbf{x}_{\parallel} = \arg \min_{\mathbf{y} \in \mathcal{K}} \|\mathbf{x} - \mathbf{y}\|.$$

For a closed convex cone \mathcal{K} , the projection \mathbf{x}_{\parallel} is well defined and is unique [18, Appendix E.9.2]. We will use \mathbf{x}_{\perp} to denote $\mathbf{x} - \mathbf{x}_{\parallel}$. We will use $x_{\parallel ij}$ to denote the $(i, j)^{\text{th}}$ component of \mathbf{x}_{\parallel} . Similarly, $x_{\perp ij}$.

Note that unlike projection on to a subspace, projection on to a cone is not linear, i.e., $(\mathbf{x} + \mathbf{y})_{\parallel} \neq \mathbf{x}_{\parallel} + \mathbf{y}_{\parallel}$. A simple counter example is the following. In \mathbb{R}^2 , let $\mathbf{x} = (2, 2)$ and $\mathbf{y} = (-1, -1)$. Consider the positive quadrant as the cone of interest. Then, $\mathbf{x}_{\parallel} = (2, 2)$, $\mathbf{y}_{\parallel} = (0, 0)$ and $(\mathbf{x} + \mathbf{y})_{\parallel} = (1, 1)$.

Since for any $\mathbf{x} \in \mathbb{R}^{n^2}$, $\mathbf{x}_{\parallel} \in \mathcal{K}$, from the definition of the cone \mathcal{K} , we have that every component of \mathbf{x}_{\parallel} is non negative, i.e., $x_{\parallel ij} \geq 0$. However, \mathbf{x}_{\perp} could have negative components.

The polar cone \mathcal{K}° of cone \mathcal{K} is defined as

$$\mathcal{K}^{\circ} = \left\{ \mathbf{x} \in \mathbb{R}^{n^2} : \langle \mathbf{x}, \mathbf{y} \rangle \leq 0 \text{ for all } \mathbf{y} \in \mathcal{K} \right\}.$$

The polar cone \mathcal{K}° is negative of the dual cone \mathcal{K}^* of the cone \mathcal{K} . For any $\mathbf{x} \in \mathbb{R}^{n^2}$, $\mathbf{x}_{\perp} \in \mathcal{K}^{\circ}$ and $\langle \mathbf{x}_{\parallel}, \mathbf{x}_{\perp} \rangle = 0$ [18, Appendix E.9.2]. Therefore, pythagoras theorem is applicable, i.e.,

$$(5) \quad \|\mathbf{x}\|^2 = \|\mathbf{x}_{\parallel}\|^2 + \|\mathbf{x}_{\perp}\|^2$$

and so, $\|\mathbf{x}_{\parallel}\| \leq \|\mathbf{x}\|$ and $\|\mathbf{x}_{\perp}\| \leq \|\mathbf{x}\|$.

Projection onto any closed convex set in \mathbb{R}^{n^2} (and so onto a closed convex cone) is nonexpansive [18, Appendix E.9.3]. Therefore, we have $\|\mathbf{x}_{\parallel} - \mathbf{y}_{\parallel}\| \leq \|\mathbf{x} - \mathbf{y}\|$. Since \mathbf{x}_{\perp} is a projection onto \mathcal{K}° , we also have

$$(6) \quad \|\mathbf{x}_{\perp} - \mathbf{y}_{\perp}\| \leq \|\mathbf{x} - \mathbf{y}\|.$$

2.3. *Moment bounds from Lyapunov drift conditions.* In this paper, we will use the Lyapunov drift based approach presented in [7] to obtain bounds of average queue length under MaxWeight. A key ingredient in this approach is to obtain moment bounds from drift conditions. A lemma from [19] was used in [7] to obtain these bounds and we first state it here as it was stated in [7].

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 $\kappa < \infty$ such that for any $t = 1, 2, \dots$ and for all $X \in \mathcal{X}$ with $Z(X) \geq \kappa$,*

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

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

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

Then, there exists a $\theta^ > 0$ and a $C^* < \infty$ such that*

$$\limsup_{t \rightarrow \infty} \mathbb{E} \left[e^{\theta^* Z(X(t))} \right] \leq C^*.$$

If we further assume that the Markov chain $\{X(t)\}_t$ is positive recurrent, then $Z(X(t))$ converges in distribution to a random variable \bar{Z} for which

$$\mathbb{E} \left[e^{\theta^* \bar{Z}} \right] \leq C^*,$$

which directly implies that all moments of \bar{Z} exist and are finite.

This lemma (and its original form in [19]) is quiet general and versatile. However, we use a different result in this paper to obtain moment bounds that are tighter than the bounds that can be obtained using Lemma 2 (or its original form in [19]). The following lemma essentially follows from [20, Theorem 1] except for some minor differences. The proof is presented in Appendix B and makes use of Lemma 2.

LEMMA 3. *Consider an irreducible and aperiodic Markov chain 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. The drift $\Delta Z(X)$ of Z at X as defined in Lemma 2 is assumed to satisfy the conditions C.1 and C.2. Further assume that the Markov chain $\{X(t)\}_t$ converges in distribution to a random variable \bar{X} . Then, for any $m = 0, 1, 2, \dots$,*

$$\mathbb{P}(Z(\bar{X}) > \kappa + 2Dm) \leq \left(\frac{D}{D + \eta} \right)^{m+1}.$$

As a result, for any $r = 1, 2, \dots$,

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

3. Universal lower bound. In this section, we will prove the following lower bound on the average queue lengths, which is valid under any scheduling policy.

PROPOSITION 1. *Consider a set of switch systems with the arrival processes $\mathbf{a}^{(\epsilon)}(t)$ described in Section 2.1, parameterized by $0 < \epsilon < 1$, such that the mean arrival rate vector is $\boldsymbol{\lambda}^\epsilon = (1 - \epsilon)\boldsymbol{\nu}$ for some $\boldsymbol{\nu} \in \mathcal{F}$ and variance is $(\boldsymbol{\sigma}^{(\epsilon)})^2$. The load is then $\rho = (1 - \epsilon)$. 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 average queue length is lower bounded by*

$$\mathbb{E} \left[\sum_{ij} \bar{q}_{ij}^{(\epsilon)} \right] \geq \frac{\|\boldsymbol{\sigma}^{(\epsilon)}\|^2}{2\epsilon} - \frac{n(1 - \epsilon)}{2}.$$

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

$$\liminf_{\epsilon \downarrow 0} \epsilon \mathbb{E} \left[\sum_{ij} \bar{q}_{ij}^{(\epsilon)} \right] \geq \frac{\|\boldsymbol{\sigma}\|^2}{2}.$$

PROOF. We will obtain a lower bound on sum of all the queue lengths by lower bounding the queue lengths at each input port, i.e., we will first bound $\mathbb{E}[\sum_j \bar{q}_{ij}^{(\epsilon)}]$ for a fixed input port i . We do this by considering a single queue that is coupled to the process $\sum_j q_{ij}^{(\epsilon)}(t)$. Consider a single server queue $\phi_i^{(\epsilon)}(t)$ in discrete time. Packets arrive into this queue to be served. Each packet needs exactly one time slot of service. The arrival process to this queue is $\alpha_i^{(\epsilon)}(t) = \sum_j a_{ij}^{(\epsilon)}(t)$. Mean arrival to this queue is $\mathbb{E}[\alpha_i^{(\epsilon)}(t)] = \sum_j \lambda_{ij}^{(\epsilon)} = (1 - \epsilon) \sum_j \nu_{ij}^{(\epsilon)} = (1 - \epsilon)$ since $\boldsymbol{\nu} \in \mathcal{F}$. As long as the queue is non empty, one packet is served in every time slot. Thus, this queue evolves as

$$\begin{aligned} \phi_i^{(\epsilon)}(t+1) &= \left[\phi_i^{(\epsilon)}(t) + \alpha_i^{(\epsilon)}(t) - 1 \right]^+ \\ &= \phi_i^{(\epsilon)}(t) + \alpha_i^{(\epsilon)}(t) - 1 + v^{(\epsilon)}(t), \end{aligned}$$

where $v^{(\epsilon)}(t)$ is the unused service and so $v^{(\epsilon)}(t)\phi_i^{(\epsilon)}(t+1) = 0$. Clearly, $\phi_i^{(\epsilon)}(t)$ is positive recurrent and let $\bar{\phi}_i^{(\epsilon)}$ denote the steady state random variable to which it converges in distribution.

CLAIM 1. *In steady state,*

$$\mathbb{E}\left[\sum_j \bar{q}_{ij}^{(\epsilon)}\right] \geq \mathbb{E}[\bar{\phi}_i^{(\epsilon)}].$$

PROOF. Suppose that at time zero, the queue $\phi_i^{(\epsilon)}$ starts with $\phi_i^{(\epsilon)}(0) = \sum_j q_{ij}^{(\epsilon)}(0)$. Then, for any time t , the queue $\phi_i^{(\epsilon)}(t)$ is stochastically no greater than $\sum_j q_{ij}^{(\epsilon)}(t)$. This can easily be seen using induction. For $t = 0$, we have $\sum_j q_{ij}(0) \geq \phi_i(0)$. Suppose that $\sum_j q_{ij}^{(\epsilon)}(t) \geq \phi_i^{(\epsilon)}(t)$. Then, at time $(t + 1)$,

$$\begin{aligned} \sum_j q_{ij}^{(\epsilon)}(t + 1) &= \sum_j \left[q_{ij}^{(\epsilon)}(t) + a_{ij}^{(\epsilon)}(t) - s_{ij}^{(\epsilon)}(t) \right]^+ \\ &\geq \left[\sum_j (q_{ij}^{(\epsilon)}(t) + a_{ij}^{(\epsilon)}(t) - s_{ij}^{(\epsilon)}(t)) \right]^+ \\ &\geq \left[\phi_i^{(\epsilon)}(t) + \alpha_i^{(\epsilon)}(t) - 1 \right]^+ \\ &= \phi_i^{(\epsilon)}(t + 1), \end{aligned}$$

where the last inequality follows from the inductive hypothesis, definition of $\alpha^{(\epsilon)}(t)$, the constraint $s_{ij}^{(\epsilon)}(t) \leq 1$ and the fact that if $x \geq y$, we have that $[x]^+ \geq [y]^+$. Thus, we have that in steady state, $\mathbb{E}[\sum_j \bar{q}_{ij}^{(\epsilon)}] \geq \mathbb{E}[\bar{\phi}_i^{(\epsilon)}]$. Since steady state distribution of $\sum_j \bar{q}_{ij}^{(\epsilon)}$ and $\bar{\phi}_i^{(\epsilon)}$ does not depend on the initial state at time zero, we have the lower bound $\mathbb{E}[\sum_j \bar{q}_{ij}^{(\epsilon)}] \geq \mathbb{E}[\bar{\phi}_i^{(\epsilon)}]$ independent of the initials states $\phi_i^{(\epsilon)}(0)$ and $\sum_j q_{ij}^{(\epsilon)}(0)$. \square

We will now bound $\mathbb{E}[\bar{\phi}_i^{(\epsilon)}]$. This result is obtained in [15]. We present it here for completeness. Consider the drift of $\mathbb{E}[(\phi_i^{(\epsilon)}(t))^2]$.

$$\begin{aligned} &\mathbb{E}[(\phi_i^{(\epsilon)}(t + 1))^2 - (\phi_i^{(\epsilon)}(t))^2] \\ &= \mathbb{E}[(\phi_i^{(\epsilon)}(t) + \alpha_i^{(\epsilon)}(t) - 1 + v^{(\epsilon)}(t))^2 - (\phi_i^{(\epsilon)}(t))^2] \\ &\stackrel{(a)}{=} \mathbb{E}[(\phi_i^{(\epsilon)}(t) + \alpha_i^{(\epsilon)}(t) - 1)^2 - (v^{(\epsilon)}(t))^2 - (\phi_i^{(\epsilon)}(t))^2] \\ &= \mathbb{E}[(\alpha_i^{(\epsilon)}(t) - 1)^2 + 2(\phi_i^{(\epsilon)}(t))(\alpha_i^{(\epsilon)}(t) - 1) - (v^{(\epsilon)}(t))^2] \\ &\stackrel{(b)}{=} \mathbb{E}[(\alpha_i^{(\epsilon)}(t) - (1 - \epsilon) - \epsilon)^2] - 2\epsilon\mathbb{E}[\phi_i^{(\epsilon)}(t)] - \mathbb{E}[v^{(\epsilon)}(t)] \\ &\stackrel{(c)}{=} \text{Var}(\alpha_i^{(\epsilon)}(t)) + \epsilon^2 - 2\epsilon\mathbb{E}[\phi_i^{(\epsilon)}(t)] - \mathbb{E}[v^{(\epsilon)}(t)] \\ &\stackrel{(d)}{=} \sum_j (\sigma_{ij}^{(\epsilon)})^2 + \epsilon^2 - 2\epsilon\mathbb{E}[\phi_i^{(\epsilon)}(t)] - \mathbb{E}[v^{(\epsilon)}(t)], \end{aligned}$$

where (a) follows from noting that $(v^{(\epsilon)}(t))(\phi_i^{(\epsilon)}(t) + \alpha_i^{(\epsilon)}(t) - 1 + v^{(\epsilon)}(t)) = 0$; (b) follows from independence of $\phi_i^{(\epsilon)}(t)$ and the arrivals $\alpha_i^{(\epsilon)}(t)$ and since $v^{(\epsilon)}(t) \in \{0, 1\}$; (c) follows from the fact that $\mathbb{E}[\phi_i^{(\epsilon)}(t)] = (1 - \epsilon)$; (d) follows from the definition of $\alpha_i^{(\epsilon)}(t)$ and independence of the arrival process $a_{ij}(t)$ across ports. It can easily be shown that $\mathbb{E}[(\bar{\phi}_i^{(\epsilon)})^2]$ is finite [15, Section 10.1]. Therefore, the steady state drift of $\mathbb{E}[(\phi_i^{(\epsilon)}(t))^2]$ is zero, i.e., in steady-state, we get

$$(7) \quad 2\epsilon\mathbb{E}[\bar{\phi}_i^{(\epsilon)}] = \sum_j \left(\sigma_{ij}^{(\epsilon)}\right)^2 + \epsilon^2 - \mathbb{E}[\bar{v}^{(\epsilon)}].$$

Consider the drift of $\mathbb{E}[\phi_i^{(\epsilon)}(t)]$.

$$\begin{aligned} \mathbb{E}[\phi_i^{(\epsilon)}(t+1) - \phi_i^{(\epsilon)}(t)] &= \mathbb{E}[\alpha_i^{(\epsilon)}(t) - 1 + v^{(\epsilon)}(t)] \\ &= -\epsilon + \mathbb{E}[v^{(\epsilon)}(t)]. \end{aligned}$$

Since $\phi_i^{(\epsilon)}(t) \in \mathbb{Z}_+$, we have $\phi_i^{(\epsilon)}(t) \leq (\phi_i^{(\epsilon)}(t))^2$, and so we get finiteness of $\mathbb{E}[\bar{\phi}_i^{(\epsilon)}]$ from that of $\mathbb{E}[(\bar{\phi}_i^{(\epsilon)})^2]$. Therefore, the drift of $\mathbb{E}[\phi_i^{(\epsilon)}(t)]$ is zero in steady state. Thus, we get that in steady state, $\mathbb{E}[\bar{v}^{(\epsilon)}] = \epsilon$. Substituting this in (7), and using the claim, we get

$$(8) \quad \mathbb{E}\left[\sum_j \bar{q}_{ij}^{(\epsilon)}\right] \geq \mathbb{E}[\bar{\phi}_i^{(\epsilon)}] = \frac{1}{2\epsilon} \sum_j \left(\sigma_{ij}^{(\epsilon)}\right)^2 - \frac{1 - \epsilon}{2}.$$

Since this lower bound is true for any input port i , summing over all the input ports, we have the proposition. Note that we could have obtained the same bound by similarly lower bounding the sum of lengths of all the queues destined to port j , i.e., $\sum_i q_{ij}^{(\epsilon)}(t)$. \square

We do not know if this lower bound is tight, i.e., if there is a scheduling policy that attains this lower bound. However, in section 5, we show that under MaxWeight scheduling algorithm, the average queue lengths are within a factor of less than 2 away from this universal lower bound, thus showing that MaxWeight has optimal scaling.

