

Asymptotically Tight Steady-State Queue Length Bounds Implied By Drift Conditions

Atila Eryilmaz · R. Srikant

the date of receipt and acceptance should be inserted later

Abstract The Foster-Lyapunov theorem and its variants serve as the primary tools for studying the stability of queueing systems. In addition, it is well known that setting the drift of the Lyapunov function equal to zero in steady-state provides bounds on the expected queue lengths. However, such bounds are often very loose due to the fact that they fail to capture resource pooling effects. The main contribution of this paper is to show that the approach of “setting the drift of a Lyapunov function equal to zero” can be used to obtain bounds on the steady-state queue lengths which are tight in the heavy-traffic limit. The key is to establish an appropriate notion of state-space collapse in terms of steady-state moments of weighted queue length differences, and use this state-space collapse result when setting the Lyapunov drift equal to zero. As an application of the methodology, we prove the steady-state equivalent of the heavy-traffic optimality result of Stolyar for wireless networks operating under the MaxWeight scheduling policy.

1 Introduction

The performance of control policies in queueing systems is evaluated by studying the sum of appropriately weighted queue lengths, either in steady-state and along almost every sample path. However, deriving optimal control policies is difficult because any stochastic optimal control formulation of the problem is often intractable. An alternative is to study the system in heavy-traffic, i.e., when the vector of exogenous arrival rates to the queueing system is close to the capacity region in the network. In such regimes, the behavior of the network often simplifies: a multi-dimensional state description of the queueing system reduces to a

A. Eryilmaz
Department of ECE
The Ohio State University
eryilmaz@osu.edu

R. Srikant
Department of ECE and CSL
University of Illinois
rsrikant@illinois.edu

single dimension (or to a small number of dimensions) and it is easier to reason about the optimality of the control policies in one dimension. This behavior of the queueing system in heavy-traffic, called *state-space collapse*, is at the heart of most heavy-traffic optimality results.

Heavy-traffic analysis of queueing systems using diffusion limits was initiated by Kingman in [21] and state-space collapse was observed for priority queues by Whitt in [42]. The use of state-space collapse to study heavy-traffic optimality was introduced by Foschini and Salz in [8] in their classic paper on join-the-shortest queue (JSQ) routing. Since then, the methodology and applicability of this technique have been extended in a number of papers; see, for example, the works of Reiman [32], Bramson [5], Williams [43], Harrison [16], Harrison and Lopez [17], and Bell and Williams [2]. This list of papers is by no means exhaustive, it is only meant to be a representative sampling of the papers in the area. Many of these papers which consider multi-queue models served by multiple resources rely on the so-called *resource pooling condition*, under which the behavior of the queueing system under study in heavy-traffic is governed by a single bottleneck resource. This results in the state-space collapse mentioned earlier, which is critical to establishing heavy-traffic optimality. In addition, these papers also implicitly assume that the scheduling policy in the queueing system is work conserving, i.e., backlogged work is served at the maximum possible rate by each station. In a seminal paper on generalized switches, Stolyar extended the notion of state-space collapse and resource pooling to systems where such per-node work-conserving policies are hard to define [36]. In particular, he showed that a class of scheduling policies called the MaxWeight policies, introduced by Tassiulas and Ephremides [38] (see [27, 31, 7, 29] for extensions) is heavy-traffic optimal in an appropriate sense. While Stolyar's work considered single-hop traffic only, the proof of heavy-traffic optimality in the multi-hop case was provided by Dai and Lin [6]. Extensions to other types of scheduling policies were presented in [35] by Shakkottai, Stolyar and Srikant. It should be noted that the MaxWeight policy was shown to be optimal at all traffic loads for a simple wireless network model (with symmetric Bernoulli arrival and service processes) by Tassiulas and Ephremides [39], who also obtained optimal policies (which are not of MaxWeight type) for wireless networks where the links are arranged in a line [40].

In general, state-space collapse does not lead to a one-dimensional state space. When the state-space collapse is to a multi-dimensional state, it is often harder to prove optimality in heavy-traffic. A model of the Internet with multi-dimensional state-space collapse has been considered by Kang, Kelly, Lee and Williams [19]. While the resource allocation policy considered there is not optimal in heavy-traffic, an important contribution there is to show that the expected workload is only a function of the number of resources in the system and not the number of flows in the system. State-space collapse is key to establishing such a result. Multi-dimensional state-space collapse has also been studied by Shah and Wischik for generalized switches [33]. In addition to recovering many of the earlier results for other models as special cases, a key contribution in [33] is the introduction and study of appropriate notions of optimality when the arrival rates lie outside the capacity region of the system. Multi-dimensional state-space collapse for a very simple four-link wireless network has been considered by Ji, Athanasopoulou and Srikant in [18] who derive the heavy-traffic optimal policy for a network which has two bottleneck resources in the heavy-traffic limit. In contrast to heavy-traffic lim-

its, Venkatramanan and Lin have shown the optimality of the MaxWeight policies in a large-deviations sense [41].

Much of work on heavy-traffic analysis of queueing systems relies on showing that a scaled version of the queue lengths in the system converges to a regulated Brownian motion. The typical result shows sample-path optimality in scaled time over a finite time interval. Often, the results allow a straightforward conjecture regarding the distribution in steady-state. Proving convergence to the steady-state distribution is an additional step which is not often undertaken. (Some exceptions are the works by Gamarnik and Zeevi [9] and the recent work of Stolyar and Yudovina [37].) For example, to establish the convergence of steady-state distributions in [36], one has to show that the limits used for the diffusion scaling and the steady-state limit (i.e., time going to ∞) can be interchanged, which can be done using the results already established in [36]. In parallel with the development of analyzing the heavy-traffic limits of queueing systems, Harrison [15] suggested directly approximating the stochastic arrival and service processes in queueing systems by Brownian motions. This theme was further developed by Laws [26], and Kelly and Laws [20]. The main idea in these papers is to study convex optimization problems suggested by flow conservation equations, along with the Brownian control problems. This theme has been influential in much of the work on heavy-traffic optimality of control policies for queueing systems.

Instead of establishing optimality of control policies, if one is simply interested in evaluating the performance of a particular policy, a common technique is to study the drift of an appropriate Lyapunov function in steady-state. Assuming that appropriate moments exist, which sometimes might be non-trivial to establish as in the work of Glynn and Zeevi [12], setting the drift equal to zero in steady-state immediately provides bounds on the steady-state moments of queue lengths. An early use of this technique was used by Kingman to derive his well-known bound on the expected waiting time in a $G/G/1$ queue [22]. This idea was pursued successfully by Kumar and Kumar [24] and by Bertsimas, Paschalidis and Tsitsiklis [4] who present many extensions of the basic idea to different types of queueing systems. The method was extended to loss models by Kumar, Srikant and Kumar in [25]. The study of Lyapunov drift to analyze performance has its roots in Markov chain stability theory using the Foster-Lyapunov theorem (see the books by Asmussen [1], Meyn [28], and Meyn and Tweedie [30]). An explicit connection between moment bounds and Lyapunov drift-based stability was provided by Kumar and Meyn in [23]. While most of these papers obtain bounds on polynomial moments, Hajek [14] and Bertsimas, Gamarnik and Tsitsiklis [3] obtained exponential-type bounds on the queue lengths. The bounds obtained by Hajek will be very useful to us in this paper.

Now that we have described prior work on heavy-traffic analysis and performance evaluation of queueing systems, we present our motivation for this paper. The Lyapunov drift-based moment bounding techniques are simple to derive since they require elementary probabilistic tools. Although the arrival and service processes are assumed to be quite simple to apply these techniques, the queueing models and control techniques can be quite complicated, thus making the techniques fairly general in their applicability. However, in many cases, the bounds obtained from these techniques are extremely loose in complex queueing systems such as wireless networks operating under the MaxWeight policy, or sometimes even in simple systems such as parallel servers to which arrivals are routed accord-

ing to the JSQ policy. Therefore, the bounds obtained from drift considerations are sometimes not very useful to evaluate control policies for queueing systems. The main reason is that, simple drift-based bounds do not exploit resource pooling effects observed in heavy-traffic. The key to introducing resource pooling effects into the drift-based arguments is to define an appropriate notion of state-space collapse that can be easily incorporated into the derivation of the drift-based moment bounds. In this paper, we present techniques for doing so, i.e., we present techniques for introducing the notion of state-space collapse into the drift-based moment bound derivations. In particular, we obtain upper bounds on the expected value of weighted functions of queue lengths for queueing systems operating under certain control policies, which coincide with lower bounds in heavy-traffic. Additionally, our techniques provide explicit bounds on moments even when the system is not in heavy-traffic, and can thus serve as a performance evaluation tool for the pre-limit systems when combined with optimization techniques such as those suggested in [34].

We consider two types of queueing systems to illustrate our methodology: parallel servers where arrivals are routed according to the JSQ policy and wireless networks operating under the MaxWeight policy. In this paper, we call these two problems, the routing problem and the scheduling problem, respectively. The scheduling problem is directly motivated by scheduling in wireless networks and high-speed networks, while the routing problem is an abstraction of multi-path routing in communication networks, wireless or wireline. The proofs of results are simpler for the JSQ problem, and so for ease of exposition, we illustrate all the steps in our derivation using JSQ first. The MaxWeight policy is harder to analyze, but once the basic idea behind the proof is presented for JSQ, the extension to MaxWeight follows when the geometrical insight in Stolyar's work [36] is translated into the Lyapunov drift framework. We note that the work of Gans and van Ryzin [10, 11] is similar in spirit to our work. They study heavy-traffic optimality in terms of steady-state moments. They first obtain a lower bound on the workload in a $G/G/1$ queue, which is similar to the Kingman-type bound [22] that we also use. However, beyond these similarities, their work is quite different in the following aspects: their upper bounds are derived for policies which make explicit use of the knowledge of the arrival rates and the topology of the capacity region (which is unknown in wireless networks with time-varying channels as we will see later). In contrast, the policies that we study do not require any of this information, and further, our Lyapunov drift-based approach appears to be quite different from the techniques used in [10, 11]. Also related is the work of Gupta and Shroff [13], who use Lyapunov-based bounds to numerically study the performance of MaxWeight algorithms. However, their bounds are not provably tight although they perform well in numerical studies and simulations. Thus, they cannot be used to establish heavy-traffic optimality.

1.1 Outline of the Methodology

The main contribution of the paper is a Lyapunov-drift based approach to obtaining bounds on steady-state queue lengths that are tight in heavy traffic. Our approach consists of three major steps:

1. Lower bound: First, we present a lower bound for the expected queue lengths in a single-server queue, and use this to identify appropriate lower bounds for both the parallel server model and the wireless network model. The lower bound that we obtain for the single-server queue is perhaps well-known and it follows the basic idea behind the Kingman bound [22]. However, we have not seen it explicitly stated, so we provide it here along with a short derivation.
2. State-space collapse: The next step is to show state-space collapse to a single dimension for both models. Unlike fluid limit proofs, in our model, the state does not actually collapse to a single dimension. What we show is that compared to the queue lengths, the differences between appropriately weighted queue lengths is small in an expected sense in steady-state.
3. Upper bound: The final step is to derive an upper bound on the expected steady-state queue lengths. For this step, as we will see, we use a natural Lyapunov function suggested by the resource pooling to be expected in the two problems. However, setting the drift of the Lyapunov function equal to zero directly does not yield a good upper bound. In addition, one has to use the state-space collapse result from the previous step to get an asymptotically tight upper bound. Once the upper bound is derived, checking that the upper and lower bounds coincide in heavy-traffic is straightforward. Even when the system is not in heavy-traffic, the upper bounds hold and can be of independent interest for performance analysis.

2 Notation, System Models and Other Preliminaries

In this section, we introduce some notation that will be used throughout the paper. We consider the control of a network of L queues that synchronously evolve in a time-slotted fashion. The evolution of the length of queue l is governed by:

$$\begin{aligned} Q_l[t+1] &= (Q_l[t] + A_l[t] - S_l[t])^+ \\ &= Q_l[t] + A_l[t] - S_l[t] + U_l[t], \quad \text{for each } l = 1, \dots, L, \end{aligned} \quad (1)$$

where $A_l[t]$ and $S_l[t]$ respectively denote the amount of arrivals and offered services to queue l in slot t , and $U_l[t] \triangleq \max(0, S_l[t] - A_l[t] - Q_l[t])$ denotes the unused service by queue l in slot t . We assume that $A_l[t]$ and $S_l[t]$ are non-negative integer valued so that $Q_l[t]$, for each l , evolves over the space of non-negative integers. For convenience, we will sometimes use boldface letters \mathbf{Q} , \mathbf{A} , \mathbf{S} , and \mathbf{U} to represent the L -dimensional vectors of queue-lengths, arrivals, offered services, and unused services, respectively. In our models in this paper, the queue length process $\{\mathbf{Q}[t]\}_{t \geq 0}$ will form a Markov chain. We will say that the queueing system is stable if this Markov chain is positive recurrent, and use $\bar{\mathbf{Q}}$ to denote the random vector whose probability distribution is the same as the steady-state distribution of $\{\mathbf{Q}[t]\}_{t \geq 0}$.

Since the queue-length vector $\mathbf{Q}[t]$ evolves over integer values in the nonnegative quadrant of the L -dimensional real vector space \mathbb{R}_+^L , we will have occasion to use the following notation in the paper and so we state them once for convenience: for two vectors $\mathbf{x} = (x_l)$ and $\mathbf{y} = (y_l)$ in \mathbb{R}^L , their inner product, Euclidean norm, and the angle between them are respectively given by

$$\langle \mathbf{x}, \mathbf{y} \rangle \triangleq \sum_{l=1}^L x_l y_l, \quad \|\mathbf{x}\| \triangleq \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle} = \sqrt{\sum_{l=1}^L x_l^2}, \quad \theta_{\mathbf{x}, \mathbf{y}} \triangleq \arccos \left(\frac{\langle \mathbf{x}, \mathbf{y} \rangle}{\|\mathbf{x}\| \|\mathbf{y}\|} \right). \quad (2)$$

We use \preceq, \prec, \succeq to denote component-wise comparison of two vectors, and $\mathbf{0}$ and $\mathbf{1}$ to denote all-zero and all-one vectors, respectively. We call two vectors *orthogonal*, denoted $\mathbf{x} \perp \mathbf{y}$, if their inner product is zero, in which case the Pythagorean Theorem applies:

$$\|\mathbf{x} + \mathbf{y}\|^2 = \|\mathbf{x}\|^2 + \|\mathbf{y}\|^2, \quad \text{for } \mathbf{x} \perp \mathbf{y}. \quad (3)$$

Additional characteristics and constraints on the arrival and service processes depend on the particular problem we will consider, which will be discussed in Sections 2.2 and 2.3 in the context of routing and scheduling, respectively. Next, we present a result that provides a bound on the steady-state moment-generating function of a random process using drift conditions.

2.1 A Useful Result

Our heavy-traffic analysis of the above models uses a result developed by Hajek [14] in a more general context. In particular, this result will be useful in proving state collapse in the sense mentioned in Section 1.1. We present it below for easy reference.

Lemma 1 *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*

$$\mathbb{E}[\Delta Z(X) | X[t] = X] \leq -\eta, \quad \text{for all } X \in \mathcal{X} \text{ with } Z(X) \geq \kappa.$$

(C2) *There exists a $D < \infty$ such that*

$$\mathbb{P}(|\Delta Z(X)| \leq D) = 1, \quad \text{for all } X \in \mathcal{X}.$$

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.

We now introduce the two queueing control problems that will be considered in this paper.

2.2 Routing Problem and the JSQ Routing Policy

In the routing problem, we consider a system of L parallel servers and their associated queues. Packets are routed to one of the queues upon arrival (see Figure 1). The goal is to find a routing policy which minimizes the total workload in the system.

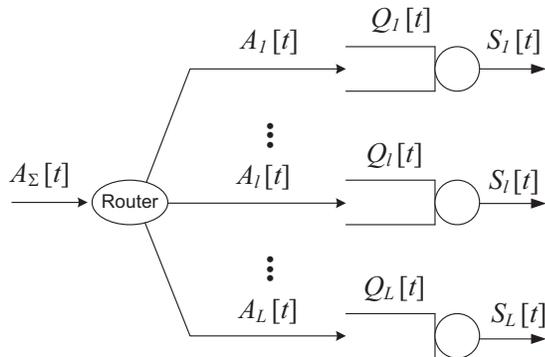


Fig. 1 System model for the routing problem.

We let $A_\Sigma[t]$ denote the random number of exogenous arrivals to the router at the beginning of slot t . The router is then required to distribute the incoming packets to L queues for service. In other words, in each time slot t , the router is expected to select a nonnegative-valued vector $\mathbf{A}[t] = (A_l[t])_l$ such that $\sum_{l=1}^L A_l[t] = A_\Sigma[t]$. Then, the queues will evolve as in (1) based on this routing decision. We make the following assumptions on the arrival and service processes.

Assumption 1 (Assumptions for the Routing Problem) *We assume that the exogenous arrival process and the service processes for different queues are independent (not necessarily identically distributed). Also, we assume that the exogenous arrival process $\{A_\Sigma[t]\}_t$ and each service process $\{S_l[t]\}_t$ is composed of a sequence of independent and identically distributed (i.i.d.) nonnegative-integer-valued and bounded random variables with $A_\Sigma[t] \leq A_{max} < \infty$ and $S_l[t] \leq S_{max} < \infty$, for all l and all t .*

We additionally use the following notation for the mean and variance of different random variables in the routing problem: for the arrival process, $\lambda_\Sigma \triangleq \mathbb{E}[A_\Sigma[1]]$, $\sigma_\Sigma^2 \triangleq \text{var}(A_\Sigma[1])$; for the service process of each queue l , $\mu_l \triangleq \mathbb{E}[S_l[1]]$, $\nu_l^2 \triangleq \text{var}(S_l[1])$; and for the hypothetical total service process, defined as $S_\Sigma[t] \triangleq \sum_{l=1}^L S_l[t]$, we let $\mu_\Sigma \triangleq \sum_{l=1}^L \mu_l$, $\nu_\Sigma^2 \triangleq \sum_{l=1}^L \nu_l^2$. Finally, we also introduce the boldface notation for the L -dimensional vector of mean and variances for the service processes: $\boldsymbol{\mu} \triangleq (\mu_l)_l$, and $\boldsymbol{\nu}^2 \triangleq (\nu_l^2)_l$. Without loss of generality¹, we assume that $\mu_{min} \triangleq \min_l \mu_l > 0$.

¹ We can eliminate any server with $\mu_l = 0$ from the system since $S_l[t] = 0$ for all t with probability one for any such server.

We are interested in the performance of a well-known and very natural routing policy, called *Join the Shortest Queue* (JSQ) Routing Policy, defined next.

Definition 1 (JSQ Routing Policy) For the routing problem introduced above (see Figure 1), in each time slot t , the *Join-Shortest-Queue (JSQ) Routing Policy* forwards all incoming packets to one of the the queues with the shortest queue-length in that time-slot, breaking ties uniformly at random, i.e., given $\mathbf{Q}[t]$, and $A_\Sigma[t]$, the arrival vector $\mathbf{A}[t]$ is selected as

$$\mathbf{A}[t] = \text{RAND} \left\{ \begin{array}{l} \arg \min \\ \{\mathbf{A} \geq \mathbf{0}; \sum_l A_l = A_\Sigma[t]\} \end{array} \langle \mathbf{A}, \mathbf{Q}[t] \rangle \right\},$$

where RAND denotes that ties are broken uniformly at random.

The principle behind the JSQ routing policy is very intuitive: it constantly tries to equalize the queue-lengths at all servers by routing all the incoming packets in each time slot to one of the smallest queues. Thus, it hopes to make sure that no server is idle when there is work to be done. If this can be achieved, then the system will behave as though all the servers pool their resources together and act like a single server queue. In this paper, we make this claim precise by using a Lyapunov-based analysis.

It is clear that the maximum achievable service rate for the parallel queue system is $\mu_\Sigma = \sum_{l=1}^L \mu_l$. A control policy is said to be throughput optimal if it can stabilize any set of arrival rates which can be stabilized by another policy. The JSQ routing policy is well known to be throughput-optimal; additionally, the following lemma states that all moments are finite in steady-state. The proof is provided in Appendix A for completeness.

Lemma 2 *If the mean exogenous arrival rate λ_Σ lies outside the capacity region \mathcal{R} , i.e., $\lambda_\Sigma > \mu_\Sigma$, then the queueing network cannot be stabilized by any feasible routing policy.*

For any λ_Σ in the interior of \mathcal{R} , i.e., $\lambda_\Sigma < \mu_\Sigma$, the JSQ Routing Policy stabilizes the network in the following strong sense: $\{\mathbf{Q}[t]\}_t$ converges in distribution to a random vector $\overline{\mathbf{Q}}$ whose all moments are bounded, i.e., there exist constants $\{M_r\}_{r=1,2,\dots}$ such that $\mathbb{E}[\|\overline{\mathbf{Q}}\|^r] = M_r$.

2.3 Scheduling Problem and the MW Scheduling Policy

In the scheduling problem, the goal is to select an instantaneous service rate vector for L queues with independent packet arrival processes, subject to certain feasibility constraints on the rates at which the queues can be simultaneously served (see Figure 2).

Here, the feasibility constraints can be used to model various types of coupling between service availability at different queues. For example, in the context of wireless networks with collision-based interference constraints, the feasibility constraints can capture interference among simultaneous wireless transmissions. Generically, we let \mathcal{S} denote the *set of feasible service rate vectors* that the controller is allowed to select from. Hence, each element $\mathbf{S} \triangleq (S_l)_l$ of \mathcal{S} is a vector of service

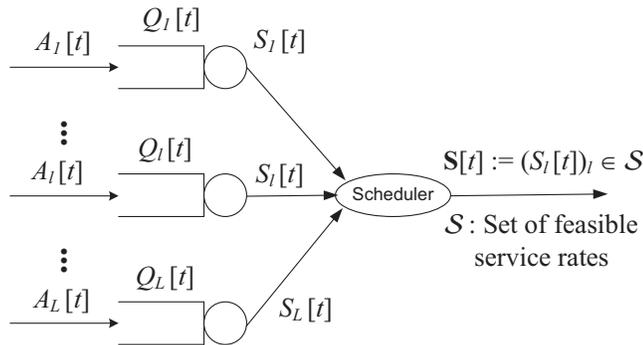


Fig. 2 System model for the scheduling problem.

rates that can be offered to the queues in one slot, and hence it contains the all zero vector $\mathbf{0}$, corresponding to no service.

In each time slot t , the scheduler is required to select a feasible service vector $\mathbf{S}[t] \in \mathcal{S}$. Then, the queues will evolve according to (1) based on this scheduling decision. We make the following assumptions on the arrival processes and the set of feasible service rates.

Assumption 2 (Assumptions for the Scheduling Problem) *We assume that the arrival processes to different queues are independent (not necessarily identically distributed). Also, we assume that each exogenous arrival process $\{A_l[t]\}_t$ is composed of a sequence of i.i.d. nonnegative-integer-valued and bounded random variables with $A_l[t] \leq A_{max} < \infty$ for all l and all t . We assume that each feasible service rate vector $\mathbf{S} \in \mathcal{S}$ is nonnegative-integer-valued and bounded with $S_l \leq S_{max} < \infty$, for all l .*

We additionally use the following notation for the mean and variance of different processes in the scheduling problem: for the exogenous arrival process for each queue l , $\lambda_l \triangleq \mathbb{E}[A_l[1]]$, $\sigma_l^2 \triangleq \text{var}(A_l[1])$; for the service process of each queue l , $\mu_l \triangleq \mathbb{E}[S_l[1]]$, $\nu_l^2 \triangleq \text{var}(S_l[1])$; for the hypothetical total arrival process, defined as $A_{\Sigma}[t] \triangleq \sum_{l=1}^L A_l[t]$, we let $\lambda_{\Sigma} \triangleq \sum_{l=1}^L \lambda_l$, $\sigma_{\Sigma}^2 \triangleq \sum_{l=1}^L \sigma_l^2$. Finally, we also introduce the boldface notation for the L -dimensional vector of mean and variances for the arrival and service processes: $\boldsymbol{\lambda} \triangleq (\lambda_l)_l$, and $\boldsymbol{\sigma}^2 \triangleq (\sigma_l^2)_l$; and $\boldsymbol{\mu} \triangleq (\mu_l)_l$, and $\boldsymbol{\nu}^2 \triangleq (\nu_l^2)_l$. Without loss of generality², we can assume that $\lambda_l > 0$, $\forall l$.

In the rest of paper, we are interested in the performance of a well-known queue-length-based scheduling policy, called *Maximum Weight* (MW) Scheduling Policy, defined next.

Definition 2 (MW Scheduling Policy) For the scheduling problem introduced above (see Figure 2), in each time slot t , the *Maximum Weight (MW) Scheduling Policy* selects the service rate vector $\mathbf{S}[t]$ from \mathcal{S} to optimize the following objective, breaking ties uniformly at random,

$$\mathbf{S}[t] = \text{RAND} \left\{ \arg \max_{\mathbf{S} \in \mathcal{S}} \langle \mathbf{Q}[t], \mathbf{S} \rangle \right\}$$

² We can eliminate any arrival with $\lambda_l = 0$ from the system since $A_l[t] = 0$ for all t with probability one for any such arrival.

where $\langle \cdot, \cdot \rangle$ is the vector inner product in \mathbb{R}^L defined in (2).

2.4 Capacity Region for the Scheduling Problem

Definition 3 (Maximum Achievable Rate (Capacity) Region for the Scheduling Problem, \mathcal{R}) For the scheduling problem described in Section 2.3 operating under Assumption 2 and the given set of feasible service rate vectors \mathcal{S} , the maximum achievable rate region \mathcal{R} is the convex hull³ of \mathcal{S} , i.e.,

$$\mathcal{R} \triangleq \text{Convex Hull}(\mathcal{S}).$$

Under the assumed finite size and nonnegative nature of the set \mathcal{S} , the region \mathcal{R} becomes a polyhedron in the nonnegative quadrant of \mathbb{R}^L , and hence can be equivalently described by

$$\mathcal{R} = \{\mathbf{r} \geq \mathbf{0} : \langle \mathbf{c}^{(k)}, \mathbf{r} \rangle \leq b^{(k)}, k = 1, \dots, K\}, \quad (4)$$

where K denotes the finite (and minimal) number of hyperplanes that fully describe the polyhedron, and the pair $(\mathbf{c}^{(k)}, b^{(k)})$ defines the k^{th} hyperplane, $\mathcal{H}^{(k)}$, through its normal vector $\mathbf{c}^{(k)} \in \mathbb{R}^L$ and its inner product value $b^{(k)} \in \mathbb{R}$ as $\mathcal{H}^{(k)} \triangleq \{\mathbf{r} : \langle \mathbf{c}^{(k)}, \mathbf{r} \rangle = b^{(k)}\}$. Further, since the allowed feasible service rates are nonnegative-valued, the pairs $(\mathbf{c}^{(k)}, b^{(k)})_k$, are assumed to satisfy:

$$\|\mathbf{c}^{(k)}\| = 1, \quad \mathbf{c}^{(k)} \succeq \mathbf{0}, \quad b^{(k)} > 0, \quad \text{for } k = 1, \dots, K.$$

Furthermore, we call the intersection of the k^{th} hyperplane with \mathcal{R} the k^{th} face $\mathcal{F}^{(k)}$ of the achievable rate region, i.e., $\mathcal{F}^{(k)} \triangleq \mathcal{H}^{(k)} \cap \mathcal{R}$. \diamond

Figures 3 and 4 illustrate the above concepts for a 2-dimensional and a 3-dimensional capacity region. The next lemma states the well-known throughput-optimality of MW Scheduling, and shows that the steady-state moments of the queue length vector are bounded. As in the routing problem, the proof of the following lemma uses Lemma 1 and is omitted.