4. State space collapse under MaxWeight policy. As mentioned earlier, a closely-related but nevertheless different notion of state-space collapse to the cone \mathcal{K} has also been shown in [4, 5], but we need the type of state-space collapse proved here to establish the upper bounds in Section 5. In this section, we will show that under the MaxWeight scheduling

algorithm, in the heavy traffic limit, the steady state queue length vector concentrates within the cone \mathcal{K} in the following sense. As the parameter ϵ approaches zero, the mean arrival rate approaches the boundary of the capacity region and we know from the lower bound that the average queue lengths go to infinity $\Omega(1/\epsilon)$. We will show that under the MaxWeight algorithm, the $\bar{\mathbf{q}}_{\perp}^{(\epsilon)}$ component of the queue length vector is upper bounded independent of ϵ . Thus the $\bar{\mathbf{q}}_{\perp}^{(\epsilon)}$ component is negligible compared to the $\bar{\mathbf{q}}_{\parallel}^{(\epsilon)}$ component of $\bar{\mathbf{q}}^{(\epsilon)}$. This is called state space collapse. We say that the state space collapses to the cone \mathcal{K} . It was shown in [21] that the state space collapses to the subspace containing the cone \mathcal{K} . Here, we show the stronger result that the state space collapses to the cone, which is essential to obtain the upper bounds in Section 5. A similar result was also shown in [22] for a different problem.

We define the following quadratic Lyapunov functions and their corresponding drifts.

$$\begin{aligned}
 V(\mathbf{q}) &\triangleq \|\mathbf{q}\|^2 = \sum_{ij} q_{ij}^2 & W_{\perp}(\mathbf{q}) &\triangleq \|\mathbf{q}_{\perp}\| \\
 V_{\perp}(\mathbf{q}) &\triangleq \|\mathbf{q}_{\perp}\|^2 = \sum_{ij} q_{\perp ij}^2 & V_{\parallel}(\mathbf{q}) &\triangleq \|\mathbf{q}_{\parallel}\|^2 = \sum_{ij} q_{\parallel ij}^2 \\
 \Delta V(\mathbf{q}) &\triangleq [V(\mathbf{q}(t+1)) - V(\mathbf{q}(t))] \mathcal{I}(\mathbf{q}(t) = \mathbf{q}) \\
 \Delta W_{\perp}(\mathbf{q}) &\triangleq [W_{\perp}(\mathbf{q}(t+1)) - W_{\perp}(\mathbf{q}(t))] \mathcal{I}(\mathbf{q}(t) = \mathbf{q}) \\
 \Delta V_{\perp}(\mathbf{q}) &\triangleq [V_{\perp}(\mathbf{q}(t+1)) - V_{\perp}(\mathbf{q}(t))] \mathcal{I}(\mathbf{q}(t) = \mathbf{q}) \\
 \Delta V_{\parallel}(\mathbf{q}) &\triangleq [V_{\parallel}(\mathbf{q}(t+1)) - V_{\parallel}(\mathbf{q}(t))] \mathcal{I}(\mathbf{q}(t) = \mathbf{q}).
 \end{aligned}$$

We will use Lemma 3 using the Lyapunov function $W_{\perp}(\mathbf{q})(\cdot)$ to bound the $\bar{\mathbf{q}}_{\perp}^{(\epsilon)}$ component in steady state. We need the following lemma, which follows from concavity of square root function and the pythagorean theorem (5). The proof of this lemma is similar to the proof of Lemma 7 in [7] and so we skip it here.

LEMMA 4. *Drift of $W_{\perp}(\cdot)$ can be bounded in terms of drift of $V(\cdot)$ and $V_{\parallel}(\cdot)$ as follows.*

$$\Delta W_{\perp}(\mathbf{q}) \leq \frac{1}{2\|\mathbf{q}_{\perp}\|} (\Delta V(\mathbf{q}) - \Delta V_{\parallel}(\mathbf{q})) \quad \forall \mathbf{q} \in \mathbb{R}^{n^2}.$$

We will now formally state the state space collapse result.

PROPOSITION 2. Consider a set of switch systems under MaxWeight scheduling algorithm, with the arrival processes $\mathbf{a}^{(\epsilon)}(t)$ described in Section 2.1, parameterized by $0 < \epsilon < 1$, such that the mean arrival rate vector is $\boldsymbol{\lambda}^\epsilon = (1 - \epsilon)\boldsymbol{\nu}$ for some $\boldsymbol{\nu} \in \mathcal{F}$ such that $\nu_{\min} \triangleq \min_{ij} \nu_{ij} > 0$. The load is then $\rho = (1 - \epsilon)$. Let the variance of the arrival process be $(\boldsymbol{\sigma}^{(\epsilon)})^2$. 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 $0 < \epsilon \leq \nu_{\min}/2\|\boldsymbol{\nu}\|$, the steady state queue lengths vector satisfies

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

where

$$M_r^{(\epsilon)} = 2^{\frac{1}{r}} \max \left(\frac{4(\|\boldsymbol{\lambda}^{(\epsilon)}\|^2 + \|\boldsymbol{\sigma}^{(\epsilon)}\|^2 + n)}{\nu_{\min}}, (\sqrt{r}e)^{1/r} 16^{\frac{r}{e}} \frac{na_{\max}}{\nu_{\min}} (na_{\max} + 1) \right).$$

REMARK 1. Note that for any r , the expressions M_r are upper bounded by a constant not dependant on ϵ whenever there exists a $\tilde{\sigma}$ which does not depend on ϵ such that $\|\boldsymbol{\sigma}^{(\epsilon)}\|^2 \leq \tilde{\sigma}$ for all ϵ . This is why we call this state space collapse. Our notion of state-space collapse considers the system in steady-state, and is hence mathematically different from the state-space collapse result in [5], although the results are similar in spirit.

PROOF. We will skip the superscript (ϵ) in this proof for ease of notation. Thus, we will use $\mathbf{q}(t)$, $\boldsymbol{\lambda}$ and $\boldsymbol{\sigma}$ to denote $\mathbf{q}^{(\epsilon)}(t)$, $\boldsymbol{\lambda}^{(\epsilon)}$ and $\boldsymbol{\sigma}^{(\epsilon)}$ respectively. We will verify both the conditions C.1 and C.2 to apply Lemma 3 for the Markov chain $\mathbf{q}(t)$ and Lyapunov function $W_{\perp}(\mathbf{q}(\cdot))$. First we consider condition C.2.

$$\begin{aligned} |\Delta W_{\perp}(\mathbf{q})| &= \left| \|\mathbf{q}_{\perp}(t+1)\| - \|\mathbf{q}_{\perp}(t)\| \right| \mathcal{I}(\mathbf{q}(t) = \mathbf{q}) \\ &\stackrel{(a)}{\leq} \|\mathbf{q}_{\perp}(t+1) - \mathbf{q}_{\perp}(t)\| \\ &\stackrel{(b)}{\leq} \|\mathbf{q}(t+1) - \mathbf{q}(t)\| \\ &= \sqrt{\sum_{ij} |q_{ij}(t+1) - q_{ij}(t)|^2} \\ &\stackrel{(c)}{\leq} \sqrt{\sum_{ij} a_{\max}^2} \\ (9) \quad &\leq na_{\max}, \end{aligned}$$

where (a) follows from triangle inequality, i.e., $\|\mathbf{x}\| - \|\mathbf{y}\| \leq \|\mathbf{x} - \mathbf{y}\|$ and $\mathcal{I}(\cdot) \leq 1$; (b) follows from nonexpansivity of projection operator (6); (c) is true because each queue lengths can increase by at most $a_{\max} \geq 1$ due to arrivals and can decrease by at most 1 due to departures. Thus condition C.2 of Lemma is true with $D = na_{\max}$.

We will now verify C.1, using Lemma 4 by bounding the drifts $\Delta V(\mathbf{q})$ and $\Delta V_{\parallel}(\mathbf{q})$.

$$\begin{aligned}
& \mathbb{E} [\Delta V(\mathbf{q}) | \mathbf{q}(t) = \mathbf{q}] \\
&= \mathbb{E} [\|\mathbf{q}(t+1)\|^2 - \|\mathbf{q}(t)\|^2 | \mathbf{q}(t) = \mathbf{q}] \\
&= \mathbb{E} [\|\mathbf{q}(t) + \mathbf{a}(t) - \mathbf{s}(t) + \mathbf{u}(t)\|^2 - \|\mathbf{q}(t)\|^2 | \mathbf{q}(t) = \mathbf{q}] \\
&= \mathbb{E} [\|\mathbf{q}(t) + \mathbf{a}(t) - \mathbf{s}(t)\|^2 + \|\mathbf{u}(t)\|^2 - \|\mathbf{q}(t)\|^2 | \mathbf{q}(t) = \mathbf{q}] \\
&\quad + \mathbb{E} [2 \langle \mathbf{q}(t+1) - \mathbf{u}(t), \mathbf{u}(t) \rangle | \mathbf{q}(t) = \mathbf{q}] \\
&\stackrel{(a)}{\leq} \mathbb{E} [\|\mathbf{a}(t) - \mathbf{s}(t)\|^2 + 2 \langle \mathbf{q}(t), \mathbf{a}(t) - \mathbf{s}(t) \rangle | \mathbf{q}(t) = \mathbf{q}] \\
&\stackrel{(b)}{=} \mathbb{E} \left[\sum_{ij} (a_{ij}^2(t) + s_{ij}(t) - 2a_{ij}(t)s_{ij}(t)) \middle| \mathbf{q}(t) = \mathbf{q} \right] \\
&\quad + 2 \langle \mathbf{q}, \boldsymbol{\lambda} - \mathbb{E}[\mathbf{s}(t) | \mathbf{q}(t) = \mathbf{q}] \rangle \\
&\stackrel{(c)}{=} \sum_{ij} (\lambda_{ij}^2 + \sigma_{ij}^2) + n - 2 \mathbb{E} \left[\sum_{ij} \lambda_{ij} s_{ij}(t) \middle| \mathbf{q}(t) = \mathbf{q} \right] \\
(10) \quad &\quad + 2 \langle \mathbf{q}, \boldsymbol{\lambda} - \mathbb{E}[\mathbf{s}(t) | \mathbf{q}(t) = \mathbf{q}] \rangle \\
&= \|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\sigma}\|^2 + n - 2(1 - \epsilon) \mathbb{E} \left[\sum_{ij} \nu_{ij} s_{ij}(t) \middle| \mathbf{q}(t) = \mathbf{q} \right] \\
&\quad + 2 \langle \mathbf{q}, (1 - \epsilon)\boldsymbol{\nu} - \mathbb{E}[\mathbf{s}(t) | \mathbf{q}(t) = \mathbf{q}] \rangle \\
&\leq \|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\sigma}\|^2 + n - 2\epsilon \langle \mathbf{q}, \boldsymbol{\nu} \rangle + 2 \langle \mathbf{q}, \boldsymbol{\nu} - \mathbb{E}[\mathbf{s}(t) | \mathbf{q}(t) = \mathbf{q}] \rangle \\
(11) \quad &= \|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\sigma}\|^2 + n - 2\epsilon \langle \mathbf{q}, \boldsymbol{\nu} \rangle + 2 \min_{\mathbf{r} \in \mathcal{C}} \langle \mathbf{q}, \boldsymbol{\nu} - \mathbf{r} \rangle,
\end{aligned}$$

where (a) follows from the fact that $\langle \mathbf{q}(t+1), \mathbf{u}(t) \rangle = 0$ and dropping the $-\|\mathbf{u}(t)\|^2$ term; (b) is true because $s_{ij} \in \{0, 1\}$. Note that $\mathbb{E}[a_{ij}^2(t)] = \mathbb{E}[a_{ij}(t)]^2 + \text{Var}(a_{ij}(t))$. Also note that the arrivals in each time slot are independent of the queue lengths and hence are also independent of the service process. These facts and (2) give (c). Since we use MaxWeight scheduling algorithm, from (1), we have (11). In order to bound the last term in (11), we present the following claim.