Lemma 3 *If the arrival rate vector $\boldsymbol{\lambda} \succeq \mathbf{0}$ lies outside the capacity region \mathcal{R} , i.e., $\boldsymbol{\lambda} \notin \mathcal{R}$, then the queueing network cannot be stabilized by any feasible scheduling policy.*

For any arrival rate vector $\boldsymbol{\lambda}$ in the interior of \mathcal{R} , i.e., $\boldsymbol{\lambda} \in \text{Int}(\mathcal{R})$, the MW Scheduling Policy stabilizes the network in the following strong sense that the queue-length vector process $\{\mathbf{Q}[t]\}_t$ converges in distribution to a random vector $\overline{\mathbf{Q}}$ whose all moments are bounded, i.e., there exist finite numbers $\{M_r\}_{r=1,2,\dots}$ such that $\mathbb{E}[\|\overline{\mathbf{Q}}\|^r] = M_r$.

3 Lower Bounds

We first consider a simple single-server queueing system, depicted in Figure 5 and obtain a lower bound on the steady-state expected queue length in this system. This lower bound will then be used to obtain lower bounds for both the routing and scheduling problems.

³ The convex hull of a set $\mathcal{S} \subset \mathbb{R}^L$ is the smallest convex set that contains \mathcal{S} .

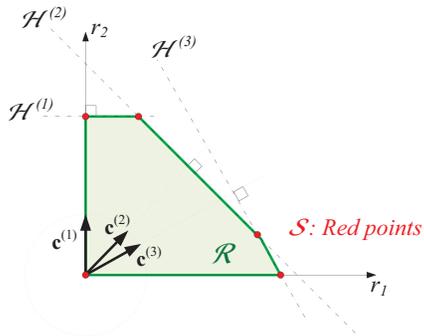


Fig. 3 The figure illustrates the capacity region for a 2-dimensional example, indicating the hyperplanes $\mathcal{H}^{(k)}$ and the associated normal vectors $\mathbf{c}^{(k)}$.

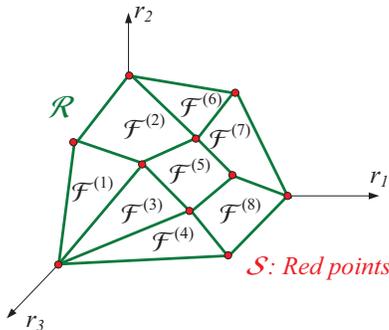


Fig. 4 The figure illustrates the capacity region for a 3-dimensional example ($L = 3$), indicating that each $\mathcal{F}^{(k)}$ is polyhedral in \mathbb{R}^{L-1} with a non-empty interior.

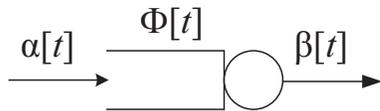


Fig. 5 Lower bounding system with i.i.d. arrival and service processes over time with bounded support.

We assume that the arrival and service processes to the single-server queue are, respectively, described by two independent sequences of i.i.d. nonnegative-valued random variables $\{\alpha[t]\}_t$ and $\{\beta[t]\}_t$. We assume that both distributions have finite support, i.e., there exist $\alpha_{max} < \infty$ and $\beta_{max} < \infty$ such that $\alpha[t] \leq \alpha_{max}$ and $\beta[t] \leq \beta_{max}$ with probability 1 for all t . We also denote the means and variances of the arrival and service processes as $\alpha \triangleq \mathbb{E}[\alpha[1]]$, $\sigma_\alpha^2 \triangleq \text{var}(\alpha[1])$, and $\beta \triangleq \mathbb{E}[\beta[1]]$, $\nu_\beta^2 \triangleq \text{var}(\beta[1])$. Then, the queue-length of the server $\Phi[t]$ evolves as

$$\Phi[t + 1] = (\Phi[t] + \alpha[t] - \beta[t])^+, \quad t \geq 0. \tag{5}$$

We are now ready to provide a useful lower bound on the steady-state performance of this system.

Lemma 4 *For the system of Figure 5 with a given service process $\{\beta[t]\}_t$, consider the arrival process $\{\alpha^{(\epsilon)}[t]\}_t$, parameterized by $\epsilon > 0$, with mean $\alpha^{(\epsilon)}$ satisfying $\epsilon = \beta - \alpha^{(\epsilon)}$, and with variance denoted as $\sigma_{\alpha^{(\epsilon)}}^2$. Let the queue-length process, denoted by $\{\Phi^{(\epsilon)}[t]\}_t$, evolve as in (5) with $\alpha[t] := \alpha^{(\epsilon)}[t]$.*

Then, $\{\Phi^{(\epsilon)}[t]\}_t$ is a positive Harris recurrent Markov Chain ([30]) for any $\epsilon > 0$, and converges in distribution to a random variable $\bar{\Phi}^{(\epsilon)}$ with all bounded moments.

Moreover, the mean of $\bar{\Phi}^{(\epsilon)}$ can be lower-bounded as

$$\mathbb{E}[\bar{\Phi}^{(\epsilon)}] \geq \frac{\zeta^{(\epsilon)}}{2\epsilon} - B_1, \quad (6)$$

where $\zeta^{(\epsilon)} \triangleq \sigma_{\alpha}^2 + \nu_{\beta}^2 + \epsilon^2$, and $B_1 \triangleq \frac{\beta_{max}}{2}$.

*Therefore, in the **heavy-traffic limit** as the mean arrival rate approaches the mean service rate from below, i.e., as $\epsilon \downarrow 0$, and assuming the variance $\sigma_{\alpha^{(\epsilon)}}^2$ converges to a constant σ_{α}^2 , the lower bounds become*

$$\liminf_{\epsilon \downarrow 0} \epsilon \mathbb{E}[\bar{\Phi}^{(\epsilon)}] \geq \frac{\zeta}{2}, \quad (7)$$

where $\zeta \triangleq \sigma_{\alpha}^2 + \nu_{\beta}^2$.

Proof In the following argument, we will omit the (ϵ) superscript for ease of exposition, and revive it when necessary. The claim that $\{\Phi[t]\}_t$ is positive Harris recurrent follows from standard negative drift conditions ([30]). Then, Lemma 1 applies to the Lyapunov function $V(\Phi) \triangleq \|\Phi\|$ as in the proofs of Lemmas 2 and 3, and is omitted here. These establish that $\Phi[t]$ converges in distribution to $\bar{\Phi}$ with all bounded moments, i.e., $\mathbb{E}[\|\bar{\Phi}\|^r] < \infty$ for each $r = 1, 2, \dots$.

To prove the lower bound (6), we first expand (5) as

$$\Phi[t+1] = \Phi[t] + \alpha[t] - \beta[t] + \chi[t], \quad t \geq 0. \quad (8)$$

where $\chi[t]$ denotes the unused service in slot t . For the quadratic Lyapunov function $W(\Phi) \triangleq \|\Phi\|^2$, the mean drift $\Delta W(\Phi) \triangleq [W(\Phi[t+1]) - W(\Phi[t])] \mathcal{I}(\Phi[t] = \Phi)$ in steady-state must be zero, i.e., $\mathbb{E}[\Delta W(\bar{\Phi})] = 0$. Next, we expand the conditional mean drift of W , omitting the time reference $[t]$ for brevity:

$$\begin{aligned} \mathbb{E}[\Delta W(\Phi) | \Phi[t] = \Phi] &= \mathbb{E}[(\Phi + \alpha - \beta + \chi)^2 - \Phi^2 | \Phi] \\ &= \mathbb{E}[(\Phi + \alpha - \beta)^2 + 2(\Phi + \alpha - \beta)\chi + \chi^2 - \Phi^2 | \Phi] \\ &= \mathbb{E}[(\alpha - \beta)^2 | \Phi] + 2\mathbb{E}[(\alpha - \beta) | \Phi]\Phi - \mathbb{E}[\chi^2 | \Phi] \end{aligned}$$

where the last step uses the fact that $\chi(\Phi + \alpha - \beta) = -\chi^2$ by definition, and the independence of the arrival and service processes from each other, and from Φ . Taking expectations of both sides with respect to the steady-state distribution of $\{\Phi[t]\}$ and using the fact that $\mathbb{E}[\Delta W(\bar{\Phi})] = 0$ yields, after re-organizing and reviving the (ϵ) notation,

$$\begin{aligned} \epsilon \mathbb{E}[\bar{\Phi}^{(\epsilon)}] &= \frac{\mathbb{E}[(\alpha^{(\epsilon)} - \beta)^2]}{2} - \frac{\mathbb{E}[\chi(\bar{\Phi}^{(\epsilon)})^2]}{2} \\ &\geq \frac{\zeta^{(\epsilon)}}{2\epsilon} - \frac{\epsilon \beta_{max}}{2}, \end{aligned}$$

where the last step follows from simple manipulations of the first term, and from the facts that $\mathbb{E}[\chi(\bar{\Phi}^{(\epsilon)})] = \epsilon$, and $\chi[t] \leq \beta_{max}$ for all t . \diamond

Next, we will discuss the implications of this lower bound on the routing and scheduling problems, respectively.

3.1 Lower Bounds for the Routing Problem

To lower bound the performance of any feasible routing policy, we assume *resource pooling* whereby we consider a hypothetical single server that serves all exogenous arrivals with a service rate of $\{S_\Sigma[t]\}$ defined as the sum of the service processes of all servers in the actual system, i.e., $S_\Sigma \triangleq \sum_{l=1}^L S_l[t]$. This results in a single-server queueing system of Figure 5 with the arrival process $\alpha[t] = A_\Sigma[t]$, for all t with $\alpha_{max} = A_{max}$, and the service process $\beta[t] = S_\Sigma[t]$, for all t with $\beta_{max} = LS_{max}$. It is then easy to see that the corresponding queue-length process $\{\Phi[t]\}_t$ is stochastically smaller than the total queue-length process $\{\sum_{l=1}^L Q_l[t]\}_t$ of the original multi-server system in Figure 1 under any feasible routing policy. Thus, utilizing the lower bounds from Lemma 4, we next establish lower bounds on the performance of any routing policy (and hence JSQ) for all arrival rates.

Lemma 5 *For the routing problem of Section 2.2 with a given service vector process $\{\mathbf{S}[t]\}_t$, consider the exogenous arrival process $\{A_\Sigma^{(\epsilon)}[t]\}_t$, parameterized by $\epsilon > 0$, with mean $\lambda_\Sigma^{(\epsilon)}$ satisfying $\epsilon = \mu_\Sigma - \lambda_\Sigma^{(\epsilon)}$, and with variance denoted as $(\sigma_\Sigma^{(\epsilon)})^2$. Accordingly, let the queue-length process under JSQ Routing with this arrival process be denoted as $\{\mathbf{Q}^{(\epsilon)}[t]\}_t$, evolving as in (1). Moreover, let $\bar{\mathbf{Q}}^{(\epsilon)}$ denote the limiting random vector with all bounded moments that the process $\{\mathbf{Q}^{(\epsilon)}[t]\}_t$ is known (by Lemma 2) to converge to, for any $\epsilon > 0$.*

Then, the first moment of the sum length of $\bar{\mathbf{Q}}^{(\epsilon)}$ can be lower-bounded, with notation $\zeta^{(\epsilon)} \triangleq (\sigma_\Sigma^{(\epsilon)})^2 + \nu_\Sigma^2 + \epsilon^2$, as:

$$\mathbb{E} \left[\sum_{l=1}^L \bar{Q}_l^{(\epsilon)} \right] \geq \frac{\zeta^{(\epsilon)}}{2\epsilon} - B_1, \quad (9)$$

where $B_1 \triangleq \frac{LS_{max}}{2}$.

*Therefore, in the **heavy-traffic limit** as the mean arrival rate approaches the mean service rate from below, i.e., as $\epsilon \downarrow 0$, and as the variance $(\sigma_\Sigma^{(\epsilon)})^2$ converges to a constant σ_Σ^2 , the lower bound becomes, with notation $\zeta \triangleq \sigma_\Sigma^2 + \nu_\Sigma^2$,*

$$\liminf_{\epsilon \downarrow 0} \epsilon \mathbb{E} \left[\sum_{l=1}^L \bar{Q}_l^{(\epsilon)} \right] \geq \frac{\zeta}{2}. \quad (10)$$

3.2 Lower Bounds for the Scheduling Problem

Unlike the routing problem, where there is a single natural lower bound on the total queue-length in the system, in the case of scheduling, the polyhedral nature

of the capacity region \mathcal{R} (see Definition 3), implies that there are many possible lower bounds.

Assume that $\lambda \in \text{Int}(\mathcal{R})$. From Definition 3, recall that the capacity region is bounded by K hyperplanes, each hyperplane $\mathcal{H}^{(k)}$ described by its normal vector $\mathbf{c}^{(k)}$ and the value $b^{(k)}$. Then, we can define the distance of λ to the k^{th} hyperplane and the closest point, respectively, as

$$\epsilon^{(k)} \triangleq \min_{\mathbf{r} \in \mathcal{H}^{(k)}} \|\lambda - \mathbf{r}\|, \quad (11)$$

$$\lambda^{(k)} \triangleq \lambda + \epsilon^{(k)} \mathbf{c}^{(k)}, \quad (12)$$

where $\epsilon^{(k)}$ is strictly positive for each k since λ is assumed to be in the interior of \mathcal{R} . We let $\epsilon \triangleq (\epsilon^{(k)})_{k=1}^K$ denote the vector of distances to all hyperplanes, which will later be used to alternatively refer to the arrival rate vector λ . Noting that $\lambda^{(k)}$ may be outside the capacity region \mathcal{R} for some hyperplanes, we next distinguish those for which $\lambda^{(k)}$ is within \mathcal{R} as

$$\mathcal{K}_\lambda \triangleq \{k \in \{1, \dots, K\} : \lambda^{(k)} \in \mathcal{R}\}. \quad (13)$$

Notice that \mathcal{K}_λ identifies the set of *dominant hyperplanes* whose closest point to λ is on the boundary of the capacity region \mathcal{R} , hence is a feasible average rate for service. We note that for any $\lambda \in \text{Int}(\mathcal{R})$, the set \mathcal{K}_λ is non-empty, and hence is well-defined. Figure 6 illustrates the distance $\epsilon^{(k)}$ and the projection $\lambda^{(k)}$ of a given arrival vector λ for the 2-dimensional rate region of Figure 3.

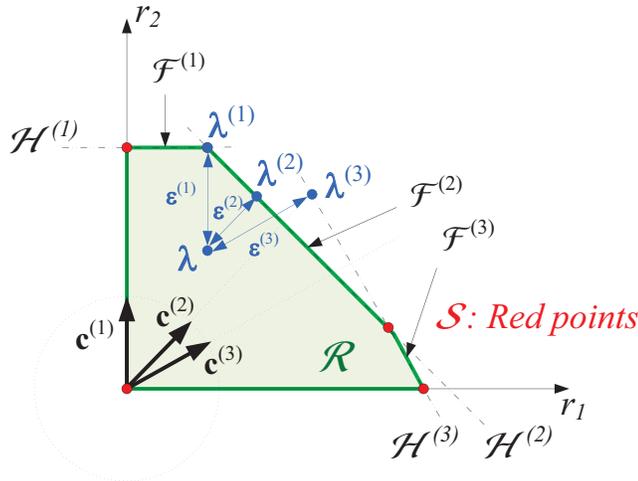


Fig. 6 Illustration of the introduced parameters $\epsilon^{(k)}$ and $\lambda^{(k)}$ in the 2-dimensional rate region of Figure 3. Note that, for the selected λ , \mathcal{K}_λ includes the first two faces, but not the third face since $\lambda^{(3)}$ is outside the rate region \mathcal{R} .

Associated with each hyperplane, say $\mathcal{H}^{(k)}$, we construct the following hypothetical single-server queue: following the notation in Figure 5 with the superscript

^(k) to distinguish different hyperplanes, we set $\alpha^{(k)}[t] = \langle \mathbf{c}^{(k)}, \mathbf{A}[t] \rangle$, where $\mathbf{A}[t]$ is the vector of arrivals to the L queues in the scheduling problem, and $\beta^{(k)}[t] = b^{(k)}$, where $b^{(k)}$ is the positive constant defining the value of the k^{th} hyperplane. Accordingly, the single server queue-length $\Phi^{(k)}[t]$ evolves as in (5) with these arrival and service process mappings. Then, $\{\Phi^{(k)}[t]\}_t$ is stochastically smaller than the queue-length process $\{\langle \mathbf{c}^{(k)}, \mathbf{Q}[t] \rangle\}_t$ under any feasible (and hence MW) scheduling policy. Hence, utilizing Lemma 4, we can find the following lower bounds on the moments of the limiting queue-length vector under MW Scheduling.

Lemma 6 *For the scheduling problem of Section 2.3 with a given set of feasible schedules \mathcal{S} , consider the exogenous arrival vector process $\{\mathbf{A}^{(\epsilon)}[t]\}_t$ with mean vector $\boldsymbol{\lambda}^{(\epsilon)} \in \text{Int}(\mathcal{R})$ and with variance vector denoted as $(\boldsymbol{\sigma}^{(\epsilon)})^2 \triangleq ((\sigma_l^{(\epsilon)})^2)_{l=1}^L$. Here, the superscript vector is given by $\boldsymbol{\epsilon} = (\epsilon^{(k)})_{k=1}^K \succ 0$, where $\epsilon^{(k)}$ and $\lambda^{(k)}$ gives the distance and the projection of $\boldsymbol{\lambda}^{(\epsilon)}$ to the k^{th} hyperplane as defined in (11) and (12), respectively. Accordingly, let the queue-length process under MW Scheduling with this arrival process be denoted as $\{\mathbf{Q}^{(\epsilon)}[t]\}_t$, evolving as in (1). Moreover, let $\overline{\mathbf{Q}}^{(\epsilon)}$ denote a random vector with the same distribution as the steady-state distribution of $\{\mathbf{Q}^{(\epsilon)}[t]\}_t$ (by Lemma 3).*

Then, for each $k \in \{1, \dots, K\}$, and with $\zeta^{(\epsilon, k)} \triangleq \langle (\mathbf{c}^{(k)})^2, (\boldsymbol{\sigma}^{(\epsilon)})^2 \rangle + (\epsilon^{(k)})^2 = \sum_{l=1}^L (c_l^{(k)})^2 (\sigma_l^{(\epsilon)})^2 + (\epsilon^{(k)})^2$,

$$\mathbb{E} \left[\langle \mathbf{c}^{(k)}, \overline{\mathbf{Q}}^{(\epsilon)} \rangle \right] \geq \frac{\zeta^{(\epsilon, k)}}{2\epsilon^{(k)}} - B_1^{(k)} \quad (14)$$

where $B_1^{(k)} \triangleq \frac{b^{(k)}}{2}$.

*Further, for each $k \in \mathcal{K}_{\boldsymbol{\lambda}^{(\epsilon)}}$, consider the **heavy-traffic limit** $\epsilon^{(k)} \downarrow 0$; and suppose that the variance vector $(\boldsymbol{\sigma}^{(\epsilon)})^2$ approaches a constant vector $\boldsymbol{\sigma}^2$. Then, defining $\zeta^{(k)} \triangleq \langle (\mathbf{c}^{(k)})^2, \boldsymbol{\sigma}^2 \rangle$, we have*

$$\liminf_{\epsilon^{(k)} \downarrow 0} \epsilon^{(k)} \mathbb{E} \left[\langle \mathbf{c}^{(k)}, \overline{\mathbf{Q}}^{(\epsilon)} \rangle \right] \geq \frac{\zeta^{(k)}}{2} \quad (15)$$

4 State-Space Collapse

In this section, we will show a moment condition which is the equivalent of state-space collapse in prior literature of fluid and diffusion limit analysis of the routing and scheduling problems. Specifically, we will show that the steady-state queue-length vector concentrates around a line within the L -dimensional state space in the following sense: the deviations from the line are bounded, independent of heavy-traffic parameter ϵ . Since the lower bounds on the queue lengths are of the order of $1/\epsilon$, this would then show that the deviations from the line are small compared to the queue lengths themselves. To establish state-space collapse, in both problems, we will first identify a vector and show that the queue lengths concentrate along this vector. For this purpose, we first introduce some notation.

Let $\mathbf{c} \succ \mathbf{0}$ be a vector with unit norm. The components of the queue length vector \mathbf{Q} parallel and perpendicular to \mathbf{c} are, respectively, given by:

$$\begin{aligned}\mathbf{Q}_{\parallel} &\triangleq \langle \mathbf{c}, \mathbf{Q} \rangle \mathbf{c}, \\ \mathbf{Q}_{\perp} &\triangleq \mathbf{Q} - \mathbf{Q}_{\parallel} = [Q_l - \langle \mathbf{c}, \mathbf{Q} \rangle c_l]_{l=1}^L.\end{aligned}$$

If \mathbf{c} represents the direction along which state-space collapse occurs, then we study the statistics of \mathbf{Q}_{\perp} to show that all its moments are bounded by constants that do not depend on the proximity of the arrival rate vector to the boundary of the capacity region. To that end, we utilize Lemma 1 by studying the drift of the Lyapunov function $V_{\perp}(\mathbf{Q}) = \|\mathbf{Q}_{\perp}\|$ to show that: (i) the drift of this Lyapunov function is always bounded; (ii) and that, when $\|\mathbf{Q}_{\perp}\|$ value is sufficiently large, it has a strictly negative drift *whose value does not depend on the distance of the arrival rate vector from the boundary of the capacity region*, denoted as ϵ and ϵ in Sections 3.1 and 3.2. This latter property will be important in establishing the heavy-traffic optimality of the schemes. It turns out that studying of the drift of $\|\mathbf{Q}_{\perp}\|$ is hard; instead, as we will see later, it is easier to study the drift of $\|\mathbf{Q}\|^2$ and $\|\mathbf{Q}_{\parallel}\|^2$. The following lemma serves as a useful preliminary step which relates the drift of these three quantities; the proof is provided in Appendix B. We will use this result in deriving our state-space collapse results later.

Lemma 7 *Consider the generic queueing system described in Section 2 with queues evolving according to (1), where the arrival and service processes are respectively bounded by A_{max} and S_{max} for each link and each time slot. Define the following Lyapunov functions:*

$$V_{\perp}(\mathbf{Q}) \triangleq \|\mathbf{Q}_{\perp}\|, \quad W(\mathbf{Q}) \triangleq \|\mathbf{Q}\|^2, \quad \text{and} \quad W_{\parallel}(\mathbf{Q}) \triangleq \|\mathbf{Q}_{\parallel}\|^2, \quad (16)$$

with the corresponding single-step drifts denoted by:

$$\Delta V_{\perp}(\mathbf{Q}) \triangleq [V_{\perp}(\mathbf{Q}[t+1]) - V_{\perp}(\mathbf{Q}[t])] \mathcal{I}(\mathbf{Q}[t] = \mathbf{Q}), \quad (17)$$

$$\Delta W(\mathbf{Q}) \triangleq [W(\mathbf{Q}[t+1]) - W(\mathbf{Q}[t])] \mathcal{I}(\mathbf{Q}[t] = \mathbf{Q}), \quad (18)$$

$$\Delta W_{\parallel}(\mathbf{Q}) \triangleq [W_{\parallel}(\mathbf{Q}[t+1]) - W_{\parallel}(\mathbf{Q}[t])] \mathcal{I}(\mathbf{Q}[t] = \mathbf{Q}). \quad (19)$$

Then, the following properties hold for the random variable $\Delta V_{\perp}(\mathbf{Q})$:

1. We can bound the drift of V_{\perp} with the drifts of W and W_{\parallel} as follows:

$$\Delta V_{\perp}(\mathbf{Q}) \leq \frac{1}{2\|\mathbf{Q}_{\perp}\|} (\Delta W(\mathbf{Q}) - \Delta W_{\parallel}(\mathbf{Q})), \quad \text{for all } \mathbf{Q} \in \mathbb{R}_+^L. \quad (20)$$

2. The drift of V_{\perp} is absolutely bounded as:

$$|\Delta V_{\perp}(\mathbf{Q})| \leq 2\sqrt{L} \max(A_{max}, S_{max}), \quad \text{for all } \mathbf{Q} \in \mathbb{R}_+^L, \quad (21)$$

where, recall from Section 2 that, A_{max} and S_{max} respectively bounds the number of arrivals to and departures from a queue, and L is the number of queues in the network.

Notice that (21) verifies Condition (C2) of Lemma 1 with

$$D = 2\sqrt{L} \max(A_{max}, S_{max})$$

for the Lyapunov function $V_{\perp}(\mathbf{Q})$. To conclude that $V_{\perp}(\mathbf{Q}[t])$ converges to $\|\overline{\mathbf{Q}}_{\perp}\|$ with finite moments, we also need to verify condition (C1) of Lemma 1. In the following subsections, we do so for both the JSQ Routing and MW Scheduling policies. Further, we will show that the moments bounds on $\|\overline{\mathbf{Q}}_{\perp}\|$ are independent of the proximity of the arrival rate vectors to the boundary of the capacity regions – a key feature to be used for the proof of heavy-traffic optimality of these policies in Section 5.

4.1 State-Space Collapse for the Routing Problem under the JSQ Policy

We recall the model introduced in Lemma 5 for the routing problem, where the exogenous arrival process $\{A_{\Sigma}^{(\epsilon)}\}_t$ is parameterized by a positive valued $\epsilon \triangleq \mu_{\Sigma} - \lambda_{\Sigma}^{(\epsilon)}$, and where the queue-length process and the limiting queue-length random vector under the JSQ Router are respectively denoted by $\{\mathbf{Q}^{(\epsilon)}[t]\}_t$ and $\overline{\mathbf{Q}}^{(\epsilon)}$. Recall that the JSQ policy attempts to equalize the queue lengths, so one would expect the state-space to collapse along the direction of a unit vector, all of whose components are equal, i.e., the vector $\mathbf{c} \triangleq \frac{\mathbf{1}}{\sqrt{L}} = (\frac{1}{\sqrt{L}})_{l=1}^L$. Then, the projection and the perpendicular of any given $\mathbf{Q}^{(\epsilon)} \in \mathbb{R}_{+}^L$ with respect to this line become:

$$\mathbf{Q}_{\parallel}^{(\epsilon)} \triangleq \frac{Q_{\Sigma}^{(\epsilon)}}{L} \mathbf{1}, \quad \text{and} \quad \mathbf{Q}_{\perp}^{(\epsilon)} \triangleq \left[Q_l - \frac{1}{L} Q_{\Sigma}^{(\epsilon)} \right]_{l=1}^L,$$

where $Q_{\Sigma}^{(\epsilon)} \triangleq \sum_{l=1}^L Q_l^{(\epsilon)}$ denotes the total number of packets in the network. We already know, from Lemma 2, that $\mathbf{Q}^{(\epsilon)}[t]$ converges to $\overline{\mathbf{Q}}^{(\epsilon)}$ in distribution for any $\epsilon > 0$. Due to continuous mapping theorem, this implies that $\mathbf{Q}_{\parallel}^{(\epsilon)}[t]$ also converges in distribution to $\overline{\mathbf{Q}}_{\parallel}^{(\epsilon)}$. The following result establishes the boundedness of all moments of the limiting $\overline{\mathbf{Q}}_{\perp}^{(\epsilon)}$ uniformly for all $\epsilon > 0$.

Proposition 1 *Consider the limiting queueing process $\overline{\mathbf{Q}}^{(\epsilon)}$ under JSQ Routing, serving the exogenous arrival process $\{A_{\Sigma}^{(\epsilon)}\}_t$ parameterized by $\epsilon \triangleq \mu_{\Sigma} - \lambda_{\Sigma}^{(\epsilon)}$. Then, for any choice of $\delta \in (0, \mu_{min})$ there exists a sequence of finite numbers $\{N_r\}_{r=1,2,\dots}$ such that $\mathbb{E} \left[\|\overline{\mathbf{Q}}_{\perp}^{(\epsilon)}\|^r \right] \leq N_r$, for all $\epsilon \in (0, (\mu_{min} - \delta)L)$ and for each $r = 1, 2, \dots$.*

Proof The proof follows from Lemma 1 applied to the Lyapunov function $V_{\perp}(\mathbf{Q}^{(\epsilon)}) \triangleq \|\mathbf{Q}_{\perp}^{(\epsilon)}\|$, where the following argument omits the superscript $^{(\epsilon)}$ for ease of exposition. Note that (21) already establishes Condition (C2) of Lemma 1. Thus, the proof is done once we verify Condition (C1) for the JSQ policy. To that end, we use the bound on $\Delta V_{\perp}(\mathbf{Q})$ in (20) and study the mean conditional drift of $\Delta W(\mathbf{Q})$ and $\Delta W_{\parallel}(\mathbf{Q})$ next.