CLAIM 2. For any $\mathbf{q} \in \mathbb{R}^{n^2}$,

$$\boldsymbol{\nu} + \frac{\nu_{\min}}{\|\mathbf{q}_{\perp}\|} \mathbf{q}_{\perp} \in \mathcal{C}.$$

PROOF. Since $|q_{\perp ij}| \leq \|\mathbf{q}_{\perp}\|$, $\nu_{ij} + \frac{\nu_{\min}}{\|\mathbf{q}_{\perp}\|} q_{\perp ij} \geq \nu_{ij} - \nu_{\min} \geq 0$ and so $\boldsymbol{\nu} + \frac{\nu_{\min}}{\|\mathbf{q}_{\perp}\|} \mathbf{q}_{\perp} \in \mathbb{R}_+^{n^2}$. We know that $\mathbf{q}_{\perp} \in \mathcal{K}^{\circ}$ and $\mathbf{e}^{(i)} \in \mathcal{K}$, and so $\langle \mathbf{q}_{\perp}, \mathbf{e}^{(i)} \rangle \leq 0$. Thus, for any i , we have

$$\begin{aligned} \left\langle \boldsymbol{\nu} + \frac{\nu_{\min}}{\|\mathbf{q}_{\perp}\|} \mathbf{q}_{\perp}, \mathbf{e}^{(i)} \right\rangle &= \langle \boldsymbol{\nu}, \mathbf{e}^{(i)} \rangle + \frac{\nu_{\min}}{\|\mathbf{q}_{\perp}\|} \langle \mathbf{q}_{\perp}, \mathbf{e}^{(i)} \rangle \\ &\leq \langle \boldsymbol{\nu}, \mathbf{e}^{(i)} \rangle \\ &= 1, \end{aligned}$$

where the last equality is due to the fact that $\boldsymbol{\nu} \in \mathcal{F}$. Similarly, we can show that $\left\langle \boldsymbol{\nu} + \frac{\nu_{\min}}{\|\mathbf{q}_{\perp}\|} \mathbf{q}_{\perp}, \tilde{\mathbf{e}}^{(j)} \right\rangle \leq 1$ for any j , proving the claim. \square

Using the claim in (11), we get

$$\begin{aligned} &\mathbb{E} [\Delta V(\mathbf{q}) | \mathbf{q}(t) = \mathbf{q}] \\ &\leq \|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\sigma}\|^2 + n - 2\epsilon \langle \mathbf{q}, \boldsymbol{\nu} \rangle + 2 \left\langle \mathbf{q}, \boldsymbol{\nu} - \left(\boldsymbol{\nu} + \frac{\nu_{\min}}{\|\mathbf{q}_{\perp}\|} \mathbf{q}_{\perp} \right) \right\rangle \\ &= \|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\sigma}\|^2 + n - 2\epsilon \langle \mathbf{q}, \boldsymbol{\nu} \rangle - \frac{2\nu_{\min}}{\|\mathbf{q}_{\perp}\|} \langle \mathbf{q}_{\parallel} + \mathbf{q}_{\perp}, \mathbf{q}_{\perp} \rangle \\ (12) \quad &= \|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\sigma}\|^2 + n - 2\epsilon \langle \mathbf{q}, \boldsymbol{\nu} \rangle - 2\nu_{\min} \|\mathbf{q}_{\perp}\|, \end{aligned}$$

where the last equality follows from the fact that $\langle \mathbf{q}_{\parallel}, \mathbf{q}_{\perp} \rangle = 0$. We will now bound the drift $\Delta V_{\parallel}(\mathbf{q})$.

$$\begin{aligned} &\mathbb{E} [\Delta V_{\parallel}(\mathbf{q}) | \mathbf{q}(t) = \mathbf{q}] \\ &= \mathbb{E} [\|\mathbf{q}_{\parallel}(t+1)\|^2 - \|\mathbf{q}_{\parallel}(t)\|^2 | \mathbf{q}(t) = \mathbf{q}] \\ &= \mathbb{E} [\langle \mathbf{q}_{\parallel}(t+1) + \mathbf{q}_{\parallel}(t), \mathbf{q}_{\parallel}(t+1) - \mathbf{q}_{\parallel}(t) \rangle | \mathbf{q}(t) = \mathbf{q}] \\ &= \mathbb{E} [\|\mathbf{q}_{\parallel}(t+1) - \mathbf{q}_{\parallel}(t)\|^2 | \mathbf{q}(t) = \mathbf{q}] \\ &\quad + 2\mathbb{E} [\langle \mathbf{q}_{\parallel}(t), \mathbf{q}_{\parallel}(t+1) - \mathbf{q}_{\parallel}(t) \rangle | \mathbf{q}(t) = \mathbf{q}] \\ &\geq 2\mathbb{E} [\langle \mathbf{q}_{\parallel}(t), \mathbf{q}_{\parallel}(t+1) - \mathbf{q}_{\parallel}(t) \rangle | \mathbf{q}(t) = \mathbf{q}] \\ &= 2\mathbb{E} [\langle \mathbf{q}_{\parallel}(t), \mathbf{q}(t+1) - \mathbf{q}(t) \rangle | \mathbf{q}(t) = \mathbf{q}] \\ &\quad - 2\mathbb{E} [\langle \mathbf{q}_{\parallel}(t), \mathbf{q}_{\perp}(t+1) - \mathbf{q}_{\perp}(t) \rangle | \mathbf{q}(t) = \mathbf{q}] \\ &\stackrel{(a)}{\geq} 2\mathbb{E} [\langle \mathbf{q}_{\parallel}(t), \mathbf{a}(t) - \mathbf{s}(t) + \mathbf{u}(t) \rangle | \mathbf{q}(t) = \mathbf{q}] \end{aligned}$$

$$\begin{aligned}
& \stackrel{(b)}{\geq} 2 \langle \mathbf{q}_{\parallel}, \boldsymbol{\lambda} \rangle - 2\mathbb{E} [\langle \mathbf{q}_{\parallel}, \mathbf{s}(t) \rangle | \mathbf{q}(t) = \mathbf{q}] \\
& = -2\epsilon \langle \mathbf{q}_{\parallel}, \boldsymbol{\nu} \rangle - 2\mathbb{E} [\langle \mathbf{q}_{\parallel}, \mathbf{s}(t) - \boldsymbol{\nu} \rangle | \mathbf{q}(t) = \mathbf{q}] \\
(13) \quad & = -2\epsilon \langle \mathbf{q}_{\parallel}, \boldsymbol{\nu} \rangle.
\end{aligned}$$

Equation (a) is true because $\langle \mathbf{q}_{\parallel}(t), \mathbf{q}_{\perp}(t) \rangle = 0$ and $\langle \mathbf{q}_{\parallel}(t), \mathbf{q}_{\perp}(t+1) \rangle \leq 0$ since $\mathbf{q}_{\parallel}(t) \in \mathcal{K}$ and $\mathbf{q}_{\perp}(t+1) \in \mathcal{K}^{\circ}$. All the components of \mathbf{q}_{\parallel} and $\mathbf{u}(t)$ are nonnegative. Using this fact with independence of the arrivals and the queue lengths gives Equation (b). The last equality follows from (4) since $\mathbf{q}_{\parallel} \in \mathcal{K} \in \mathcal{V}_{\mathcal{K}}$ and $\mathbf{s}(t), \boldsymbol{\nu} \in \mathcal{F}$ from (2). Now substituting (12) and (13) in Lemma 4, we get

$$\begin{aligned}
& \mathbb{E} [\Delta W_{\perp}(\mathbf{q}) | \mathbf{q}(t) = \mathbf{q}] \\
& \leq \frac{1}{2\|\mathbf{q}_{\perp}\|} (\|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\sigma}\|^2 + n - 2\epsilon \langle \mathbf{q}, \boldsymbol{\nu} \rangle - 2\nu_{\min}\|\mathbf{q}_{\perp}\| + 2\epsilon \langle \mathbf{q}_{\parallel}, \boldsymbol{\nu} \rangle) \\
& = \frac{\|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\sigma}\|^2 + n}{2\|\mathbf{q}_{\perp}\|} - \nu_{\min} - \frac{\epsilon}{\|\mathbf{q}_{\perp}\|} \langle \mathbf{q}_{\perp}, \boldsymbol{\nu} \rangle \\
& \stackrel{(a)}{\leq} \frac{\|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\sigma}\|^2 + n}{2\|\mathbf{q}_{\perp}\|} - \nu_{\min} + \epsilon\|\boldsymbol{\nu}\| \\
& \leq \frac{\|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\sigma}\|^2 + n}{2\|\mathbf{q}_{\perp}\|} - \frac{\nu_{\min}}{2} \text{ whenever } \epsilon \leq \frac{\nu_{\min}}{2\|\boldsymbol{\nu}\|} \\
& \leq -\frac{\nu_{\min}}{4} \text{ for all } \mathbf{q} \text{ such that } W_{\perp}(\mathbf{q}) \geq \frac{2(\|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\sigma}\|^2 + n)}{\nu_{\min}},
\end{aligned}$$

where (a) is due to the Cauchy Schwartz inequality $\langle \frac{-\mathbf{q}_{\perp}}{\|\mathbf{q}_{\perp}\|}, \boldsymbol{\nu} \rangle \leq \frac{\|\mathbf{q}_{\perp}\|}{\|\mathbf{q}_{\perp}\|} \|\boldsymbol{\nu}\|$. Thus condition C.1 is valid with $\kappa = \frac{2(\|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\sigma}\|^2 + n)}{\nu_{\min}}$ and $\eta = \frac{\nu_{\min}}{4}$. Then from Lemma 3, we get for $r = 1, 2, \dots$,

$$\begin{aligned}
\mathbb{E} \left[\|\bar{\mathbf{q}}_{\perp}^{(\epsilon)}\|^r \right] & \leq \left(\frac{4(\|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\sigma}\|^2 + n)}{\nu_{\min}} \right)^r + r! \left(16 \frac{na_{\max}}{\nu_{\min}} \right)^r \left(na_{\max} + \frac{\nu_{\min}}{4} \right)^r \\
& \stackrel{(a)}{\leq} \left(\frac{4(\|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\sigma}\|^2 + n)}{\nu_{\min}} \right)^r + \sqrt{r}e \left(16 \frac{r}{e} \frac{na_{\max}}{\nu_{\min}} (na_{\max} + 1) \right)^r \\
& \leq 2 \max \left(\frac{4(\|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\sigma}\|^2 + n)}{\nu_{\min}}, (\sqrt{r}e)^{1/r} 16 \frac{r}{e} \frac{na_{\max}}{\nu_{\min}} (na_{\max} + 1) \right)^r.
\end{aligned}$$

where (a) follows from Stirling's upper bound of the factorial function, $r! \leq e^{1-r} r^{r+\frac{1}{2}}$ and noting that $\nu_{\min} \leq 1$ follows from the definition of ν_{\min} and the capacity region \mathcal{C} . The last inequality follows from $a^r + b^r \leq 2 \max(a, b)^r$, proving the proposition. \square

Recall that there are $n!$ maximal schedules (perfect matchings or permutations). For each of them, MaxWeight assigns a weight which is the sum of corresponding queue lengths and then picks the one with the maximum weight. In this process, it is equalizing the weights of all the schedules by serving the matching with maximum weight and thereby decreasing it. The cone \mathcal{K} has the property that if the queue lengths vector \mathbf{q} is in the cone \mathcal{K} , we have w_i and \tilde{w}_j such that $q_{ij} = w_i + \tilde{w}_j$. This means that all the maximal schedules have the same weight $\sum_i w_i + \sum_j \tilde{w}_j$ and the MaxWeight algorithm is agnostic between them. Thus, the state space collapse result states that in steady state, MaxWeight is (almost) successful in being able to equalize the weights of all maximal schedules in the heavy traffic limit. This behavior is very similar to Join-the-shortest queue (JSQ) routing policy in a supermarket checkout system. In such a system, there are a few servers, each with a queue. When a customer arrives to be served, under JSQ policy, (s)he picks the server with the shortest queue. It was shown in [7] that in the heavy traffic limit, the state of this system collapses to a state where all the queues are equal, and thus, JSQ is agnostic between all the queues when such a state space collapse occurs. Here JSQ policy is trying to equalize all the queues by increasing the shortest one, and it is (almost) successful in doing that in steady state in heavy traffic limit.

A natural question in this context is if there is any interpretation to the variables w_i and \tilde{w}_j . These variables are the optimal dual variables for the maximum weight matching problem. The maximum weighted perfect matching problem in bipartite graphs (that MaxWeight solves in every time slot) can be written as the integer program (14) and its linear program (LP) relaxation is the linear program (15).

$$\begin{aligned}
 (14) \quad & \max \sum_{ij} q_{ij} s_{ij} \\
 & \text{subject to: } \sum_i s_{ij} = 1 \quad \forall j \\
 & \quad \quad \sum_j s_{ij} = 1 \quad \forall i \\
 & \quad \quad s_{ij} \in \{0, 1\} \quad \forall i, j.
 \end{aligned}
 \qquad
 \begin{aligned}
 (15) \quad & \max \sum_{ij} q_{ij} s_{ij} \\
 & \text{subject to: } \sum_i s_{ij} = 1 \quad \forall j \\
 & \quad \quad \sum_j s_{ij} = 1 \quad \forall i \\
 & \quad \quad s_{ij} \geq 0 \quad \forall i, j.
 \end{aligned}$$

It can be proved that the optimal solution of the LP relaxation (15) is identical to the optimal solution of the original integer program (14) [23]. The dual of the LP (15) is the following.

$$\min \sum_i w_i + \sum_j \tilde{w}_j$$

$$(16) \quad \text{subject to: } w_i + \tilde{w}_j \geq q_{ij} \quad \forall \quad i, j.$$

For any perfect matching π and its corresponding schedule s_{ij} , and for any dual feasible w_i, \tilde{w}_j , we have that $\sum_i q_{i\pi(i)} = \sum_{ij} q_{ij} s_{ij} \leq \sum_i w_i + \sum_j \tilde{w}_j$. Suppose s_{ij}^* is an optimal solution of (14) and corresponds to a permutation π^* , and suppose w_i^*, \tilde{w}_j^* is an optimal solution of (16). Then, from strong duality, we have that

$$(17) \quad \sum_i q_{i\pi^*(i)} = \sum_{ij} q_{ij} s_{ij}^* = \sum_i w_i^* + \sum_j \tilde{w}_j^*.$$

Moreover, any π^* and w_i^*, \tilde{w}_j^* that satisfy (17) are optimal solutions for problems (14) (or (15)) and (16) respectively. This means that any optimal perfect matching consists of only links (i, j) such that $q_{ij} = w_i + \tilde{w}_j$. This property is also called complementary slackness. The Hungarian assignment algorithm for solving the MaxWeight matching problem is based on this property. The cone \mathcal{K} , has the special property that if w_i, \tilde{w}_j is the optimal solution, then for any (i, j) , we have $q_{ij} = w_i + \tilde{w}_j$ and so any perfect matching is an optimal matching and all perfect matchings have the same weight.

5. Asymptotically tight upper and lower bounds under the Max-Weight policy. In the previous section, we have shown that the queue length vector collapses within the cone \mathcal{K} in the steady state. We will use this result to obtain lower and upper bounds on the average queue lengths under MaxWeight algorithm. The lower and upper bounds differ only in $o(1/\epsilon)$ and so match in the heavy traffic limit.

We will obtain these bounds by equating the drift of certain carefully chosen functions equal to zero in steady-state. We first define a few Lyapunov-type functions and their drifts, in addition to the already defined $V(\mathbf{q}) = \|\mathbf{q}\|^2$. The following lemma states that all these Lyapunov functions have finite expectations in steady state.

$$\begin{aligned} V_1(\mathbf{q}) &\triangleq \sum_i \left(\sum_j q_{ij} \right)^2 & V_2(\mathbf{q}) &\triangleq \sum_j \left(\sum_i q_{ij} \right)^2 & V_3(\mathbf{q}) &\triangleq \left(\sum_{ij} q_{ij} \right)^2 \\ \Delta V_1(\mathbf{q}) &\triangleq [V_1(\mathbf{q}(t+1)) - V_1(\mathbf{q}(t))] \mathcal{I}(\mathbf{q}(t) = \mathbf{q}) \\ \Delta V_2(\mathbf{q}) &\triangleq [V_2(\mathbf{q}(t+1)) - V_2(\mathbf{q}(t))] \mathcal{I}(\mathbf{q}(t) = \mathbf{q}) \\ \Delta V_3(\mathbf{q}) &\triangleq [V_3(\mathbf{q}(t+1)) - V_3(\mathbf{q}(t))] \mathcal{I}(\mathbf{q}(t) = \mathbf{q}). \end{aligned}$$

LEMMA 5. Consider the switch under MaxWeight scheduling algorithm. For any arrival rate vector λ in the interior of the capacity region $\lambda \in \text{int}(\mathcal{C})$, the steady state means $\mathbb{E}[V(\bar{\mathbf{q}})]$, $\mathbb{E}[V_1(\bar{\mathbf{q}})]$, $\mathbb{E}[V_2(\bar{\mathbf{q}})]$ and $\mathbb{E}[V_3(\bar{\mathbf{q}})]$ are finite.

The lemma is proved in Appendix C. We will now state and prove the main result of this paper.

THEOREM 1. Consider a set of switch systems under the MaxWeight scheduling algorithm, with arrival processes $\mathbf{a}^{(\epsilon)}(t)$ described in Section 2.1, parameterized by $0 < \epsilon < 1$, such that the mean arrival rate vector is $\lambda^\epsilon = (1 - \epsilon)\boldsymbol{\nu}$ for some $\boldsymbol{\nu} \in \mathcal{F}$ such that $\nu_{\min} \triangleq \min_{ij} \nu_{ij} > 0$. The load is then $\rho = (1 - \epsilon)$. Let the variance of the arrival process be $(\boldsymbol{\sigma}^{(\epsilon)})^2$. The queue length process $\mathbf{q}^{(\epsilon)}(t)$ for each system converges in distribution to the steady state random vector $\bar{\mathbf{q}}^{(\epsilon)}$. For each system with $0 < \epsilon \leq \nu_{\min}/2\|\boldsymbol{\nu}\|$, the steady state average queue length satisfies

$$\left(1 - \frac{1}{2n}\right) \frac{\|\boldsymbol{\sigma}^{(\epsilon)}\|^2}{\epsilon} - B_1(\epsilon, n) \leq \mathbb{E} \left[\sum_{ij} \bar{q}_{ij}^{(\epsilon)} \right] \leq \left(1 - \frac{1}{2n}\right) \frac{\|\boldsymbol{\sigma}^{(\epsilon)}\|^2}{\epsilon} + B_2(\epsilon, n),$$

where

$$B_1(\epsilon, n) = -\frac{n\epsilon}{2} + n + 3n^{(2-\frac{1}{r})} \epsilon^{(-\frac{1}{r})} M_r^{(\epsilon)} \quad \text{and}$$

$$B_2(\epsilon, n) = \frac{n(1+\epsilon)}{2} + 2n^{(2-\frac{1}{r})} \epsilon^{(-\frac{1}{r})} M_r^{(\epsilon)}$$

for any $r \in \{2, 3, \dots\}$. The terms $B_1(\epsilon, n)$ and $B_2(\epsilon, n)$ are both $o(\frac{1}{\epsilon})$, i.e., $\lim_{\epsilon \downarrow 0} \epsilon B_1(\epsilon, n) = 0$ and $\lim_{\epsilon \downarrow 0} \epsilon B_2(\epsilon, n) = 0$. Therefore, in the heavy traffic limit as $\epsilon \downarrow 0$ which means as the mean arrival rate $\lambda^\epsilon \rightarrow \frac{1}{n}\mathbf{1}$, if $(\boldsymbol{\sigma}^{(\epsilon)})^2 \rightarrow \boldsymbol{\sigma}^2$, we have

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

PROOF. Fix an $0 < \epsilon \leq \nu_{\min}/2\|\boldsymbol{\nu}\|$ and we consider the system with index ϵ . For simplicity of notation, we again skip the superscript (ϵ) in this proof and use $\bar{\mathbf{q}}$ to denote the steady state queue length vector. We will use $\bar{\mathbf{a}}$ to denote the arrival vector in steady state, which is identically distributed to the random vector $\mathbf{a}(t)$ for any time t . We will use $\mathbf{s}(\bar{\mathbf{q}})$ and $\mathbf{u}(\bar{\mathbf{q}})$ to denote the schedule and unused service to show their dependence on the queue lengths. We will use $\bar{\mathbf{q}}^+$ to denote $\bar{\mathbf{q}} + \bar{\mathbf{a}} - \mathbf{s}(\bar{\mathbf{q}}) + \mathbf{u}(\bar{\mathbf{q}})$, which is the queue

lengths vector at time $t + 1$ if it was $\bar{\mathbf{q}}$ at time t . Clearly, $\bar{\mathbf{q}}^+$ and $\bar{\mathbf{q}}$ have the same distribution.

Define a new function $V_4(\mathbf{q})$ and its drift as follows.

$$\begin{aligned} V_4(\mathbf{q}) &= V_1(\mathbf{q}) + V_2(\mathbf{q}) - \frac{1}{n}V_3(\mathbf{q}) \\ &= \sum_i \left(\sum_j q_{ij} \right)^2 + \sum_j \left(\sum_i q_{ij} \right)^2 - \frac{1}{n} \left(\sum_{ij} q_{ij} \right)^2 \\ \Delta V_4(\mathbf{q}) &\triangleq [V_4(\mathbf{q}(t+1)) - V_4(\mathbf{q}(t))] \mathcal{I}(\mathbf{q}(t) = \mathbf{q}) \\ &= \Delta V_1(\mathbf{q}) + \Delta V_2(\mathbf{q}) - \frac{1}{n} \Delta V_3(\mathbf{q}). \end{aligned}$$

Since $-\frac{1}{n}V_3(\mathbf{q}) \leq V_4(\mathbf{q}) \leq V_1(\mathbf{q}) + V_2(\mathbf{q})$, the steady state mean $\mathbb{E}[V_4(\bar{\mathbf{q}})]$ is finite from Lemma 5. Therefore, the mean drift of $V_4(\cdot)$ in steady state is zero, i.e.,

$$\begin{aligned} \mathbb{E}[\Delta V_4(\bar{\mathbf{q}})] &= \mathbb{E}[[V_4(\mathbf{q}(t+1)) - V_4(\mathbf{q}(t))] \mathcal{I}(\mathbf{q}(t) = \bar{\mathbf{q}})] \\ &= \mathbb{E}[V_4(\bar{\mathbf{q}}^+) - V_4(\bar{\mathbf{q}})] \\ &= \mathbb{E}[V_4(\bar{\mathbf{q}})] - \mathbb{E}[V_4(\bar{\mathbf{q}})] \\ (18) \quad &= 0. \end{aligned}$$

Therefore,

$$(19) \quad \mathbb{E}[\Delta V_1(\bar{\mathbf{q}})] + \mathbb{E}[\Delta V_2(\bar{\mathbf{q}})] - \frac{1}{n} \mathbb{E}[\Delta V_3(\bar{\mathbf{q}})] = 0.$$

Expanding the drift of $V_1(\cdot)$, we get

$$\begin{aligned} &\mathbb{E}[\Delta V_1(\bar{\mathbf{q}})] \\ &= \mathbb{E}[V_1(\bar{\mathbf{q}} + \bar{\mathbf{a}} - \mathbf{s}(\bar{\mathbf{q}}) + \mathbf{u}(\bar{\mathbf{q}})) - V_1(\bar{\mathbf{q}})] \\ &= \mathbb{E} \left[\sum_i \left(\sum_j (\bar{q}_{ij} + \bar{a}_{ij} - s_{ij}(\bar{\mathbf{q}}) + u_{ij}(\bar{\mathbf{q}})) \right)^2 - \sum_i \left(\sum_j \bar{q}_{ij} \right)^2 \right] \\ &= \mathbb{E} \left[\sum_i \left(\sum_j (\bar{a}_{ij} - s_{ij}(\bar{\mathbf{q}})) \right)^2 + \sum_i \left(\sum_j u_{ij}(\bar{\mathbf{q}}) \right)^2 \right] \\ &\quad + 2\mathbb{E} \left[\sum_i \left(\sum_j (\bar{q}_{ij} + \bar{a}_{ij} - s_{ij}(\bar{\mathbf{q}})) \right) \left(\sum_{j'} u_{ij'}(\bar{\mathbf{q}}) \right) \right] \end{aligned}$$

$$\begin{aligned}
& + 2\mathbb{E} \left[\sum_i \left(\sum_j \bar{q}_{ij} \right) \left(\sum_{j'} (\bar{a}_{ij'} - s_{ij'}(\bar{\mathbf{q}})) \right) \right] \\
& = \mathbb{E} \left[\sum_i \left(\sum_j (\bar{a}_{ij} - s_{ij}(\bar{\mathbf{q}})) \right)^2 + 2 \sum_i \left(\sum_j \bar{q}_{ij}^+ \right) \left(\sum_{j'} u_{ij'}(\bar{\mathbf{q}}) \right) \right] \\
& + \mathbb{E} \left[- \sum_i \left(\sum_j u_{ij}(\bar{\mathbf{q}}) \right)^2 + 2 \sum_i \left(\sum_j \bar{q}_{ij} \right) \left(\sum_{j'} (\bar{a}_{ij'} - s_{ij'}(\bar{\mathbf{q}})) \right) \right].
\end{aligned}$$

Similarly expanding drifts of $V_2(\cdot)$ and $V_3(\cdot)$ and substituting in (19), we get the following expression. Since this is a lengthy equation, we split into various terms which we denote by $\mathcal{T}_1, \mathcal{T}_2, \mathcal{T}_3$ and \mathcal{T}_4 . For simplicity of notation, we suppress all the dependencies in terms of $\bar{\mathbf{q}}, \bar{\mathbf{a}}, \mathbf{s}(\bar{\mathbf{q}})$ and $\mathbf{u}(\bar{\mathbf{q}})$.

$$(20) \quad \mathcal{T}_1 = \mathcal{T}_2 + \mathcal{T}_3 + \mathcal{T}_4,$$

where

$$\begin{aligned}
\mathcal{T}_1 & = 2\mathbb{E} \left[\sum_i \left(\sum_j \bar{q}_{ij} \right) \left(\sum_{j'} (s_{ij'}(\bar{\mathbf{q}}) - \bar{a}_{ij'}) \right) \right] \\
& + 2\mathbb{E} \left[\sum_j \left(\sum_i \bar{q}_{ij} \right) \left(\sum_{i'} (s_{i'j}(\bar{\mathbf{q}}) - \bar{a}_{i'j}) \right) \right] \\
& - \frac{2}{n} \mathbb{E} \left[\left(\sum_{ij} \bar{q}_{ij} \right) \left(\sum_{i'j'} (s_{i'j'}(\bar{\mathbf{q}}) - \bar{a}_{i'j'}) \right) \right] \\
\mathcal{T}_2 & = \mathbb{E} \left[\sum_i \left(\sum_j (\bar{a}_{ij} - s_{ij}(\bar{\mathbf{q}})) \right)^2 \right] + \mathbb{E} \left[\sum_j \left(\sum_i (\bar{a}_{ij} - s_{ij}(\bar{\mathbf{q}})) \right)^2 \right] \\
& - \frac{1}{n} \mathbb{E} \left[\left(\sum_{ij} (\bar{a}_{ij} - s_{ij}(\bar{\mathbf{q}})) \right)^2 \right] \\
\mathcal{T}_3 & = - \mathbb{E} \left[\sum_i \left(\sum_j u_{ij}(\bar{\mathbf{q}}) \right)^2 \right] - \mathbb{E} \left[\sum_j \left(\sum_i u_{ij}(\bar{\mathbf{q}}) \right)^2 \right] \\
& + \frac{1}{n} \mathbb{E} \left[\left(\sum_{ij} u_{ij}(\bar{\mathbf{q}}) \right)^2 \right]
\end{aligned}$$

$$\begin{aligned} \mathcal{T}_4 = & 2\mathbb{E} \left[\sum_i \left(\sum_j \bar{q}_{ij}^+ \right) \left(\sum_{j'} u_{ij'}(\bar{\mathbf{q}}) \right) \right] + 2\mathbb{E} \left[\sum_j \left(\sum_i \bar{q}_{ij}^+ \right) \left(\sum_{i'} u_{i'j}(\bar{\mathbf{q}}) \right) \right] \\ & - \frac{2}{n} \mathbb{E} \left[\left(\sum_{ij} \bar{q}_{ij}^+ \right) \left(\sum_{i'j'} u_{i'j'}(\bar{\mathbf{q}}) \right) \right]. \end{aligned}$$

We will now bound each of the four terms. The schedule in each time slot is maximal (2) and so $\sum_i s_{ij} = 1, \sum_j s_{ij} = 1$ and $\sum_{ij} s_{ij} = n$. Noting that the arrivals are independent of queue lengths, we can simplify the term \mathcal{T}_1 as follows.

$$\begin{aligned} \mathcal{T}_1 = & 2\mathbb{E} \left[\sum_i \left(\sum_j \bar{q}_{ij} \right) \left(1 - \sum_{j'} \lambda_{ij'} \right) \right] + 2\mathbb{E} \left[\sum_j \left(\sum_i \bar{q}_{ij} \right) \left(1 - \sum_{i'} \lambda_{i'j} \right) \right] \\ & - \frac{2}{n} \mathbb{E} \left[\left(\sum_{ij} \bar{q}_{ij} \right) \left(n - \sum_{i'j'} \lambda_{i'j'} \right) \right] \\ \stackrel{(a)}{=} & 2\mathbb{E} \left[\sum_i \epsilon \left(\sum_j \bar{q}_{ij} \right) \right] + 2\mathbb{E} \left[\sum_j \epsilon \left(\sum_i \bar{q}_{ij} \right) \right] - \frac{2}{n} \mathbb{E} \left[n\epsilon \left(\sum_{ij} \bar{q}_{ij} \right) \right] \\ = & 2\epsilon \mathbb{E} \left[\sum_{ij} \bar{q}_{ij} \right], \end{aligned}$$

where (a) follows from the fact that $\sum_j \lambda_{ij} = 1 - \epsilon$ and $\sum_i \lambda_{ij} = 1 - \epsilon$ since $\boldsymbol{\lambda}^\epsilon = (1 - \epsilon)\boldsymbol{\nu}$ and $\boldsymbol{\nu} \in \mathcal{F}$.