We start by upper-bounding $\mathbb{E}[\Delta W(\mathbf{Q}) \mid \mathbf{Q}] \triangleq \mathbb{E}[\Delta W(\mathbf{Q}) \mid \mathbf{Q}[t] = \mathbf{Q}]$. In what follows, we will omit the time reference $[t]$ after the first step for brevity.

$$\begin{aligned}
\mathbb{E}[\Delta W(\mathbf{Q}) \mid \mathbf{Q}] &= \mathbb{E} \left[\|\mathbf{Q}[t+1]\|^2 - \|\mathbf{Q}\|^2 \mid \mathbf{Q} \right] \\
&= \mathbb{E} \left[\|\mathbf{Q} + \mathbf{A} - \mathbf{S} + \mathbf{U}\|^2 - \|\mathbf{Q}\|^2 \mid \mathbf{Q} \right] \\
&= \mathbb{E} \left[\|\mathbf{Q} + \mathbf{A} - \mathbf{S}\|^2 + 2\langle \mathbf{Q} + \mathbf{A} - \mathbf{S}, \mathbf{U} \rangle + \|\mathbf{U}\|^2 - \|\mathbf{Q}\|^2 \mid \mathbf{Q} \right] \\
&\stackrel{(a)}{\leq} \mathbb{E} \left[\|\mathbf{Q} + \mathbf{A} - \mathbf{S}\|^2 - \|\mathbf{Q}\|^2 \mid \mathbf{Q} \right] \\
&= \mathbb{E} \left[2\langle \mathbf{Q}, \mathbf{A} - \mathbf{S} \rangle + \|\mathbf{A} - \mathbf{S}\|^2 \mid \mathbf{Q} \right] \\
&\leq 2\mathbb{E}[\langle \mathbf{Q}, \mathbf{A} - \mathbf{S} \rangle \mid \mathbf{Q}] + K_1, \tag{22}
\end{aligned}$$

where the inequality (a) follows from the fact that $U_l(Q_l + A_l - S_l) = -U_l^2 \leq 0$, for each l , and $K_1 \triangleq L \max(A_{max}, S_{max})^2$ is bounded since both the arrival and service processes are bounded.

Next, we bound the first term in (22) by first defining a hypothetical arrival rate vector $\boldsymbol{\lambda} = (\lambda_l)_l$ with respect to the given service rate vector $\boldsymbol{\mu}$ and the $\epsilon \in (0, (\mu_{min} - \delta)L)$ as $\boldsymbol{\lambda} \triangleq \boldsymbol{\mu} - \frac{\epsilon}{\sqrt{L}}\mathbf{c}$. Note that $\lambda_{min} \triangleq \min_l \lambda_l > \delta$, where δ is a fixed constant in $(0, \mu_{min})$, and that $\sum_{l=1}^L \lambda_l = \mu_\Sigma - \epsilon$, which is identical to the assumed rate λ_Σ of the exogenous arrival process $\{A_\Sigma[t]\}_t$. The following derivation starts by adding and subtracting $\boldsymbol{\lambda}$ to the first term in (22):

$$\begin{aligned}
\mathbb{E}[\langle \mathbf{Q}, \mathbf{A} - \mathbf{S} \rangle \mid \mathbf{Q}] &= \langle \mathbf{Q}, \mathbb{E}[\mathbf{A} \mid \mathbf{Q}] - \boldsymbol{\lambda} \rangle - \langle \mathbf{Q}, \boldsymbol{\mu} - \boldsymbol{\lambda} \rangle \\
&\stackrel{(a)}{=} \langle \mathbf{Q}, \mathbb{E}[\mathbf{A} \mid \mathbf{Q}] - \boldsymbol{\lambda} \rangle - \frac{\epsilon}{\sqrt{L}} \langle \mathbf{Q}, \mathbf{c} \rangle \\
&\stackrel{(b)}{=} \mathbb{E}[A_\Sigma \mid \mathbf{Q}] Q_{min} - \langle \mathbf{Q}, \boldsymbol{\lambda} \rangle - \frac{\epsilon}{\sqrt{L}} \|\mathbf{Q}\| \\
&= \lambda_\Sigma Q_{min} - \sum_{l=1}^L \lambda_l Q_l - \frac{\epsilon}{\sqrt{L}} \|\mathbf{Q}\| \\
&= -\sum_{l=1}^L \lambda_l (Q_l - Q_{min}) - \frac{\epsilon}{\sqrt{L}} \|\mathbf{Q}\| \\
&\leq -\lambda_{min} \sum_{l=1}^L |Q_l - Q_{min}| - \frac{\epsilon}{\sqrt{L}} \|\mathbf{Q}\| \\
&= -\|\mathbf{Q} - Q_{min} \mathbf{1}\|_1 \lambda_{min} - \frac{\epsilon}{\sqrt{L}} \|\mathbf{Q}\| \\
&\stackrel{(c)}{\leq} -\|\mathbf{Q} - Q_{min} \mathbf{1}\| \lambda_{min} - \frac{\epsilon}{\sqrt{L}} \|\mathbf{Q}\| \\
&\stackrel{(d)}{\leq} -\|\mathbf{Q} - \frac{1}{L} Q_\Sigma \mathbf{1}\| \lambda_{min} - \frac{\epsilon}{\sqrt{L}} \|\mathbf{Q}\| \\
&\stackrel{(e)}{\leq} -\delta \|\mathbf{Q}_\perp\| - \frac{\epsilon}{\sqrt{L}} \|\mathbf{Q}\|
\end{aligned}$$

where

$$Q_{\min} \triangleq \min_{1 \leq m \leq L} Q_m \geq 0; \quad \|\mathbf{x}\|_1 \triangleq \sum_{l=1}^L |x_l| \text{ is the } l_1\text{-norm of } \mathbf{x},$$

and step (a) follows from the definition of λ ; (b) follows from the definitions of the JSQ policy (see Definition 1) and of \mathbf{Q}_{\parallel} ; (c) is true since l_1 norm $\|\mathbf{x}\|_1$ of any vector $\mathbf{x} \in \mathbb{R}^L$ is no smaller than its l_2 (Euclidean) norm $\|\mathbf{x}\|$; (d) is true since $\frac{1}{L}Q_{\Sigma}$ minimizes the convex function $\|\mathbf{Q} - y \mathbf{1}\|$ over $y \in \mathbb{R}$; (e) is true since $\lambda_{\min} > \delta$ by construction of λ . This bound, when substituted in (22), yields:

$$\mathbb{E}[\Delta W(\mathbf{Q}) | \mathbf{Q}] \leq -2\delta \|\mathbf{Q}_{\perp}\| - \frac{2\epsilon}{\sqrt{L}} \|\mathbf{Q}_{\parallel}\| + K_1. \quad (23)$$

Next, we lower bound $\mathbb{E}[\Delta W_{\parallel}(\mathbf{Q}) | \mathbf{Q}] \triangleq \mathbb{E}[\Delta W_{\parallel}(\mathbf{Q}) | \mathbf{Q}[t] = \mathbf{Q}]$:

$$\begin{aligned} \mathbb{E}[\Delta W_{\parallel}(\mathbf{Q}) | \mathbf{Q}] &= \mathbb{E} \left[\langle \mathbf{c}, \mathbf{Q}[t+1] \rangle^2 - \langle \mathbf{c}, \mathbf{Q}[t] \rangle^2 | \mathbf{Q} \right] \\ &= \mathbb{E} \left[\langle \mathbf{c}, \mathbf{Q} + \mathbf{A} - \mathbf{S} + \mathbf{U} \rangle^2 - \langle \mathbf{c}, \mathbf{Q} \rangle^2 | \mathbf{Q} \right] \\ &= \mathbb{E} \left[\langle \mathbf{c}, \mathbf{Q} + \mathbf{A} - \mathbf{S} \rangle^2 + 2\langle \mathbf{c}, \mathbf{Q} + \mathbf{A} - \mathbf{S} \rangle \langle \mathbf{c}, \mathbf{U} \rangle \right. \\ &\quad \left. + \langle \mathbf{c}, \mathbf{U} \rangle^2 - \langle \mathbf{c}, \mathbf{Q} \rangle^2 | \mathbf{Q} \right] \\ &= \mathbb{E} \left[2\langle \mathbf{c}, \mathbf{Q} \rangle \langle \mathbf{c}, \mathbf{A} - \mathbf{S} \rangle + \langle \mathbf{c}, \mathbf{A} - \mathbf{S} \rangle^2 \right. \\ &\quad \left. + 2\langle \mathbf{c}, \mathbf{Q} + \mathbf{A} - \mathbf{S} \rangle \langle \mathbf{c}, \mathbf{U} \rangle + \langle \mathbf{c}, \mathbf{U} \rangle^2 | \mathbf{Q} \right] \\ &= 2\langle \mathbf{c}, \mathbf{Q} \rangle \langle \mathbf{c}, \mathbb{E}[\mathbf{A} | \mathbf{Q}] - \boldsymbol{\mu} \rangle - 2\mathbb{E}[\langle \mathbf{c}, \mathbf{S} \rangle \langle \mathbf{c}, \mathbf{U} \rangle | \mathbf{Q}] \\ &\quad + \mathbb{E} \left[\langle \mathbf{c}, \mathbf{A} - \mathbf{S} \rangle^2 + 2\langle \mathbf{c}, \mathbf{Q} + \mathbf{A} \rangle \langle \mathbf{c}, \mathbf{U} \rangle + \langle \mathbf{c}, \mathbf{U} \rangle^2 | \mathbf{Q} \right] \quad (24) \\ &\geq 2\langle \mathbf{c}, \mathbf{Q} \rangle \langle \mathbf{c}, \mathbb{E}[\mathbf{A} | \mathbf{Q}] - \boldsymbol{\mu} \rangle - 2\mathbb{E}[\langle \mathbf{c}, \mathbf{S} \rangle \langle \mathbf{c}, \mathbf{U} \rangle | \mathbf{Q}], \quad (25) \end{aligned}$$

where the inequality follows from the observation that (24) ≥ 0 as $\mathbf{c} \succ \mathbf{0}$, $\mathbf{U} \succeq \mathbf{0}$, $\mathbf{Q} \succeq \mathbf{0}$. We further lower-bound the remaining terms in (25) individually. For the first term, we first recall that $\langle \mathbf{c}, \mathbf{Q} \rangle = \|\mathbf{Q}_{\parallel}\|$. Also noting that $\mathbf{c} = \mathbf{1}/\sqrt{L}$ here and that $\lambda_{\Sigma} = \mu_{\Sigma} - \epsilon$,

$$\langle \mathbf{c}, \mathbb{E}[\mathbf{A} | \mathbf{Q}] - \boldsymbol{\mu} \rangle = \frac{1}{\sqrt{L}} \sum_{l=1}^L (\mathbb{E}[A_l | \mathbf{Q}] - \mu_l) = \frac{\lambda_{\Sigma} - \mu_{\Sigma}}{\sqrt{L}} = -\frac{\epsilon}{\sqrt{L}}.$$

Finally, we can lower-bound the second term in (25) easily by $-K_2$, where $K_2 \triangleq 2LS_{max}^2$, since $S_l \leq S_{max}$ and $U_l \leq S_{max}$ for all l . Thus, using these lower bounds in (25) yields:

$$\mathbb{E}[\Delta W_{\parallel}(\mathbf{Q}) | \mathbf{Q}] \geq -\frac{2\epsilon}{\sqrt{L}} \|\mathbf{Q}_{\parallel}\| - K_2. \quad (26)$$

We now substitute the bounds (23) and (26) in (20), and cancel common terms, to bound the conditional mean drift of $V_{\perp}(\mathbf{Q})$ as

$$\mathbb{E}[\Delta V_{\perp}(\mathbf{Q}) | \mathbf{Q}] \leq -\delta + \frac{K_1 + K_2}{2\|\mathbf{Q}_{\perp}\|},$$

where δ , K_1 , and K_2 are positive constants defined above, all independent of ϵ . Note that this inequality verifies Condition (C1) of Lemma 1, and hence establishes the claimed existence of finite constants $\{N_r\}_{r=1,2,\dots}$ for which $\mathbb{E}[\|\overline{\mathbf{Q}}_{\perp}^{(\epsilon)}\|^r] \leq N_r$, for all $\epsilon \in (0, (\mu_{\min} - \delta)L)$, and for each $r = 1, 2, \dots$. \diamond

This proposition establishes that even for a heavily loaded network, where ϵ approaches zero, the steady-state queue-length vector concentrates around the line along \mathbf{c} , and the difference between the queue lengths have bounded moments. In contrast, the lower bound on the total queue length goes to infinity as ϵ goes to zero. Thus, in comparison to the total queue length, the differences in the queue lengths are negligible. This observation will be critical in Section 5 for arguing that the queueing network performs close to the lower-bounding resource-pooled system of Section 3. Since relative to the queue lengths, the differences in the queue lengths are negligible, the queue lengths can be thought as being attracted towards the line defined by the vector \mathbf{c} . We will call this line, the *line of attraction*.

4.2 State-Space Collapse for the Scheduling Problem under the MW Policy

We recall the model introduced in Lemma 6 for the scheduling problem, where the exogenous arrival vector process $\{\mathbf{A}_{\Sigma}^{(\epsilon)}\}_t$ is parameterized by the vector $\epsilon \succ \mathbf{0}$ that measures the distance of the arrival rate vector $\boldsymbol{\lambda}^{(\epsilon)} \in \text{Int}(\mathcal{R})$ to each of the K hyperplanes of \mathcal{R} . The corresponding queue-length vector process and the limiting queue-length random vector under the MW Scheduler are respectively denoted by $\{\mathbf{Q}^{(\epsilon)}[t]\}_t$ and $\overline{\mathbf{Q}}^{(\epsilon)}$.