Thus, from (20), we have

$$(21) \quad 2\epsilon \mathbb{E} \left[\sum_{ij} \bar{q}_{ij} \right] = \mathcal{T}_2 + \mathcal{T}_3 + \mathcal{T}_4.$$

Now the rest of the proof involves bounding the term $\mathcal{T}_2, \mathcal{T}_3$ and \mathcal{T}_4 . We start with the term \mathcal{T}_2 . Consider the first term of \mathcal{T}_2 . Again noting that the schedules are maximal (2), we get

$$\begin{aligned} & \mathbb{E} \left[\sum_i \left(\sum_j (\bar{a}_{ij} - s_{ij}(\bar{\mathbf{q}})) \right)^2 \right] \\ = & \sum_i \mathbb{E} \left[\left(\sum_j \bar{a}_{ij} - 1 \right)^2 \right] \end{aligned}$$

$$\begin{aligned}
&= \sum_i \mathbb{E} \left[\left(\sum_j \bar{a}_{ij} - (1 - \epsilon) - \epsilon \right)^2 \right] \\
&= \sum_i \mathbb{E} \left[\left(\sum_j \bar{a}_{ij} - (1 - \epsilon) \right)^2 \right] + \sum_i \epsilon^2 - \sum_i 2\epsilon \mathbb{E} \left[\left(\sum_j \bar{a}_{ij} - (1 - \epsilon) \right) \right] \\
&\stackrel{(a)}{=} n\epsilon^2 + \sum_i \text{Var} \left(\sum_j \bar{a}_{ij} \right) \\
&\stackrel{(b)}{=} n\epsilon^2 + \sum_{ij} \sigma_{ij}^2 \\
&= n\epsilon^2 + \|\boldsymbol{\sigma}\|^2,
\end{aligned}$$

where (a) is true because $\mathbb{E}[\sum_j \bar{a}_{ij}] = (1 - \epsilon)$; (b) follows from the independence of the arrival processes across ports. Similarly, we can show that the second term in \mathcal{T}_2 evaluates to $\mathbb{E} \left[\sum_i \left(\sum_j (\bar{a}_{ij} - s_{ij}(\bar{\mathbf{q}})) \right)^2 \right] = n\epsilon^2 + \|\boldsymbol{\sigma}\|^2$.

The last term can likewise be evaluated as follows.

$$\begin{aligned}
&\frac{1}{n} \mathbb{E} \left[\left(\sum_{ij} (\bar{a}_{ij} - s_{ij}(\bar{\mathbf{q}})) \right)^2 \right] \\
&= \frac{1}{n} \mathbb{E} \left[\left(\sum_{ij} \bar{a}_{ij} - n \right)^2 \right] \\
&= \frac{1}{n} \mathbb{E} \left[\left(\sum_{ij} \bar{a}_{ij} - n(1 - \epsilon) - n\epsilon \right)^2 \right] \\
&= \frac{1}{n} \mathbb{E} \left[\left(\sum_{ij} \bar{a}_{ij} - n(1 - \epsilon) \right)^2 \right] + n\epsilon^2 - 2\epsilon \mathbb{E} \left[\left(\sum_{ij} \bar{a}_{ij} - n(1 - \epsilon) \right) \right] \\
&= n\epsilon^2 + \frac{1}{n} \text{Var} \left(\sum_{ij} \bar{a}_{ij} \right) \\
&= n\epsilon^2 + \frac{1}{n} \sum_{ij} \sigma_{ij}^2 \\
&= n\epsilon^2 + \frac{1}{n} \|\boldsymbol{\sigma}\|^2.
\end{aligned}$$

Putting all the terms of \mathcal{T}_2 together, we get

$$(22) \quad \mathcal{T}_2 = n\epsilon^2 + \left(2 - \frac{1}{n}\right) \|\boldsymbol{\sigma}\|^2.$$

Since $\sum_{ij} q_{ij} \in \mathbb{Z}_+$, we have $\sum_{ij} q_{ij} \leq (\sum_{ij} q_{ij})^2$. Using the fact that $\mathbb{E} \left[(\sum_{ij} \bar{q}_{ij})^2 \right]$ is finite from Lemma 5, we have that $\mathbb{E} \left[\sum_{ij} \bar{q}_{ij} \right]$ is finite and so its drift is zero in steady state. Thus, we get

$$\begin{aligned} 0 &= \mathbb{E} \left[\left[\sum_{ij} q_{ij}(t+1) - \sum_{ij} q_{ij}(t) \right] \mathcal{I}(\mathbf{q}(t) = \bar{\mathbf{q}}) \right] \\ &= \mathbb{E} \left[\sum_{ij} \bar{a}_{ij} - \sum_{ij} s_{ij}(\bar{\mathbf{q}}) + \sum_{ij} u_{ij}(\bar{\mathbf{q}}) \right] \\ \mathbb{E} \left[\sum_{ij} u_{ij}(\bar{\mathbf{q}}) \right] &= n - n(1 - \epsilon) \\ (23) \quad &= n\epsilon. \end{aligned}$$

We will now bound the term \mathcal{T}_3 . Since $u_{ij}(t) \leq s_{ij}(t)$, we have $\sum_i u_{ij} \leq 1$, $\sum_j u_{ij} \leq 1$ and $\sum_{ij} u_{ij} \leq n$. Therefore,

$$\begin{aligned} -\mathbb{E} \left[\sum_i \left(\sum_j u_{ij}(\bar{\mathbf{q}}) \right)^2 \right] - \mathbb{E} \left[\sum_j \left(\sum_i u_{ij}(\bar{\mathbf{q}}) \right)^2 \right] &\leq \mathcal{T}_3 \leq \frac{1}{n} \mathbb{E} \left[\left(\sum_{ij} u_{ij}(\bar{\mathbf{q}}) \right)^2 \right] \\ -\mathbb{E} \left[\sum_i \left(\sum_j u_{ij}(\bar{\mathbf{q}}) \right) \right] - \mathbb{E} \left[\sum_j \left(\sum_i u_{ij}(\bar{\mathbf{q}}) \right) \right] &\leq \mathcal{T}_3 \leq \frac{1}{n} \mathbb{E} \left[n \left(\sum_{ij} u_{ij}(\bar{\mathbf{q}}) \right) \right] \\ (24) \quad &-2n\epsilon \leq \mathcal{T}_3 \leq n\epsilon. \end{aligned}$$

We now consider the term \mathcal{T}_4 . It can be rewritten as follows, and can be split into two parts, one each corresponding to $\bar{\mathbf{q}}_{\parallel}^+$ and $\bar{\mathbf{q}}_{\perp}^+$, where $\bar{\mathbf{q}}_{\parallel}^+$ means $(\bar{\mathbf{q}}^+)_{\parallel}$ and similarly $\bar{\mathbf{q}}_{\perp}^+$.

$$\begin{aligned} \mathcal{T}_4 &= 2\mathbb{E} \left[\sum_{ij} u_{ij}(\bar{\mathbf{q}}) \left(\sum_{j'} \bar{q}_{ij'}^+ + \sum_{i'} \bar{q}_{i'j}^+ - \frac{1}{n} \sum_{i'j'} \bar{q}_{i'j'}^+ \right) \right] \\ &= 2\mathbb{E} \left[\sum_{ij} u_{ij}(\bar{\mathbf{q}}) \left(\sum_{j'} \bar{q}_{\parallel ij'}^+ + \sum_{i'} \bar{q}_{\parallel i'j}^+ - \frac{1}{n} \sum_{i'j'} \bar{q}_{\parallel i'j'}^+ \right) \right] \end{aligned}$$

$$+ 2\mathbb{E} \left[\sum_{ij} u_{ij}(\bar{\mathbf{q}}) \left(\sum_{j'} \bar{q}_{\perp ij'}^+ + \sum_{i'} \bar{q}_{\perp i'j}^+ - \frac{1}{n} \sum_{i'j'} \bar{q}_{\perp i'j'}^+ \right) \right].$$

Since the vector $\bar{\mathbf{q}}_{\parallel}^+$ is in cone \mathcal{K} by definition, (3) in Lemma 1 is applicable. Recall that when $u_{ij}(t) = 1, q_{ij}(t + 1) = 0$. Thus, when $u_{ij}(\bar{\mathbf{q}}) = 1$, we have

$$\begin{aligned} \bar{q}_{ij}^+ &= 0 \\ \bar{q}_{\parallel ij}^+ &= -\bar{q}_{\perp ij}^+ \\ \frac{1}{n} \sum_{j'=1}^n \bar{q}_{\parallel ij'}^+ + \frac{1}{n} \sum_{i'=1}^n \bar{q}_{\parallel i'j}^+ - \frac{1}{n^2} \sum_{i'=1}^n \sum_{j'=1}^n \bar{q}_{\parallel i'j'}^+ &= -\bar{q}_{\perp ij}^+. \end{aligned}$$

Therefore, we get

$$u_{ij}(\bar{\mathbf{q}}) \left(\sum_{j'} \bar{q}_{\parallel ij'}^+ + \sum_{i'} \bar{q}_{\parallel i'j}^+ - \frac{1}{n} \sum_{i'j'} \bar{q}_{\parallel i'j'}^+ \right) = -n u_{ij}(\bar{\mathbf{q}}) \bar{q}_{\perp ij}^+,$$

and the term \mathcal{T}_4 reduces to

$$\begin{aligned} \mathcal{T}_4 &= 2\mathbb{E} \left[\sum_{ij} u_{ij}(\bar{\mathbf{q}}) \left(-n \bar{q}_{\perp ij}^+ + \sum_{j'} \bar{q}_{\perp ij'}^+ + \sum_{i'} \bar{q}_{\perp i'j}^+ - \frac{1}{n} \sum_{i'j'} \bar{q}_{\perp i'j'}^+ \right) \right] \\ &= 2\mathbb{E} \left[\left\langle \mathbf{u}(\bar{\mathbf{q}}), -n \bar{\mathbf{q}}_{\perp}^+ + \sum_i \langle \bar{\mathbf{q}}_{\perp}^+, \mathbf{e}^{(i)} \rangle \mathbf{e}^{(i)} + \sum_j \langle \bar{\mathbf{q}}_{\perp}^+, \tilde{\mathbf{e}}^{(j)} \rangle \tilde{\mathbf{e}}^{(j)} \right. \right. \\ (25) \quad &\quad \left. \left. - \frac{1}{n} \langle \bar{\mathbf{q}}_{\perp}^+, \mathbf{1} \rangle \mathbf{1} \right\rangle \right]. \end{aligned}$$

Term \mathcal{T}_4 is a critical term to bound and our choice of the Lyapunov function $V_4(\cdot)$ is motivated primarily to obtain (25). We explain the motivation in detail at the end of this section. From state space collapse, we know that $\bar{\mathbf{q}}_{\perp}^+$ is bounded. We will now use this result to show that \mathcal{T}_4 is $o(\epsilon)$. Since $\bar{\mathbf{q}}_{\perp}^+ \in \mathcal{K}^\circ$ and $\mathbf{e}^{(i)}, \tilde{\mathbf{e}}^{(j)}, \mathbf{1} \in \mathcal{K}$ for all i, j , we have that $\langle \bar{\mathbf{q}}_{\perp}^+, \mathbf{e}^{(i)} \rangle \leq 0, \langle \bar{\mathbf{q}}_{\perp}^+, \tilde{\mathbf{e}}^{(j)} \rangle \leq 0$ and $\langle \bar{\mathbf{q}}_{\perp}^+, \mathbf{1} \rangle \leq 0$. Moreover all components of $\mathbf{u}, \mathbf{e}^{(i)}, \tilde{\mathbf{e}}^{(j)}$ and $\mathbf{1}$ take values 0 and 1. Therefore,

$$\begin{aligned} \mathcal{T}_4 &\leq 2\mathbb{E} \left[\left\langle \mathbf{u}(\bar{\mathbf{q}}), -n \bar{\mathbf{q}}_{\perp}^+ - \frac{1}{n} \langle \bar{\mathbf{q}}_{\perp}^+, \mathbf{1} \rangle \mathbf{1} \right\rangle \right] \\ &\stackrel{(a)}{\leq} 2 \left(\mathbb{E} \left[\|\mathbf{u}(\bar{\mathbf{q}})\|_{\tilde{r}}^2 \right] \right)^{\frac{1}{r}} \left(\mathbb{E} \left[\left\| -n \bar{\mathbf{q}}_{\perp}^+ - \frac{1}{n} \langle \bar{\mathbf{q}}_{\perp}^+, \mathbf{1} \rangle \mathbf{1} \right\|_r^r \right] \right)^{\frac{1}{r}} \end{aligned}$$