Different from the routing case, in the scheduling problem, depending upon the arrival rate vector, the line of attraction can be different. In general, associated with each hyperplane enclosing the capacity region \mathcal{R} (cf. Definition 3) is a possible line of attraction. Recall that the pair $(\mathbf{c}^{(k)}, b^{(k)})$ defines the k^{th} hyperplane $\mathcal{H}^{(k)}$ through its unit norm normal vector $\mathbf{c}^{(k)}$ and its value $b^{(k)} > 0$. Accordingly, we define the projection and the perpendicular vector of any given $\mathbf{Q}^{(\epsilon)}$ with respect to the k^{th} normal as:

$$\begin{aligned} \mathbf{Q}_{\parallel}^{(\epsilon,k)} &\triangleq \langle \mathbf{c}^{(k)}, \mathbf{Q}^{(\epsilon)} \rangle \mathbf{c}^{(k)} \\ \mathbf{Q}_{\perp}^{(\epsilon,k)} &\triangleq \mathbf{Q}^{(\epsilon)} - \mathbf{Q}_{\parallel}^{(\epsilon,k)}. \end{aligned}$$

Also recall the notion of dominant hyperplanes $\mathcal{K}_{\boldsymbol{\lambda}^{(\epsilon)}}$ from (13) to note that for each $k \in \mathcal{K}_{\boldsymbol{\lambda}^{(\epsilon)}}$, the corresponding projection $\boldsymbol{\lambda}^{(k)}$ as defined in (12) lies on the k^{th} face $\mathcal{F}^{(k)}$ (see Definition 3) of the capacity region \mathcal{R} . Next, we further distinguish those faces in $\mathcal{K}_{\boldsymbol{\lambda}^{(\epsilon)}}$ for which $\boldsymbol{\lambda}^{(k)}$ lies in the relative interior of the face, i.e., $\boldsymbol{\lambda}^{(k)}$ is not shared by more than one hyperplane:

$$\mathcal{K}_{\boldsymbol{\lambda}^{(\epsilon)}}^{\circ} \triangleq \{k \in \mathcal{K}_{\boldsymbol{\lambda}^{(\epsilon)}} : \boldsymbol{\lambda}^{(k)} \in \text{Relint}(\mathcal{F}^{(k)})\}, \quad (27)$$

where $\text{Relint}(\mathcal{F}^{(k)})$ indicates the relative interior of the polyhedral set $\mathcal{F}^{(k)}$, which is non-empty for each k by construction⁴. It is also easy to verify that, for each

⁴ This follows from the assumed minimality of the set of hyperplanes that define the capacity region \mathcal{R} .

$\lambda^{(\epsilon)} \in \text{Int}(\mathcal{R})$, the set $\mathcal{K}_{\lambda^{(\epsilon)}}^\circ$ is non-empty. In the example of Figure 6, we can see that $\mathcal{K}_\lambda = \{1, 2\}$, while $\mathcal{K}_\lambda^\circ = \{2\}$ since the projection $\lambda^{(1)}$ is at the corner of face $\mathcal{F}^{(1)}$, hence is not in its relative interior. With this terminology, we can now present the state-space collapse result for the MW Scheduler.

Proposition 2 *Assume $\lambda^{(\epsilon)} \in \text{Int}(\mathcal{R})$ in the scheduling problem with the corresponding set of $\{\epsilon^{(k)}\}_k$ as defined in (11). Then, under the MW Scheduling Policy, for each $k \in \mathcal{K}_{\lambda^{(\epsilon)}}^\circ$, there exist finite constants $\{N_r^{(k)}\}_{r=1,2,\dots}$ such that $\mathbb{E} \left[\|\overline{\mathbf{Q}}_\perp^{(\epsilon,k)}\|^r \right] \leq N_r^{(k)}$, for all $\epsilon \succ \mathbf{0}$, each Face k belonging to $\mathcal{K}_{\lambda^{(\epsilon)}}^\circ$ as $\epsilon \downarrow \mathbf{0}$, and each $r = 1, 2, \dots$.*

Proof Assuming a given $\epsilon \succ \mathbf{0}$, the following argument omits the superscript (ϵ) for notational convenience so that $\mathbf{Q}^{(\epsilon)}, \lambda^{(\epsilon)}, \mathbf{Q}^{(\epsilon,k)}$ are respectively denoted as $\mathbf{Q}, \lambda, \mathbf{Q}^{(k)}$. For any $k \in \mathcal{K}_\lambda^\circ$, let us recall the Lyapunov functions in (16) to define their counterparts associated with the normal vector $\mathbf{c}^{(k)}$:

$$V_\perp^{(k)}(\mathbf{Q}) \triangleq \|\mathbf{Q}_\perp^{(k)}\|, \quad W(\mathbf{Q}) \triangleq \|\mathbf{Q}\|^2, \quad \text{and} \quad W_\parallel^{(k)}(\mathbf{Q}) \triangleq \|\mathbf{Q}_\parallel^{(k)}\|^2,$$

The rest of the proof follows the same line of reasoning as in the proof of Proposition 1 to bound $\Delta V_\perp^{(k)}(\mathbf{Q})$ by utilizing (20). However, before we start analyzing the mean conditional drift of $\Delta W(\mathbf{Q}^{(k)})$ and $\Delta W_\parallel(\mathbf{Q}^{(k)})$, we need to establish some basics related to the geometry of the scheduling capacity region \mathcal{R} .

Since λ is in the interior of \mathcal{R} , for each $k \in \mathcal{K}_\lambda^\circ$, we have $\lambda^{(k)}$ in the relative interior of $\mathcal{F}^{(k)}$ by definition. This means that, for each $k \in \mathcal{K}_\lambda^\circ$, there exists a small enough $\delta^{(k)} > 0$ such that the set

$$\mathcal{B}_{\delta^{(k)}}^{(k)} \triangleq \mathcal{H}^{(k)} \cap \{\mathbf{r} \in \mathbb{R}_+^L : \|\mathbf{r} - \lambda^{(k)}\| \leq \delta^{(k)}\},$$

denoting the set of vectors on the hyperplane $\mathcal{H}^{(k)}$ that are within $\delta^{(k)}$ distance from $\lambda^{(k)}$, lies strictly within the face $\mathcal{F}^{(k)}$.

We are now ready to upper-bound $\mathbb{E}[\Delta W(\mathbf{Q}) \mid \mathbf{Q}] \triangleq \mathbb{E}[\Delta W(\mathbf{Q}) \mid \mathbf{Q}[t] = \mathbf{Q}]$ exactly as in the derivation of (22) to obtain:

$$\mathbb{E}[\Delta W(\mathbf{Q}) \mid \mathbf{Q}] \leq 2 \mathbb{E}[\langle \mathbf{Q}, \mathbf{A} - \mathbf{S} \rangle \mid \mathbf{Q}] + K_1, \quad (28)$$

where $K_1 \triangleq L \max(A_{max}, S_{max})^2$. Next, we use the definition of projection $\lambda^{(k)}$ from (12) to expand the first term as:

$$\begin{aligned} \mathbb{E}[\langle \mathbf{Q}, \mathbf{A} - \mathbf{S} \rangle \mid \mathbf{Q}] &= \langle \mathbf{Q}, \lambda^{(k)} - \epsilon^{(k)} \mathbf{c}^{(k)} \rangle - \mathbb{E}[\langle \mathbf{Q}, \mathbf{S} \rangle \mid \mathbf{Q}] \\ &= -\epsilon^{(k)} \langle \mathbf{Q}, \mathbf{c}^{(k)} \rangle + \langle \mathbf{Q}, \lambda^{(k)} - \mathbb{E}[\mathbf{S} \mid \mathbf{Q}] \rangle \\ &= -\epsilon^{(k)} \|\mathbf{Q}_\parallel^{(k)}\| + \langle \mathbf{Q}, \lambda^{(k)} - \mathbb{E}[\mathbf{S} \mid \mathbf{Q}] \rangle \end{aligned} \quad (29)$$

We note that the definition of the MW Scheduler (see Definition 2) directly implies that

$$\langle \mathbf{Q}, \mathbb{E}[\mathbf{S} \mid \mathbf{Q}] \rangle = \max_{\mathbf{r} \in \mathcal{R}} \langle \mathbf{Q}, \mathbf{r} \rangle, \quad (30)$$

where \mathcal{R} is described in Definition 3. Thus, noting that the set $\mathcal{B}_{\delta^{(k)}}^{(k)} \subset \mathcal{R}$, we can upper-bound (29) as follows:

$$\begin{aligned} \mathbb{E}[\langle \mathbf{Q}, \mathbf{A} - \mathbf{S} \rangle | \mathbf{Q}] &\leq -\epsilon^{(k)} \|\mathbf{Q}_{\parallel}^{(k)}\| + \min_{\mathbf{r} \in \mathcal{B}_{\delta^{(k)}}^{(k)}} \langle \mathbf{Q}, \boldsymbol{\lambda}^{(k)} - \mathbf{r} \rangle \\ &\stackrel{(a)}{=} -\epsilon^{(k)} \|\mathbf{Q}_{\parallel}^{(k)}\| + \min_{\mathbf{r} \in \mathcal{B}_{\delta^{(k)}}^{(k)}} \langle \mathbf{Q}_{\perp}^{(k)}, \boldsymbol{\lambda}^{(k)} - \mathbf{r} \rangle \\ &\stackrel{(b)}{=} -\epsilon^{(k)} \|\mathbf{Q}_{\parallel}^{(k)}\| - \delta^{(k)} \|\mathbf{Q}_{\perp}^{(k)}\| \end{aligned}$$

where (a) follows from the facts that $\mathbf{Q}_{\perp}^{(k)} = \mathbf{Q} - \langle \mathbf{c}^{(k)}, \mathbf{Q} \rangle \mathbf{c}^{(k)}$ by definition, and that the vector $\boldsymbol{\lambda}^{(k)} - \mathbf{r}$ is perpendicular to the normal vector $\mathbf{c}^{(k)}$ since both $\boldsymbol{\lambda}^{(k)}$ and \mathbf{r} are restricted to be on the hyperplane $\mathcal{H}^{(k)}$; (b) follows from a key idea in [36] where it is noted that the inner product between $\mathbf{Q}_{\perp}^{(k)}$ is minimized when \mathbf{r} is selected on the boundary of $\mathcal{B}_{\delta^{(k)}}^{(k)}$, i.e., with $\|\boldsymbol{\lambda}^{(k)} - \mathbf{r}\| = \delta^{(k)}$, such that $\boldsymbol{\lambda}^{(k)} - \mathbf{r}$ points in the opposite direction to $\mathbf{Q}_{\perp}^{(k)}$. Substituting the above result in (28) yields our first bound:

$$\mathbb{E}[\Delta W(\mathbf{Q}) | \mathbf{Q}] \leq -2\epsilon^{(k)} \|\mathbf{Q}_{\parallel}^{(k)}\| - 2\delta^{(k)} \|\mathbf{Q}_{\perp}^{(k)}\| + K_1. \quad (31)$$

The lower-bounding argument on $\mathbb{E}[\Delta W_{\parallel}^{(k)}(\mathbf{Q}) | \mathbf{Q}] \triangleq \mathbb{E}[\Delta W_{\parallel}^{(k)}(\mathbf{Q}) | \mathbf{Q}[t] = \mathbf{Q}]$ follows the exact same steps until (24) as in the JSQ case, with the minor modification in the final expression (25) that the conditioning on \mathbf{Q} remains for the service vector \mathbf{S} instead of the arrival vector \mathbf{A} , since the MW Scheduler controls \mathbf{S} while the JSQ Router controls \mathbf{A} . This yields

$$\begin{aligned} \mathbb{E}[\Delta W_{\parallel}^{(k)}(\mathbf{Q}) | \mathbf{Q}] &\geq 2\langle \mathbf{c}^{(k)}, \mathbf{Q} \rangle \langle \mathbf{c}^{(k)}, \boldsymbol{\lambda} - \mathbb{E}[\mathbf{S} | \mathbf{Q}] \rangle - 2\mathbb{E}[\langle \mathbf{c}^{(k)}, \mathbf{S} \rangle \langle \mathbf{c}^{(k)}, \mathbf{U} \rangle | \mathbf{Q}] \\ &\stackrel{(a)}{\geq} 2\langle \mathbf{c}^{(k)}, \mathbf{Q} \rangle \langle \mathbf{c}^{(k)}, \boldsymbol{\lambda} - \mathbb{E}[\mathbf{S} | \mathbf{Q}] \rangle - K_2 \\ &\stackrel{(b)}{=} -2\epsilon^{(k)} \|\mathbf{Q}_{\parallel}^{(k)}\| + 2\|\mathbf{Q}_{\parallel}^{(k)}\| \min_{\mathbf{r} \in \mathcal{R}} \langle \mathbf{c}, \boldsymbol{\lambda}^{(k)} - \mathbf{r} \rangle - K_2 \\ &\stackrel{(c)}{\geq} -2\epsilon^{(k)} \|\mathbf{Q}_{\parallel}^{(k)}\| + 2\|\mathbf{Q}_{\parallel}^{(k)}\| \left(\langle \mathbf{c}, \boldsymbol{\lambda}^{(k)} \rangle - b^{(k)} \right) - K_2 \\ &\stackrel{(d)}{\geq} -2\epsilon^{(k)} \|\mathbf{Q}_{\parallel}^{(k)}\| - K_2, \end{aligned} \quad (32)$$

where step: (a) is true for $K_2 \triangleq 2LS_{max}^2$; (b) uses the definitions of projection $\boldsymbol{\lambda}^{(k)}$ and $\mathbf{Q}_{\parallel}^{(k)}$ together with the fact that the MW Scheduler achieves (30); (c) is true since $\{\mathbf{r} \succeq \mathbf{0} : \langle \mathbf{c}^{(k)}, \mathbf{r} \rangle \leq b^{(k)}\} \supset \mathcal{R}$ from (4); and (d) follows from the fact that $\langle \mathbf{c}^{(k)}, \boldsymbol{\lambda}^{(k)} \rangle = b^{(k)}$ since $\boldsymbol{\lambda}^{(k)} \in \mathcal{H}^{(k)}$.

Finally, utilizing the bounds (31) and (32) in (20), and canceling common terms, yields the following bound on the conditional mean drift of $V_{\perp}^{(k)}(\mathbf{Q})$ as:

$$\mathbb{E}[\Delta V_{\perp}^{(k)}(\mathbf{Q}) | \mathbf{Q}] \leq -\delta^{(k)} + \frac{K_1 + K_2}{2\|\mathbf{Q}_{\perp}^{(k)}\|}, \quad \text{for each } k \in \mathcal{K}_{\boldsymbol{\lambda}}^{\circ},$$

where $\delta^{(k)}$, K_1 , and K_2 are positive constants defined above, all independent of ϵ . By verifying Condition (C1) of Lemma 1, this result proves the claimed existence

of finite constants $\{N_r^{(k)}\}_{r=1,2,\dots}$ such that $\mathbb{E} \left[\|\overline{\mathbf{Q}}_{\perp}^{(\epsilon,k)}\|^r \right] \leq N_r^{(k)}$, for all $\epsilon \succ \mathbf{0}$, each $k \in \mathcal{K}_{\lambda}^{\circ}(\epsilon)$, and each $r = 1, 2, \dots$. \diamond

The above proposition provides bounds on the deviation of $\overline{\mathbf{Q}}$ from different lines of attraction corresponding to the normals of different hyperplanes in $\mathcal{K}_{\lambda}^{\circ}(\epsilon)$. However, in the heavy-traffic limit as $\lambda^{(\epsilon)}$ approaches the relative interior of one of the faces, only one of the hyperplanes becomes a dominant face and therefore, there will be a single line of attraction associated with the normal of that face only, as one would expect.

5 Upper Bounds and Heavy-Traffic Optimality

The purpose of this section is to derive upper-bounds on the steady-state queue-length through simple Lyapunov-based methods that will be shown to be asymptotically tight under heavy-traffic, i.e., when the arrival rates will approach the boundary of the capacity regions. In particular, we use the state-space collapse results established in Section 4.

As in the case of the state-space collapse results, here too the analysis is similar for both the JSQ and MW Policies. The derivation of the upper bounds for both cases use the following lemma.

Lemma 8 *Consider the generic queueing system described in Section 2 with queues evolving according to (1), where the arrival $\mathbf{A}[t]$ and service $\mathbf{S}[t]$ vectors at time t are allowed to depend on $\mathbf{Q}[t]$. To indicate this dependence, we will use the notation $\mathbf{A}(\mathbf{Q})$ and $\mathbf{S}(\mathbf{Q})$ to refer to these two processes. Suppose $\{\mathbf{Q}[t]\}_t$ converges in distribution to a valid random vector $\overline{\mathbf{Q}}$, with all bounded moments, i.e., $\mathbb{E}[\|\overline{\mathbf{Q}}\|^r] < \infty$ for each $r = 1, 2, \dots$. Then, for any positive vector $\mathbf{c} \in \mathbb{R}_{++}^L$, we have*

$$\mathbb{E} [\langle \mathbf{c}, \overline{\mathbf{Q}} \rangle \langle \mathbf{c}, \mathbf{S}(\overline{\mathbf{Q}}) - \mathbf{A}(\overline{\mathbf{Q}}) \rangle] = \frac{\mathbb{E} [\langle \mathbf{c}, \mathbf{A}(\overline{\mathbf{Q}}) - \mathbf{S}(\overline{\mathbf{Q}}) \rangle^2]}{2} + \frac{\mathbb{E} [\langle \mathbf{c}, \mathbf{U}(\overline{\mathbf{Q}}) \rangle^2]}{2} \quad (33)$$

$$+ \mathbb{E} [\langle \mathbf{c}, \overline{\mathbf{Q}} + \mathbf{A}(\overline{\mathbf{Q}}) - \mathbf{S}(\overline{\mathbf{Q}}) \rangle \langle \mathbf{c}, \mathbf{U}(\overline{\mathbf{Q}}) \rangle], \quad (34)$$

where $\mathbf{U}(\overline{\mathbf{Q}})$ is the random vector of unused service (cf. (1)) in steady state.

Proof Recall the definition of $W_{\parallel}(\mathbf{Q}) \triangleq \|\mathbf{Q}_{\parallel}\|^2$ from (16) and the subsequent definition of its one-step drift $\Delta W_{\parallel}(\mathbf{Q})$. We can expand this drift expression as follows:

$$\begin{aligned} \Delta W_{\parallel}(\mathbf{Q}) &= \left[\langle \mathbf{c}, \mathbf{Q}[t+1] \rangle^2 - \langle \mathbf{c}, \mathbf{Q}[t] \rangle^2 \right] \mathcal{I}(\mathbf{Q}[t] = \mathbf{Q}) \\ &= \langle \mathbf{c}, \mathbf{Q} + \mathbf{A} - \mathbf{S} + \mathbf{U} \rangle^2 - \langle \mathbf{c}, \mathbf{Q} \rangle^2 \\ &= \langle \mathbf{c}, \mathbf{Q} + \mathbf{A} - \mathbf{S} \rangle^2 + 2\langle \mathbf{c}, \mathbf{Q} + \mathbf{A} - \mathbf{S} \rangle \langle \mathbf{c}, \mathbf{U} \rangle + \langle \mathbf{c}, \mathbf{U} \rangle^2 - \langle \mathbf{c}, \mathbf{Q} \rangle^2 \\ &= 2\langle \mathbf{c}, \mathbf{Q} \rangle \langle \mathbf{c}, \mathbf{A} - \mathbf{S} \rangle + \langle \mathbf{c}, \mathbf{A} - \mathbf{S} \rangle^2 + 2\langle \mathbf{c}, \mathbf{Q} + \mathbf{A} - \mathbf{S} \rangle \langle \mathbf{c}, \mathbf{U} \rangle \\ &\quad + \langle \mathbf{c}, \mathbf{U} \rangle^2 \end{aligned}$$

Now, since $\|\mathbf{Q}_{\parallel}\|^2 \leq \|\mathbf{Q}\|^2$ and $\mathbb{E}[\|\overline{\mathbf{Q}}\|^2] < \infty$, we clearly have $\mathbb{E}[W_{\parallel}(\overline{\mathbf{Q}})] < \infty$. Hence, in steady state, ΔW_{\parallel} must have a zero mean. Thus setting $\mathbb{E}[\Delta W_{\parallel}(\overline{\mathbf{Q}})] = 0$ gives the desired result. \diamond

The state-space collapse result is crucial in bounding the expectation in (34). To provide some intuition, we first note the following useful property concerning the unused service.

Lemma 9 *For a queueing network evolving according to (1) and for any given vector $\mathbf{c} \succeq \mathbf{0}$ with $\|\mathbf{c}\| = 1$, the following property always holds for any policy:*

$$\langle \mathbf{c}, \mathbf{Q}[t+1] \rangle \langle \mathbf{c}, \mathbf{U}[t] \rangle = \langle -\mathbf{Q}_\perp[t+1], \mathbf{U}[t] \rangle, \quad (35)$$

where $\mathbf{Q}_\perp[t+1] = \mathbf{Q}[t+1] - \langle \mathbf{c}, \mathbf{Q}[t+1] \rangle \mathbf{c}$.

Proof We use the notation $\mathbf{Q}^+ \triangleq \mathbf{Q}[t+1]$ and omit $[t]$ for the remaining parameters, $\mathbf{Q}[t]$, $\mathbf{A}[t]$, $\mathbf{S}[t]$, $\mathbf{U}[t]$. Accordingly, we define the projections and perpendiculars of \mathbf{Q}^+ and \mathbf{U} with respect to \mathbf{c} as:

$$\mathbf{Q}_\parallel^+ \triangleq \langle \mathbf{c}, \mathbf{Q}^+ \rangle \mathbf{c}, \quad \mathbf{Q}_\perp^+ \triangleq \mathbf{Q}^+ - \mathbf{Q}_\parallel^+, \quad \mathbf{U}_\parallel \triangleq \langle \mathbf{c}, \mathbf{U} \rangle \mathbf{c}, \quad \mathbf{U}_\perp \triangleq \mathbf{U} - \mathbf{U}_\parallel.$$

The following argument uses the fact that $Q_l^+ U_l = 0$ for all l , since either $U_l[t]$ or $Q_l[t+1]$ must be zero for each t . Hence, \mathbf{U} is orthogonal to \mathbf{Q}^+ . We can now prove the claim:

$$\begin{aligned} \langle \mathbf{c}, \mathbf{Q}^+ \rangle \langle \mathbf{c}, \mathbf{U} \rangle &= \|\mathbf{Q}_\parallel^+\| \|\mathbf{U}_\parallel\| \stackrel{(a)}{=} \langle \mathbf{Q}_\parallel^+, \mathbf{U}_\parallel \rangle \\ &= \langle \mathbf{Q}^+ - \mathbf{Q}_\perp^+, \mathbf{U}_\parallel \rangle \stackrel{(b)}{=} \langle \mathbf{Q}^+, \mathbf{U} - \mathbf{U}_\perp \rangle \\ &\stackrel{(c)}{=} \langle \mathbf{Q}_\perp^+ + \mathbf{Q}_\parallel^+, -\mathbf{U}_\perp \rangle \stackrel{(d)}{=} \langle -\mathbf{Q}_\perp^+, \mathbf{U}_\perp \rangle, \\ &\stackrel{(e)}{=} \langle -\mathbf{Q}_\perp^+, \mathbf{U} \rangle, \end{aligned}$$

where step: (a) corresponds to the equality case of the Cauchy-Schwartz inequality since $\mathbf{Q}_\parallel^+ \succeq \mathbf{0}$ and $\mathbf{U}_\parallel \succeq \mathbf{0}$ are aligned by definition; (b) follows from the facts that $\mathbf{Q}_\perp^+ \perp \mathbf{U}_\parallel$ and $\mathbf{U}_\parallel = \mathbf{U} - \mathbf{U}_\perp$; (c) utilizes the facts that $\mathbf{Q}^+ \perp \mathbf{U}$ and that $\mathbf{Q}^+ = \mathbf{Q}_\perp^+ + \mathbf{Q}_\parallel^+$; (d) uses the fact that $\mathbf{Q}_\perp^+ \perp \mathbf{U}_\perp$; and (e) uses the facts that $\mathbf{U} = \mathbf{U}_\perp + \mathbf{U}_\parallel$ and $\mathbf{Q}_\perp^+ \perp \mathbf{U}_\parallel$. \diamond

Now, we present some intuition regarding the usefulness of the state-space collapse results (cf. Propositions 1 and 2). Using (35) and noting that $\mathbf{Q}[t+1] = \mathbf{Q}[t] + \mathbf{A}[t] - \mathbf{S}[t] + \mathbf{U}[t]$, we can re-write (34) as:

$$\begin{aligned} (34) &= \mathbb{E} [\langle \mathbf{c}, \overline{\mathbf{Q}}[t+1] \rangle \langle \mathbf{c}, \mathbf{U}(\overline{\mathbf{Q}}) \rangle] - \mathbb{E} [\langle \mathbf{c}, \mathbf{U}(\overline{\mathbf{Q}}) \rangle^2] \\ &\leq \mathbb{E} [\langle \mathbf{c}, \overline{\mathbf{Q}}[t+1] \rangle \langle \mathbf{c}, \mathbf{U}(\overline{\mathbf{Q}}) \rangle] \\ &= \mathbb{E} [\langle -\overline{\mathbf{Q}}_\perp[t+1], \mathbf{U}(\overline{\mathbf{Q}}) \rangle] \\ &\stackrel{(a)}{\leq} \sqrt{\mathbb{E} [\|\overline{\mathbf{Q}}_\perp[t+1]\|^2] \mathbb{E} [\|\mathbf{U}(\overline{\mathbf{Q}})\|^2]} \\ &\stackrel{(b)}{=} \sqrt{\mathbb{E} [\|\overline{\mathbf{Q}}_\perp\|^2] \mathbb{E} [\|\mathbf{U}(\overline{\mathbf{Q}})\|^2]}, \end{aligned} \quad (36)$$

where step (a) follows from Cauchy-Schwartz inequality and step (b) is true since the distributions of $\mathbf{Q}[t+1]$ and $\mathbf{Q}[t]$ are the same in steady state. The last expression reveals the intuition behind why the state-space collapse results are useful: the expression vanishes as $\epsilon \downarrow 0$ since $\mathbb{E} [\|\overline{\mathbf{Q}}_\perp\|^2]$ is uniformly bounded, and unused service goes to zero as $\epsilon \downarrow 0$. Next, we will build on this intuition to prove heavy-traffic optimality of both the JSQ Routing and MW Scheduling Policies.

5.1 Upper Bounds and Heavy-Traffic Optimality of JSQ Routing

The following proposition presents a bound on the steady-state total queue-length under JSQ routing, and establishes the first-moment heavy-traffic optimality of JSQ as the network load approaches the capacity of the network.

Proposition 3 *Consider the routing problem under the exogenous arrival process $\{A_\Sigma^{(\epsilon)}[t]\}_t$ with mean rate $\lambda_\Sigma^{(\epsilon)} \in \text{Int}(\mathcal{R})$ satisfying $\epsilon \triangleq \mu_\Sigma - \lambda_\Sigma > 0$, and with variance $(\sigma_\Sigma^{(\epsilon)})^2$. Then, under the JSQ Routing Policy, the limiting steady-state queue-length vector $\bar{\mathbf{Q}}^{(\epsilon)}$ satisfies*

$$\mathbb{E} \left[\sum_{l=1}^L \bar{Q}_l^{(\epsilon)} \right] \leq \frac{\zeta^{(\epsilon)}}{2\epsilon} + \bar{B}_1^{(\epsilon)}, \quad (37)$$

where we recall that $\zeta^{(\epsilon)} \triangleq (\sigma_\Sigma^{(\epsilon)})^2 + \nu_\Sigma^2$ as it is defined in Section 3.1, and $\bar{B}_1^{(\epsilon)}$ is $o(\frac{1}{\epsilon})$, i.e., $\lim_{\epsilon \downarrow 0} \epsilon \bar{B}_1^{(\epsilon)} = 0$.