$$\begin{aligned}
 &\stackrel{(b)}{\leq} 2(n\epsilon)^{\frac{1}{\tilde{r}}} \left(\mathbb{E} \left[\left(n\|\bar{\mathbf{q}}_{\perp}^+\|_r + \frac{1}{n} |\langle \bar{\mathbf{q}}_{\perp}^+, \mathbf{1} \rangle| \|\mathbf{1}\|_r \right)^r \right] \right)^{\frac{1}{\tilde{r}}} \\
 &\stackrel{(c)}{\leq} 2(n\epsilon)^{\frac{1}{\tilde{r}}} \left(\mathbb{E} \left[\left(n\|\bar{\mathbf{q}}_{\perp}^+\|_r + \frac{1}{n} \|\bar{\mathbf{q}}_{\perp}^+\|_r \|\mathbf{1}\|_{\tilde{r}} \|\mathbf{1}\|_r \right)^r \right] \right)^{\frac{1}{\tilde{r}}} \\
 &\stackrel{(d)}{=} 2(n\epsilon)^{\frac{1}{\tilde{r}}} \left(\mathbb{E} \left[\left(n\|\bar{\mathbf{q}}_{\perp}^+\|_r + \frac{(n^2)^{\left(\frac{1}{\tilde{r}} + \frac{1}{r}\right)}}{n} \|\bar{\mathbf{q}}_{\perp}^+\|_r \right)^r \right] \right)^{\frac{1}{\tilde{r}}} \\
 &\stackrel{(e)}{=} 4n^{(1+\frac{1}{\tilde{r}})} \epsilon^{\frac{1}{\tilde{r}}} \left(\mathbb{E} [\|\bar{\mathbf{q}}_{\perp}^+\|_r^r] \right)^{\frac{1}{\tilde{r}}} \\
 &\stackrel{(f)}{\leq} 4n^{(1+\frac{1}{\tilde{r}})} \epsilon^{\frac{1}{\tilde{r}}} \left(\mathbb{E} [\|\bar{\mathbf{q}}_{\perp}^+\|_2^r] \right)^{\frac{1}{\tilde{r}}} \quad \text{for } r \geq 2 \\
 &\stackrel{(g)}{\leq} 4n^{(1+\frac{1}{\tilde{r}})} \epsilon^{\frac{1}{\tilde{r}}} M_r^{(\epsilon)} \quad \text{for } r \geq 2 \\
 &\leq 4n^{(2-\frac{1}{\tilde{r}})} \epsilon^{(1-\frac{1}{\tilde{r}})} M_r^{(\epsilon)} \quad \text{for } r \geq 2,
 \end{aligned}$$

where $\|\mathbf{x}\|_r$ denotes the ℓ_r norm of a vector \mathbf{x} , and $r, \tilde{r} \in (1, \infty)$ satisfy $1/r + 1/\tilde{r} = 1$. Inequality (a) follows from the Hölder's inequality for random vectors. Cauchy-Schwartz inequality (which is a special case of Hölder's inequality) may also be used to obtain the same bound in heavy traffic limit. However, in the non-heavy traffic limit, Hölder's inequality gives a tighter bound. Since $u_{ij} \in \{0, 1\}$, from (23), we have $\mathbb{E} [\|\mathbf{u}(\bar{\mathbf{q}})\|_{\tilde{r}}^r] = \mathbb{E} [\sum_{ij} (u_{ij}(\bar{\mathbf{q}}))^{\tilde{r}}] = \mathbb{E} [\sum_{ij} u_{ij}(\bar{\mathbf{q}})] = n\epsilon$. This fact along with using triangle inequality on the second term gives (b). Inequality (c) again follows using Hölder's inequality for vectors. The ℓ_r norm of vector $\mathbf{1}$ is $\|\mathbf{1}\|_r = n^{2/r}$, this gives (d). Since $\frac{1}{\tilde{r}} + \frac{1}{r} = 1$, we have (e). For any vector \mathbf{x} , if $0 < r < r'$, we have $\|\mathbf{x}\|_{r'} \leq \|\mathbf{x}\|_r$, and this gives (f) and (g) follows from state space collapse in Proposition 2. The last inequality follows from $1/r + 1/\tilde{r} = 1$. Similarly, we can lower bound \mathcal{T}_4 as follows.

$$\begin{aligned}
 \mathcal{T}_4 &\geq 2\mathbb{E} \left[\left\langle \mathbf{u}(\bar{\mathbf{q}}), -n\bar{\mathbf{q}}_{\perp}^+ + \sum_i \langle \bar{\mathbf{q}}_{\perp}^+, \mathbf{e}^{(i)} \rangle \mathbf{e}^{(i)} + \sum_j \langle \bar{\mathbf{q}}_{\perp}^+, \tilde{\mathbf{e}}^{(j)} \rangle \tilde{\mathbf{e}}^{(j)} \right\rangle \right] \\
 &\geq -2 \left(\mathbb{E} [\|\mathbf{u}(\bar{\mathbf{q}})\|_{\tilde{r}}^r] \right)^{\frac{1}{\tilde{r}}} \\
 &\quad \times \left(\mathbb{E} \left[\left\| -n\bar{\mathbf{q}}_{\perp}^+ + \sum_i \langle \bar{\mathbf{q}}_{\perp}^+, \mathbf{e}^{(i)} \rangle \mathbf{e}^{(i)} + \sum_j \langle \bar{\mathbf{q}}_{\perp}^+, \tilde{\mathbf{e}}^{(j)} \rangle \tilde{\mathbf{e}}^{(j)} \right\|_r^r \right] \right)^{\frac{1}{r}} \\
 &\geq -2(n\epsilon)^{\frac{1}{\tilde{r}}} \left(\mathbb{E} \left[\left(\|n\bar{\mathbf{q}}_{\perp}^+\|_r + \left\| \sum_i \langle \bar{\mathbf{q}}_{\perp}^+, \mathbf{e}^{(i)} \rangle \mathbf{e}^{(i)} \right\|_r \right)^r \right] \right)^{\frac{1}{\tilde{r}}}
 \end{aligned}$$

$$(26) \quad + \left\| \left[\sum_j \left\langle \bar{\mathbf{q}}_{\perp}^+, \tilde{\mathbf{e}}^{(j)} \right\rangle \tilde{\mathbf{e}}^{(j)} \right] \right\|_r^{\frac{1}{r}}.$$

Let's now focus on the middle term in the expectation above. From the definition of $\mathbf{e}^{(i)}$, we have

$$\begin{aligned} \left\| \sum_i \left\langle \bar{\mathbf{q}}_{\perp}^+, \mathbf{e}^{(i)} \right\rangle \mathbf{e}^{(i)} \right\|_r &= \left(\sum_i n \left| \left\langle \bar{\mathbf{q}}_{\perp}^+, \mathbf{e}^{(i)} \right\rangle \right|^r \right)^{\frac{1}{r}} \\ &= \left(\sum_i n \left(\sum_j q_{\perp ij}^+ \right)^r \right)^{\frac{1}{r}} \\ &\stackrel{(a)}{\leq} \left(\sum_i n^r \sum_j \left(q_{\perp ij}^+ \right)^r \right)^{\frac{1}{r}} \\ &= n \left\| \bar{\mathbf{q}}_{\perp}^+ \right\|_r. \end{aligned}$$

For any $(x_1, \dots, x_n) \in \mathbb{R}^n$ and $r \geq 1$, from Jensen's inequality, we have $\left(\frac{\sum_i x_i}{n}\right)^r \leq \frac{\sum_i x_i^r}{n}$. This gives inequality (a) above. We have a similar bound for the last term in expectation in (26). Using both these bounds, the lower bound on \mathcal{T}_4 becomes,

$$\begin{aligned} \mathcal{T}_4 &\geq -6n^{(1+\frac{1}{r})} \epsilon^{\frac{1}{r}} \left(\mathbb{E} \left[\left(\left\| \bar{\mathbf{q}}_{\perp}^+ \right\|_r \right)^r \right] \right)^{\frac{1}{r}} \\ &\geq -6n^{(1+\frac{1}{r})} \epsilon^{\frac{1}{r}} \left(\mathbb{E} \left[\left(\left\| \bar{\mathbf{q}}_{\perp}^+ \right\|_2 \right)^r \right] \right)^{\frac{1}{r}} && \text{for } r \geq 2 \\ &\geq -6n^{(1+\frac{1}{r})} \epsilon^{\frac{1}{r}} M_r^{(\epsilon)} && \text{for } r \geq 2 \\ &\geq -6n^{(2-\frac{1}{r})} \epsilon^{(1-\frac{1}{r})} M_r^{(\epsilon)} && \text{for } r \geq 2. \end{aligned}$$

Combining the lower and upper bounds on \mathcal{T}_4 , for $r \geq 2$, we have

$$(27) \quad -6n^{(2-\frac{1}{r})} \epsilon^{(1-\frac{1}{r})} M_r^{(\epsilon)} \leq \mathcal{T}_4 \leq 4n^{(2-\frac{1}{r})} \epsilon^{(1-\frac{1}{r})} M_r^{(\epsilon)}.$$

Using (22),(24) and (27) to bound (21) and reintroducing the superscript $^{(\epsilon)}$, we get the theorem. □

We will now present the motivation for the choice of the function $V_4(\cdot)$. First consider a discrete-time single server (G/G/1) queue, $q(t)$ that evolves according to $q(t+1) = q(t) + a(t) - s(t) + u(t)$. The queue $\phi(t)$ in Section 3 is an example. Similar to (8), we can obtain tight lower and upper bounds

on mean queue length in steady state by setting the drift of $\mathbb{E}[\bar{q}^2]$ to be zero in steady state, i.e. $\mathbb{E}[\bar{q}^2(t+1)] = \mathbb{E}[\bar{q}^2(t)]$. Such a bound is called Kingman bound. See [15, Section 10.1]. When expanded, this equation again gives four terms, similar to the terms $\mathcal{T}_1, \mathcal{T}_2, \mathcal{T}_3$ and \mathcal{T}_4 . The fourth term \mathcal{T}_4 then is $u(\bar{q})\bar{q}^+$, which is zero from the definition of unused service. This is an important step in obtaining tight bounds.

Next, consider a load balancing system, similar to a super market checkout lanes. There are n servers with a separate queue for each server. Whenever a user arrives into the system, (s)he picks one of the servers and joins the corresponding queue. We consider ‘Join the shortest queue’(JSQ) policy, in which each user joins the queue with the shortest length. Ties are broken uniformly at random. The queue length at server i then evolves according to $q_i(t+1) = q_i(t) + a_i(t) - s_i(t) + u_i(t)$. It was shown in [7] that the JSQ policy has minimum steady state sum queue lengths in heavy traffic. This was done by first showing that the queue lengths collapse to a single dimension where they are all equal. A tight upper bound is then obtained by setting the drift of the quadratic function $\mathbb{E}[(\sum_i \bar{q}_i)^2]$ to be zero in steady state. When this equation is expanded, we again have four terms and the fourth one being of the form $(\sum_i u_i(\bar{q}))(\sum_{i'} \bar{q}_{i'}^+)$. This is not zero in general because of the cross terms. However, when the state is such that all the queue lengths are equal, this term is zero. This is easy to see by considering the term $u_i(\bar{q})(\sum_{i'} \bar{q}_{i'}^+)$. When $u_i = 1$, we have that $\bar{q}_i^+ = 0$ and when all the queues are equal, for any i' , $\bar{q}_{i'}^+ = 0$.

Therefore, in all these systems, when using a quadratic Lyapunov function, the fourth term \mathcal{T}_4 is the most important and challenging one to bound correctly. Usually, it should be zero if state space collapse is such that $\bar{\mathbf{q}}_{\perp}^+ = \mathbf{0}$. However, for the switch system, if we use Lyapunov functions $V_1(\cdot)$ or $V_2(\cdot)$ or $V_3(\cdot)$ or $V_1(\cdot) + V_2(\cdot)$, we do not have the property that $\mathcal{T}_4 = 0$ when $\bar{\mathbf{q}}_{\perp}^+ = \mathbf{0}$. Armed with (3) in Lemma 1, we add the additional $-V_3(\cdot)$ to $V_1(\cdot) + V_2(\cdot)$ to obtain the Lyapunov function $V_4(\cdot)$. We have shown in (25) that \mathcal{T}_4 is zero when $\bar{\mathbf{q}} \in \mathcal{K}$ (since $\bar{\mathbf{q}}_{\perp}^+ = \mathbf{0}$). The key idea in our upper bound proof is the choice of the function $V_4(\cdot)$. Essentially, we picked the function $V_4(\cdot)$ so that it matches with the geometry of the cone \mathcal{K} in the sense that if the queue length vector is in the cone \mathcal{K} , the fourth term \mathcal{T}_4 is zero.

6. Uniformly loaded switch under Bernoulli traffic. In this section, we consider the switch system when all the ports have Bernoulli traffic with same arrival rate. The lower and upper bound expressions then have much simple form. More precisely, for the system with index ϵ , for every input-output pair (i, j) , the arrival process $a_{ij}^{(\epsilon)}(t)$ is a Bernoulli process

with rate $\lambda_{ij} = (1 - \epsilon)/n$. In other words, the rate vector approaches the vector $\boldsymbol{\nu} = \mathbf{1}/n \in \mathcal{F}$ on the face \mathcal{F} as $\epsilon \rightarrow 0$. Then, clearly the variance vector for the system with index ϵ is $(\boldsymbol{\sigma}^{(\epsilon)})^2 = \frac{1-\epsilon}{n}(1 - \frac{1-\epsilon}{n})\mathbf{1}$ with $\|\boldsymbol{\sigma}^{(\epsilon)}\|^2 = (1 - \epsilon)(n - (1 - \epsilon))$ and it converges to $\boldsymbol{\sigma}^2 = \frac{n-1}{n^2}\mathbf{1}$. Moreover, $a_{\max} = 1$ and $\nu_{\min} = \frac{1}{n}$. Using these values, we can restate Propositions 1 and 2, and Theorem 1 as follows:

THEOREM 2. *Consider a set of switch systems with the Bernoulli arrival processes $\mathbf{a}^{(\epsilon)}(t)$ parameterized by $0 < \epsilon < 1$, such that the mean arrival rate vector is $\boldsymbol{\lambda}^\epsilon = \frac{1-\epsilon}{n}\mathbf{1}$. 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 average queue length is lower bounded by*

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

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

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

Now consider the same switch systems operating under the MaxWeight scheduling algorithm. The queue length process $\mathbf{q}^{(\epsilon)}(t)$ of each system is positive recurrent and so converges to a steady state random vector in distribution $\bar{\mathbf{q}}^{(\epsilon)}$. Then, for each system with $0 < \epsilon \leq 1/2n$, the steady state queue lengths vector collapses into the cone \mathcal{K} in the sense that it satisfies

$$\mathbb{E} \left[\|\bar{\mathbf{q}}_{\perp}^{(\epsilon)}\|^r \right] \leq (\widetilde{M}_r)^r \quad \forall r \in \{1, 2, \dots\}, \quad \text{where } \widetilde{M}_r = (2\sqrt{r}\epsilon)^{1/r} 16 \frac{r}{e} n^2 (n + 1).$$

Therefore, the steady state average queue length satisfies

$$\frac{1}{\epsilon} \left(n - \frac{3}{2} + \frac{1}{2n} \right) - \widetilde{B}_1(\epsilon, n) \leq \mathbb{E} \left[\sum_{ij} \bar{q}_{ij}^{(\epsilon)} \right] \leq \frac{1}{\epsilon} \left(n - \frac{3}{2} + \frac{1}{2n} \right) + \widetilde{B}_2(\epsilon, n),$$

where

$$\begin{aligned} \widetilde{B}_1(\epsilon, n) &= \left(1 - \frac{\epsilon}{2}\right) \left(n - 2 + \frac{1}{n}\right) + n - \frac{1}{2} + 3n^{(2-\frac{1}{r})} \epsilon^{(-\frac{1}{r})} \widetilde{M}_r \quad \text{and} \\ \widetilde{B}_2(\epsilon, n) &= - \left(1 - \frac{\epsilon}{2}\right) \left(n - 2 + \frac{1}{n}\right) + \frac{n + 1}{2} + 2n^{(2-\frac{1}{r})} \epsilon^{(-\frac{1}{r})} \widetilde{M}_r \end{aligned}$$

for any $r \in \{2, 3, \dots\}$. The terms $\widetilde{B}_1(\epsilon, n)$ and $\widetilde{B}_2(\epsilon, n)$ are both $o\left(\frac{1}{\epsilon}\right)$. In the heavy traffic limit as $\epsilon \downarrow 0$ which means as the mean arrival rate $\lambda^\epsilon \rightarrow \frac{1}{n}\mathbf{1}$, we have

$$\lim_{\epsilon \downarrow 0} \epsilon \mathbb{E} \left[\sum_{ij} \bar{q}_{ij}^{(\epsilon)} \right] = \left(n - \frac{3}{2} + \frac{1}{2n} \right).$$

Thus, MaxWeight algorithm has optimal queue length scaling in the heavy traffic limit.

Thus, in the heavy traffic limit, we have a universal lower bound on the (ϵ scaled) average queue lengths that is $\Omega(n)$ and the MaxWeight policy achieves this bound within a factor less than 2. Since we are interested in the asymptotics both in term of number of ports, n and distance from boundary of the capacity region, ϵ , there are several possible limits in which the system can be studied. Heavy traffic limit is one such asymptotic, where we first let the arrival rate approach the boundary of the capacity region and look at the scaling of average queue length in terms of n . Another set of asymptotic regimes is when $\epsilon \rightarrow 0$ and $n \rightarrow \infty$ simultaneously. This can be studied by setting $\epsilon = n^{-\beta}$ for $\beta > 0$. Such a limit was studied in [13, 14] for scheduling algorithms that are different from the MaxWeight algorithms studied here. The universal lower bound in such a limit is $\Omega(n^{(1+\beta)})$. It is now easy to see the following corollary.

COROLLARY 1. *Consider a sequence of switch systems with Bernoulli arrivals, indexed by n . The n^{th} system has mean arrival rate vector $\lambda^{(n)} = \frac{1-\gamma_n n^{-\beta}}{n} \mathbf{1}$ with $\beta > 0$ and $\gamma_n > 0$ is a sequence that is $\Theta(1)$. The load is $\rho^{(n)} = 1 - \gamma_n n^{-\beta}$. Fix a scheduling policy under which the switch system is stable for any $n > 0$. Suppose that the queue lengths process $\mathbf{q}^{(n)}(t)$ process converges in distribution to a steady state random vector $\bar{\mathbf{q}}^{(n)}$. Then, for each of these systems, the average queue length is lower bounded by*

$$\mathbb{E} \left[\sum_{ij} \bar{q}_{ij}^{(n)} \right] \geq \frac{(1 - \gamma_n n^{-\beta})^2}{2\gamma_n} n^\beta (n - 1),$$

and so is $\Omega(n^{(1+\beta)})$.

Under the MaxWeight scheduling policy, the queue lengths process $\mathbf{q}^{(n)}(t)$ process is positive recurrent and so converges to a steady state random vector in distribution $\bar{\mathbf{q}}^{(n)}$. When $2\gamma_n \leq n^{(\beta-1)}$, the steady state average queue

length satisfies

$$(28) \quad \frac{n^{(1+\beta)}}{\gamma_n} - B_3(n) \leq \mathbb{E} \left[\sum_{ij} \bar{q}_{ij}^{(n)} \right] \leq \frac{n^{(1+\beta)}}{\gamma_n} + B_4(n) \quad \text{for } \beta > 4,$$

where $B_3(n)$ and $B_4(n)$ are $o(n^{(1+\beta)})$. Thus, under the MaxWeight algorithm, the average sum queue lengths is $\Theta(n^{(1+\beta)})$ and so has optimal scaling.

PROOF. The universal lower bound directly follows from Theorem 2 using $\epsilon^{(n)} = \gamma_n n^{-\beta}$. We will now prove the second part of the corollary which is under the MaxWeight policy. Since $2\gamma_n \leq n^{(\beta-1)}$, we have $0 < \epsilon^{(n)} \leq 1/2n$ and Theorem 2 is applicable. Therefore, we have (28) with

$$\begin{aligned} B_3(n) &= \left(\frac{3n^\beta - n^{\beta-1}}{2\gamma_n} \right) + \left(1 - \frac{\gamma_n n^{-\beta}}{2} \right) \left(n - 2 + \frac{1}{n} \right) + n - \frac{1}{2} \\ &\quad + 48 \left(\frac{2\sqrt{r}e}{\gamma_n} \right)^{1/r} \frac{r}{e} n^{(2-\frac{1}{r}+\frac{\beta}{r})} n^2 (n+1), \\ B_4(n) &= \left(\frac{-3n^\beta + n^{\beta-1}}{2\gamma_n} \right) - \left(1 - \frac{\gamma_n n^{-\beta}}{2} \right) \left(n - 2 + \frac{1}{n} \right) + \frac{n+1}{2} \\ &\quad + 32 \left(\frac{2\sqrt{r}e}{\gamma_n} \right)^{1/r} \frac{r}{e} n^{(2-\frac{1}{r}+\frac{\beta}{r})} n^2 (n+1). \end{aligned}$$

Clearly all but the last terms above are $o(n^{(1+\beta)})$. The last terms are $\Theta\left(n^{(5+\frac{\beta-1}{r})}\right)$. For any $\beta > 4$, we can pick r large enough so that $4 + \frac{\beta-1}{r} < \beta$ and so we have that $B_3(n)$ and $B_4(n)$ are $o(n^{(1+\beta)})$. \square

7. Conclusion. We have obtained a characterization of the heavy-traffic behavior of the sum queue length in steady-state in an $n \times n$ switch operating under the MaxWeight scheduling policy when all ports are saturated. We then considered the special case of uniform Bernoulli traffic and studied the switch in an asymptotic regime where the load increases simultaneously with the number of ports. We showed that the steady-state average queue lengths are within a factor less than 2 of a universal lower bound. The result settles one version of a conjecture regarding the performance of the MaxWeight policy. A number of extensions can be considered:

- Extensions of the result to more general traffic patterns when only a few ports are saturated or when some of the arrival rates are zero is an open problem.

- We believe that one may be also be able to allow correlations across time slots by making an assumption similar to the assumption in Section II.C of [24], and considering the drift of the Lyapunov function over multiple time slots. This extension may require a bit of additional work.
- A Brownian limit has been established in the heavy-traffic regime in [6], but a characterization of the behavior of this limit in steady-state is not known. We expect the mean of the sum queue lengths (multiplied by ϵ and in the limit $\epsilon \rightarrow 0$) in steady-state that we have derived to be equal to the sum of the steady-state expectations of the components of the Brownian motion in [6]. This would be interesting to verify.
- Verifying whether the MaxWeight algorithm achieves optimal queue-length scaling in the size of the switch in non-heavy-traffic regimes is still an open problem.

APPENDIX A: PROOF OF LEMMA 1

PROOF. Clearly, from the definition of the cone \mathcal{K} , (i) \iff (ii). We will now show that (ii) \implies (iii) \implies (iv) \implies (ii).

(ii) \implies (iii): Since all the maximal matchings have exactly one element from each row and column, when a vector \mathbf{x} such that $x_{ij} = w_i + \tilde{w}_j$ for all i, j , is used as the weight, all the matchings have the same weight, $\sum_i w_i + \sum_j \tilde{w}_j$. Moreover, since $w_i \geq 0$ and $\tilde{w}_j \geq 0$ for all i, j , we have that $x_{ij} \geq 0$ for all i, j .

(iii) \implies (iv): There are $n!$ perfect matchings and each edge appears in $(n - 1)!$ of them. Therefore, when \mathbf{x} is used as weight vector, the average weight of all the perfect matchings is $\frac{1}{n} \sum_{i',j'} x_{i'j'}$. Now consider the $(n - 1)!$ matchings that contain the edge ij . The total weight of all these matchings is $(n - 1)!x_{ij} + \sum_{i' \neq i} \sum_{j' \neq j} (n - 2)!x_{i'j'}$, because every edge $i'j'$ appears in $(n - 2)!$ of these $(n - 1)!$ matchings. Since all the matchings have same weight, equating the average weight of these $(n - 1)!$ matchings to the average of all the matchings, we have

$$\begin{aligned}
 x_{ij} + \frac{1}{n - 1} \sum_{i' \neq i} \sum_{j' \neq j} x_{i'j'} &= \frac{1}{n} \sum_{i',j'} x_{i'j'} \\
 x_{ij} + \frac{1}{n - 1} \left(\sum_{i',j'} x_{i'j'} - \sum_{j'=1}^n x_{ij'} - \sum_{i'=1}^n x_{i'j} + x_{ij} \right) &= \frac{1}{n} \sum_{i',j'} x_{i'j'} \\
 x_{ij} \left(1 + \frac{1}{n - 1} \right) - \frac{1}{n - 1} \left(\sum_{j'=1}^n x_{ij'} + \sum_{i'=1}^n x_{i'j} \right) &= \sum_{i',j'} x_{i'j'} \left(\frac{1}{n} - \frac{1}{n - 1} \right)
 \end{aligned}$$

$$nx_{ij} - \left(\sum_{j'=1}^n x_{ij'} + \sum_{i'=1}^n x_{i'j} \right) = -\frac{1}{n} \sum_{i',j'} x_{i'j'} \text{ for } n > 1.$$

This gives (3) since the $n = 1$ case is trivial.

(iv) \implies (ii): We now assume that (3) is true. For each i, j , define

$$\omega_i \triangleq \frac{1}{n} \sum_{j'=1}^n x_{ij'} - \frac{1}{2n^2} \sum_{i'=1}^n \sum_{j'=1}^n x_{i'j'}$$

and

$$\tilde{\omega}_j \triangleq \frac{1}{n} \sum_{i'=1}^n x_{i'j} - \frac{1}{2n^2} \sum_{i'=1}^n \sum_{j'=1}^n x_{i'j'}.$$

Then clearly, $x_{ij} = \omega_i + \tilde{\omega}_j$ for all i, j . If $\omega_i \geq 0$ for all i and $\tilde{\omega}_j \geq 0$, for all j , we have (ii). If not, let $\zeta_1 \triangleq \min_i \omega_i$ and $\zeta_2 \triangleq \min_j \tilde{\omega}_j$. Note that $\zeta_1 + \zeta_2 \geq 0$ since $x_{ij} \geq 0$ for all i, j . Suppose that $\zeta_1 < 0$. Then, let $w_i = \omega_i - \zeta_1$ and $\tilde{w}_j = \tilde{\omega}_j + \zeta_1$. Then clearly, $w_i \geq 0$ for all i and $\tilde{w}_j \geq 0$ for all j and $x_{ij} = \omega_i + \tilde{\omega}_j = w_i + \tilde{w}_j$. Thus, we have (ii). One can similarly define $w_i \geq 0$ and $\tilde{w}_j \geq 0$ when $\zeta_2 < 0$. This proves the lemma. \square

APPENDIX B: PROOF OF LEMMA 3

PROOF. Lemma 2 is applicable here and so we have that $\mathbb{E}[Z(\bar{X})] < \infty$. Recall that $\Delta Z(X)$ is a random variable for any X , so define

$$\tilde{D} \triangleq \sup_{X \in \mathcal{X}} \text{ess sup} |\Delta Z(X)| = \sup_{X, X' \in \mathcal{X}, \mathbb{P}(X(t+1)=X'|X(t)=X) > 0} |Z(X') - Z(X)|.$$

Also define

$$p_{\max} = \sup_{X \in \mathcal{X}} \mathbb{P}(X(t+1) > X | X(t) = X).$$

Then, from Theorem 1 in [20], we have

$$\mathbb{P}\left(Z(\bar{X}) > \kappa + 2\tilde{D}m\right) \leq \left(\frac{\tilde{D}p_{\max}}{\tilde{D}p_{\max} + \eta}\right)^{m+1}.$$

Clearly, $\tilde{D} \leq D$ and $p_{\max} \leq 1$. Therefore, we get

$$\begin{aligned} \mathbb{P}\left(Z(\bar{X}) > \kappa + 2Dm\right) &\leq \mathbb{P}\left(Z(\bar{X}) > \kappa + 2\tilde{D}m\right) \\ &\leq \left(\frac{\tilde{D}p_{\max}}{\tilde{D}p_{\max} + \eta}\right)^{m+1} \\ &\leq \left(\frac{D}{D + \eta}\right)^{m+1}, \end{aligned}$$

where the last inequality follows from $\tilde{D}p_{\max} \leq D$ and $m + 1 \geq 1$. This proves the first part of the lemma. We will now use this result to obtain moment bounds. Since $r > 0$ and $Z(\cdot) \geq 0$, we have

$$\begin{aligned}
 \mathbb{E}[Z(\bar{X})^r] &= r \int_{t=0}^{\infty} t^{r-1} \mathbb{P}(Z(\bar{X}) > t) dt \\
 &= r \int_{t=0}^{\kappa} t^{r-1} \mathbb{P}(Z(\bar{X}) > t) dt + r \int_{t=\kappa}^{\infty} t^{r-1} \mathbb{P}(Z(\bar{X}) > t) dt \\
 &\leq r \int_{t=0}^{\kappa} t^{r-1} dt + r \sum_{m=0}^{\infty} \int_{t=\kappa+2Dm}^{\kappa+2D(m+1)} t^{r-1} \mathbb{P}(Z(\bar{X}) > t) dt \\
 &\leq \kappa^r + r \sum_{m=0}^{\infty} \int_{t=\kappa+2Dm}^{\kappa+2D(m+1)} t^{r-1} \mathbb{P}(Z(\bar{X}) > \kappa + 2Dm) dt \\
 &\leq \kappa^r + \sum_{m=0}^{\infty} \left(\frac{D}{D+\eta}\right)^{m+1} \int_{t=\kappa+2Dm}^{\kappa+2D(m+1)} r t^{r-1} dt \\
 &= \kappa^r + \sum_{m=0}^{\infty} \left(\frac{D}{D+\eta}\right)^{m+1} (\kappa + 2D(m+1))^r - (\kappa + 2Dm)^r \\
 &= \kappa^r \left(1 - \frac{D}{D+\eta}\right) + \sum_{m=1}^{\infty} (\kappa + 2Dm)^r \left[\left(\frac{D}{D+\eta}\right)^m - \left(\frac{D}{D+\eta}\right)^{m+1}\right] \\
 &= \left(\frac{\eta}{D+\eta}\right) \left[\sum_{m=0}^{\infty} (\kappa + 2Dm)^r \left(\frac{D}{D+\eta}\right)^m\right] \\
 &\stackrel{(a)}{\leq} \left(\frac{\eta}{D+\eta}\right) \left[\sum_{m=0}^{\infty} (2\kappa)^r \left(\frac{D}{D+\eta}\right)^m + \sum_{m=0}^{\infty} (4Dm)^r \left(\frac{D}{D+\eta}\right)^m\right] \\
 &= (2\kappa)^r + (4D)^r \left(\frac{\eta}{D+\eta}\right) \sum_{m=0}^{\infty} m^r \left(\frac{D}{D+\eta}\right)^m,
 \end{aligned}$$

where (a) follows from the relation, $(a + b)^r \leq 2^r \max(a, b)^r \leq 2^r (a^r + b^r)$. It is known [25] that for $x < 1$ and $r = 1, 2, \dots$ $\sum_{m=0}^{\infty} m^r x^m = \frac{1}{(1-x)^{r+1}} \sum_{k=0}^{r-1} A(r, k) x^{k+1}$, where $A(r, k)$ are called the Eulerian numbers. It is also known that $\sum_{k=0}^{r-1} A(r, k) = r!$. Therefore, when $x < 1$, we have that $\sum_{m=0}^{\infty} m^r x^m \leq \frac{1}{(1-x)^{r+1}} r!$. Using this relation, we get

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

□

APPENDIX C: PROOF OF LEMMA 5

PROOF. We will use Lemma 2 to first show that $\mathbb{E}[V(\bar{\mathbf{q}})]$ is finite. Define the Lyapunov function $W(\mathbf{q}) \triangleq \|\mathbf{q}\| = \sqrt{V(\mathbf{q})}$, and its drift

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

We will first verify condition C.2 of Lemma 2. Using the same arguments as in (9), we get

$$\begin{aligned} |\Delta W(\mathbf{q})| &= \left| \|\mathbf{q}(t+1)\| - \|\mathbf{q}(t)\| \right| \mathcal{I}(\mathbf{q}(t) = \mathbf{q}) \\ &\leq \|\mathbf{q}(t+1) - \mathbf{q}(t)\| \mathcal{I}(\mathbf{q}(t) = \mathbf{q}) \\ &\leq n^2 \max_{ij} |q_{ij}(t+1) - q_{ij}(t)| \mathcal{I}(\mathbf{q}(t) = \mathbf{q}) \\ &\leq n^2 a_{\max}, \end{aligned}$$

thus verifying condition C.2. We will now verify condition C.1.

$$\begin{aligned} &\mathbb{E}[\Delta W(\mathbf{q}) | \mathbf{q}(t) = \mathbf{q}] \\ &= \mathbb{E}[\|\mathbf{q}(t+1)\| - \|\mathbf{q}(t)\| | \mathbf{q}(t) = \mathbf{q}] \\ &= \mathbb{E}\left[\sqrt{\|\mathbf{q}(t+1)\|^2} - \sqrt{\|\mathbf{q}(t)\|^2} \mid \mathbf{q}(t) = \mathbf{q}\right] \\ &\stackrel{(a)}{\leq} \mathbb{E}\left[\frac{1}{2\|\mathbf{q}(t)\|} \|\mathbf{q}(t+1)\|^2 - \|\mathbf{q}(t)\|^2 \mid \mathbf{q}(t) = \mathbf{q}\right] \\ &= \frac{1}{2\|\mathbf{q}\|} \mathbb{E}[\Delta V(\mathbf{q}) | \mathbf{q}(t) = \mathbf{q}] \\ &\stackrel{(b)}{\leq} \frac{1}{2\|\mathbf{q}\|} (\|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\sigma}\|^2 + n + 2 \langle \mathbf{q}, \boldsymbol{\lambda} - \mathbb{E}[\mathbf{s}(t) | \mathbf{q}(t) = \mathbf{q}] \rangle) \\ &\stackrel{(c)}{\leq} \frac{1}{2\|\mathbf{q}\|} \left(\|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\sigma}\|^2 + n + 2 \min_{\mathbf{r} \in \mathcal{C}} \langle \mathbf{q}, \boldsymbol{\lambda} - \mathbf{r} \rangle \right) \\ &\stackrel{(d)}{\leq} \frac{1}{2\|\mathbf{q}\|} (\|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\sigma}\|^2 + n + 2 \langle \mathbf{q}, \boldsymbol{\lambda} - (\boldsymbol{\lambda} + \epsilon_1 \mathbf{1}) \rangle) \\ &= \frac{\|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\sigma}\|^2 + n}{2\|\mathbf{q}\|} - \epsilon_1 \frac{\|\mathbf{q}\|_1}{\|\mathbf{q}\|} \\ &\stackrel{(e)}{\leq} \frac{\|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\sigma}\|^2 + n}{2\|\mathbf{q}\|} - \epsilon_1 \\ &\leq -\frac{\epsilon_1}{2} \text{ for all } \mathbf{q} \text{ such that } W(\mathbf{q}) \geq \frac{\|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\sigma}\|^2 + n}{\epsilon_1}, \end{aligned}$$

where $\boldsymbol{\sigma}$ denotes the variance vector and $\|\mathbf{q}\|_1 \triangleq \sum_{ij} q_{ij}$ denotes the ℓ_1 norm of \mathbf{q} . Inequality (a) follows from the concavity of square root function, due

to which we have that $\sqrt{y} - \sqrt{x} \leq \frac{1}{2\sqrt{x}}(y - x)$. Inequality (b) follows from the bound on drift of $V(\cdot)$ obtained in (10) in the proof of the proof of Proposition 2; (c) follows from the fact that we use MaxWeight scheduling. Since $\boldsymbol{\lambda} \in \text{int}(\mathcal{C})$, there exists a $\epsilon_1 > 0$ such that $\boldsymbol{\lambda} + \epsilon_1 \mathbf{1} \in \mathcal{C}$. This gives (d). For any vector \mathbf{x} , its ℓ_1 norm is at least its ℓ_2 norm, i.e., $\|\mathbf{x}\|_1 \geq \|\mathbf{x}\|$. This gives inequality (e). Thus, condition C.1 is verified and we have that all moments of $W(\bar{\mathbf{q}})$ exist in steady state. In particular, we have that $\mathbb{E}[V(\bar{\mathbf{q}})]$ is finite.

Now, note that

$$V_3(\mathbf{q}) = \left(\sum_{ij} q_{ij} \right)^2 \leq \left(\sum_{ij} \max_{ij} q_{ij} \right)^2 = n^4 \max_{ij} q_{ij}^2 \leq n^4 \sum_{ij} q_{ij}^2 = n^4 V(\mathbf{q}).$$

Thus, $\mathbb{E}[V_3(\bar{\mathbf{q}})]$ is also finite. The lemma follows by noting that $V_1(\mathbf{q}) \leq V_3(\mathbf{q})$ and $V_2(\mathbf{q}) \leq V_3(\mathbf{q})$. \square

REFERENCES

- [1] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Transactions on Automatic Control*, vol. 37, no. 12, pp. 1936–1948, 1992. [MR1200609](#)
- [2] N. McKeown, V. Anantharam, and J. Walrand, "Achieving 100% throughput in an input queued switch," in *Proceedings of IEEE INFOCOM*, 1996, pp. 296–302.
- [3] A. L. Stolyar, "Maxweight scheduling in a generalized switch: State space collapse and workload minimization in heavy traffic," *Annals of Applied Probability*, pp. 1–53, 2004. [MR2023015](#)
- [4] M. Andrews, K. Jung, and A. Stolyar, "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*, ser. STOC '07, 2007, pp. 145–154. [MR2402438](#)
- [5] D. Shah and D. Wischik, "Switched networks with maximum weight policies: Fluid approximation and multiplicative state space collapse," *The Annals of Applied Probability*, vol. 22, no. 1, pp. 70–127, 2012. [MR2932543](#)
- [6] W. N. Kang and R. J. Williams, "Diffusion approximation for an input-queued packet switch operating under a maximum weight algorithm," *Stochastic Systems*, 2012.
- [7] A. Eryilmaz and R. Srikant, "Asymptotically tight steady-state queue length bounds implied by drift conditions," *Queueing Systems*, vol. 72, no. 3-4, pp. 311–359, 2012. [MR2989493](#)
- [8] J. M. Harrison and R. J. Williams, "Brownian models of open queueing networks with homogeneous customer populations," *Stochastics*, vol. 22, no. 2, pp. 77–115, 1987. [MR0912049](#)
- [9] W. N. Kang, F. P. Kelly, N. H. Lee, and R. J. Williams, "State space collapse and diffusion approximation for a network operating under a fair bandwidth sharing policy," *The Annals of Applied Probability*, pp. 1719–1780, 2009. [MR2569806](#)
- [10] T. Ji, E. Athanasopoulou, and R. Srikant, "On optimal scheduling algorithms for small generalized switches," *IEEE/ACM Transactions on Networking*, vol. 18, no. 5, pp. 1585–1598, 2010.

- [11] D. Shah, J. Tsitsiklis, and Y. Zhong, "Optimal scaling of average queue sizes in an input-queued switch: an open problem," *Queueing Systems*, vol. 68, no. 3-4, pp. 375–384, 2011. [MR2834209](#)
- [12] M. J. Neely, E. Modiano, and Y.-S. Cheng, "Logarithmic delay for $n \times n$ packet switches under the crossbar constraint," *IEEE/ACM Transactions on Networking*, vol. 15, no. 3, pp. 657–668, 2007.
- [13] D. Shah, N. S. Walton, and Y. Zhong, "Optimal queue-size scaling in switched networks," *Ann. Appl. Probab.*, vol. 24, no. 6, pp. 2207–2245, 12 2014. [MR3262502](#)
- [14] D. Shah, J. N. Tsitsiklis, and Y. Zhong, "On queue-size scaling for input-queued switches," 2014, arxiv. [MR3161642](#)
- [15] R. Srikant and L. Ying, *Communication Networks: An Optimization, Control and Stochastic Networks Perspective*. Cambridge University Press, 2014. [MR3202391](#)
- [16] M. A. Readdy, "Polytopes," Lecture Notes, http://www.ms.uky.edu/~readdy/Papers/readdy_WAM_lectures.pdf.
- [17] G. Ziegler, *Lectures on Polytopes*, ser. Graduate Texts in Mathematics. Springer New York, 1995. [MR1311028](#)
- [18] J. Dattorro, *Convex Optimization & Euclidean Distance Geometry*. Meboo Publishing, 2005.
- [19] B. Hajek, "Hitting-time and occupation-time bounds implied by drift analysis with applications," *Advances in Applied Probability*, pp. 502–525, 1982. [MR0665291](#)
- [20] D. Bertsimas, D. Gamarnik, and J. N. Tsitsiklis, "Performance of multiclass markovian queueing networks via piecewise linear lyapunov functions," *Ann. Appl. Probab.*, vol. 11, no. 4, pp. 1384–1428, 11 2001. [MR1878302](#)
- [21] R. Wu, course project, ECE 567 Communication Network Analysis, Fall 2012, University of Illinois at Urbana Champaign.
- [22] R. Singh and A. Stolyar, "Maxweight scheduling: Asymptotic behavior of unscaled queue-differentials in heavy traffic," *arXiv preprint [arXiv:1502.03793](#)*, 2015.
- [23] A. Schrijver, *Combinatorial Optimization: Polyhedra and Efficiency*, ser. Algorithms and Combinatorics. Springer, 2003, no. v. 1.
- [24] A. Eryilmaz, R. Srikant, and J. R. Perkins, "Stable scheduling policies for fading wireless channels," *IEEE/ACM Trans. Network.*, vol. 13, no. 2, pp. 411–424, 2005.
- [25] L. Comtet, *Advanced Combinatorics: The Art of Finite and Infinite Expansions*. Springer Netherlands, 1974. [MR0460128](#)

SIVA THEJA MAGULURI
 MATHEMATICAL SCIENCES DEPARTMENT
 IBM T. J. WATSON RESEARCH CENTER
 YORKTOWN HEIGHTS, NY 10598
 E-MAIL: siva.theja@us.ibm.com
siva.theja@gmail.com

R. SRIKANT
 DEPARTMENT OF ECE AND CSL
 UNIVERSITY OF ILLINOIS
 AT URBANA-CHAMPAIGN
 URBANA, IL 61801
 E-MAIL: rsrikant@illinois.edu