Also, in the **heavy traffic limit**, where we consider a sequence of exogenous arrival processes $\{A_\Sigma^{(\epsilon)}[t]\}_t$ with $\epsilon \downarrow 0$ so that $\lambda_\Sigma^{(\epsilon)}$ approaches μ_Σ and $(\sigma_\Sigma^{(\epsilon)})^2$ approaches a constant σ_Σ^2 , we have

$$\limsup_{\epsilon \downarrow 0} \epsilon \mathbb{E} \left[\sum_{l=1}^L \bar{Q}_l^{(\epsilon)} \right] \leq \frac{\zeta}{2}, \quad (38)$$

where $\zeta \triangleq \sigma_\Sigma^2 + \nu_\Sigma^2$.

Hence, comparing the heavy-traffic lower-bound (7) for any feasible policy to the heavy-traffic upper-bound (38) for JSQ Router establishes the first moment heavy-traffic optimality of JSQ Routing Policy.

Proof Recalling the definition of $\mathbf{c} \triangleq \frac{\mathbf{1}}{\sqrt{L}}$ in the JSQ case, we first note a useful fact:

$$\mathbb{E}[\langle \mathbf{c}, \mathbf{U}(\bar{\mathbf{Q}}^{(\epsilon)}) \rangle] = \frac{\mu_\Sigma - \lambda_\Sigma}{\sqrt{L}} = \frac{\epsilon}{\sqrt{L}}, \quad (39)$$

which follows from the fact that the mean drift of $\langle \mathbf{c}, \mathbf{Q}^{(\epsilon)} \rangle$ must be zero in steady state. Next, we will temporarily omit the superscript $^{(\epsilon)}$ for ease of exposition and study the terms in (33) and (34) under JSQ operation.

$$\mathbb{E}[\langle \mathbf{c}, \bar{\mathbf{Q}} \rangle \langle \mathbf{c}, \mathbf{S}(\bar{\mathbf{Q}}) - \mathbf{A}(\bar{\mathbf{Q}}) \rangle] = \left(\frac{\mu_\Sigma - \lambda_\Sigma}{\sqrt{L}} \right) \mathbb{E}[\langle \mathbf{c}, \bar{\mathbf{Q}} \rangle] = \frac{\epsilon}{L} \mathbb{E} \left[\sum_{l=1}^L \bar{Q}_l \right], \quad (40)$$

which follows from the independence of the *total* exogenous arrival process and individual service rate processes from the queue-length levels.

$$\mathbb{E} \left[\langle \mathbf{c}, \mathbf{A}(\bar{\mathbf{Q}}) - \mathbf{S}(\bar{\mathbf{Q}}) \rangle^2 \right] = \frac{1}{L} \mathbb{E} \left[(A_\Sigma - S_\Sigma)^2 \right] = \frac{(\sigma_\Sigma^2 + \nu_\Sigma^2 + \epsilon^2)}{L}, \quad (41)$$

where we recall that σ_Σ^2 and ν_Σ^2 are respectively the variances of the exogenous arrival process $\{A_\Sigma[t]\}_t$ and the hypothetical total service process defined as $S_\Sigma[t] \triangleq \sum_{l=1}^L S_l[t]$.

$$\mathbb{E} \left[\langle \mathbf{c}, \mathbf{U}(\bar{\mathbf{Q}}) \rangle^2 \right] \leq \langle \mathbf{c}, S_{max} \mathbf{1} \rangle \mathbb{E} [\langle \mathbf{c}, \mathbf{U}(\bar{\mathbf{Q}}) \rangle] = \frac{\epsilon S_{max}}{L}, \quad (42)$$

where we use the bound $U_l \leq S_{max}$ for all l , and the identity (39).

Finally, to bound (34), we take the same steps as in the argument leading to (36), and use the facts that $\mathbf{c} = \mathbf{1}/\sqrt{L}$ and $U_l \leq S_{max}$ for all l , to get:

$$\begin{aligned} (34) &\leq \sqrt{\mathbb{E} [\|\bar{\mathbf{Q}}_\perp\|^2]} \sqrt{\mathbb{E} [\|\mathbf{U}(\bar{\mathbf{Q}})\|^2]} \\ &\leq \sqrt{\mathbb{E} [\|\bar{\mathbf{Q}}_\perp\|^2]} S_{max} \sqrt{L} \mathbb{E} [\langle \mathbf{c}, \mathbf{U}(\bar{\mathbf{Q}}) \rangle] \\ &\leq \sqrt{\epsilon N_2 S_{max}}, \end{aligned} \quad (43)$$

where, in the last step, we used (39) and N_2 from Proposition 1.

We reintroduce the superscript (ϵ) to highlight the dependence on ϵ and substitute (40)-(43) in (33)-(34) to get, after minor algebraic manipulations,

$$\mathbb{E} \left[\sum_{l=1}^L \bar{Q}_l^{(\epsilon)} \right] \leq \frac{((\sigma_\Sigma^{(\epsilon)})^2 + \nu_\Sigma^2 + \epsilon^2)}{2\epsilon} + \frac{S_{max}}{2} + L \sqrt{\frac{N_2 S_{max}}{\epsilon}} = \frac{\zeta^{(\epsilon)}}{2\epsilon} + \bar{B}_1^{(\epsilon)},$$

where $\zeta^{(\epsilon)} \triangleq (\sigma_\Sigma^{(\epsilon)})^2 + \nu_\Sigma^2 + \epsilon^2$ and $\bar{B}_1^{(\epsilon)} \triangleq L \sqrt{\frac{N_2 S_{max}}{\epsilon}} + \frac{S_{max}}{2}$, which is $o(1/\epsilon)$ as claimed. Then, (38) follows immediately by taking the limit of both sides. \diamond

5.2 Upper Bounds and Heavy-Traffic Optimality of MW Scheduling

The following proposition yields upper bounds on the steady-state weighted total queue-length under MW Scheduling, and then establishes the first-moment heavy-traffic optimality of MWS as the network load approaches the boundary of the capacity region of the network.

Proposition 4 *Consider the scheduling problem under the exogenous arrival vector process $\{\mathbf{A}^{(\epsilon)}[t]\}_t$ with mean vector $\boldsymbol{\lambda}^{(\epsilon)} \in \text{Int}(\mathcal{R})$ and with variance vector $(\boldsymbol{\sigma}^{(\epsilon)})^2 \triangleq ((\sigma_l^{(\epsilon)})^2)_{l=1}^L$. The parameter $\boldsymbol{\epsilon} = (\epsilon^{(k)})_{k=1}^K \succ 0$, where $\epsilon^{(k)}$ gives the distance and the projection of $\boldsymbol{\lambda}^{(\epsilon)}$ to the k^{th} hyperplane as defined in (11). Then, under the MW Scheduling, the limiting steady-state queue-vector $\bar{\mathbf{Q}}^{(\epsilon)}$ satisfies, for each $k \in \mathcal{K}_{\boldsymbol{\lambda}^{(\epsilon)}}^\circ$,*

$$\mathbb{E} \left[\langle \mathbf{c}^{(k)}, \bar{\mathbf{Q}}^{(\epsilon)} \rangle \right] \leq \frac{\zeta^{(\epsilon,k)}}{2\epsilon^{(k)}} + \bar{B}_1^{(\epsilon,k)}, \quad (44)$$

where we recall that $\zeta^{(\epsilon,k)} \triangleq \langle \mathbf{c}^{(k)} \rangle^2, (\boldsymbol{\sigma}^{(\epsilon)})^2$ as it is defined in Lemma 6, and $\bar{B}_1^{(\epsilon,k)}$ is $o\left(\frac{1}{\epsilon^{(k)}}\right)$, i.e., $\lim_{\epsilon^{(k)} \downarrow 0} \epsilon^{(k)} \bar{B}_1^{(\epsilon,k)} = 0$.

Also, in the **heavy traffic limit**, where we consider a sequence of exogenous arrival processes $\{\mathbf{A}^{(\epsilon,k)}[t]\}_t$ with their mean vector $\boldsymbol{\lambda}_A^{(\epsilon)}$ approaching the k^{th} dominant face

in $\mathcal{K}_{\lambda^{(\epsilon)}}^{\circ}$ along its normal $\mathbf{c}^{(k)}$, and variance vectors $(\boldsymbol{\sigma}_{\mathbf{A}}^{(\epsilon)})^2$ approaching a constant vector $\boldsymbol{\sigma}^2$, we have

$$\limsup_{\epsilon^{(k)} \downarrow 0} \epsilon^{(k)} \mathbb{E} \left[\langle \mathbf{c}^{(k)}, \overline{\mathbf{Q}}^{(\epsilon)} \rangle \right] \leq \frac{\zeta^{(k)}}{2}, \quad (45)$$

where $\zeta^{(k)} \triangleq \langle (\mathbf{c}^{(k)})^2, \boldsymbol{\sigma}^2 \rangle$.

Hence, comparing the heavy-traffic lower-bound (15) for any feasible policy to the heavy-traffic upper-bound (45) for MW Scheduler establishes the first moment heavy-traffic optimality of MW Scheduling Policy.

Proof We temporarily omit the superscript (ϵ) associated with the arrival and queue-length processes for ease of exposition, and consider any $k \in \mathcal{K}_{\lambda}^{\circ}$. Before we embark upon the investigation of (33) and (34) under MW Scheduler operation, we make several useful remarks. We first utilize the fact that the mean drift of $\langle \mathbf{c}^{(k)}, \mathbf{Q} \rangle$ must be zero under steady-state operation to get:

$$\begin{aligned} \mathbb{E}[\langle \mathbf{c}^{(k)}, \mathbf{U}(\overline{\mathbf{Q}}) \rangle] &= \langle \mathbf{c}^{(k)}, \mathbb{E}[\mathbf{S}(\overline{\mathbf{Q}})] \rangle - \langle \mathbf{c}^{(k)}, \boldsymbol{\lambda} \rangle \\ &\stackrel{(a)}{=} \langle \mathbf{c}^{(k)}, \mathbb{E}[\mathbf{S}(\overline{\mathbf{Q}})] \rangle - (b^{(k)} - \epsilon^{(k)}) \\ &\stackrel{(b)}{\leq} \epsilon^{(k)}, \end{aligned} \quad (46)$$

where (a) follows from (12) and the fact that $\langle \mathbf{c}^{(k)}, \boldsymbol{\lambda}^{(k)} \rangle = b^{(k)}$ since $\boldsymbol{\lambda}^{(k)} \in \mathcal{H}^{(k)}$; and (b) follows from the facts that $\mathbb{E}[\mathbf{S}(\overline{\mathbf{Q}})]$ must be in \mathcal{R} and that $\langle \mathbf{c}^{(k)}, \mathbf{r} \rangle \leq b^{(k)}$ for all $\mathbf{r} \in \mathcal{R}$ by (4).

Next, we define

$$\pi^{(k)} \triangleq \mathbb{P} \left(\langle \mathbf{c}^{(k)}, \mathbf{S}(\overline{\mathbf{Q}}) \rangle = b^{(k)} \right), \quad (47)$$

to be the fraction of time that the service rate vector is selected from face $\mathcal{F}^{(k)}$ by the MW Scheduler in steady-state. Also, we define

$$\gamma^{(k)} \triangleq \min \{ b^{(k)} - \langle \mathbf{c}^{(k)}, \mathbf{r} \rangle : \text{for all } \mathbf{r} \in \mathcal{S} \setminus \mathcal{F}^{(k)} \}. \quad (48)$$

Since the set \mathcal{S} is discrete and finite, $\gamma^{(k)}$ is a *strictly positive* number with a constant value (independent of ϵ) associated with the geometry of the capacity region \mathcal{R} . This constant helps us establish the following useful claim associated with $\pi^{(k)}$.

Claim 1 For any $\epsilon^{(k)} \in (0, \gamma^{(k)})$, we have

$$(1 - \pi^{(k)}) \leq \frac{\epsilon^{(k)}}{\gamma^{(k)}} \quad (49)$$

where the upper-bound is $O(\epsilon^{(k)})$, i.e., vanishes as $\epsilon^{(k)} \downarrow 0$.

Proof (Claim 1) We start with the observation that

$$\mathbb{E}[\langle \mathbf{c}^{(k)}, \mathbf{S}(\overline{\mathbf{Q}}) \rangle] \geq \langle \mathbf{c}^{(k)}, \boldsymbol{\lambda} \rangle = b^{(k)} - \epsilon^{(k)}$$

where the inequality follows from the stability of the queueing network (cf. Proposition 3), and the equality follows from the utilization of (12). We can split the left-hand-side into two parts by using the definition of $\pi^{(k)}$ to write

$$\pi^{(k)}b^{(k)} + \mathbb{E} \left[\langle \mathbf{c}^{(k)}, \mathbf{S}(\overline{\mathbf{Q}}) \rangle \mathcal{I} \left(\langle \mathbf{c}^{(k)}, \mathbf{S}(\overline{\mathbf{Q}}) \rangle \neq b^{(k)} \right) \right] \geq (b^{(k)} - \epsilon^{(k)}),$$

which, when re-arranged, leads to the following lower-bound on the expectation:

$$\mathbb{E} \left[\langle \mathbf{c}^{(k)}, \mathbf{S}(\overline{\mathbf{Q}}) \rangle \mathcal{I} \left(\langle \mathbf{c}^{(k)}, \mathbf{S}(\overline{\mathbf{Q}}) \rangle \neq b^{(k)} \right) \right] \geq b^{(k)} (1 - \pi^{(k)}) - \epsilon^{(k)}, \quad (50)$$

Separately, we can upper-bound the same expectation as

$$\begin{aligned} \mathbb{E} \left[\langle \mathbf{c}^{(k)}, \mathbf{S}(\overline{\mathbf{Q}}) \rangle \mathcal{I} \left(\langle \mathbf{c}^{(k)}, \mathbf{S}(\overline{\mathbf{Q}}) \rangle = b^{(k)} \right) \right] &\leq (b^{(k)} - \gamma^{(k)}) \mathbb{E} \left[\mathcal{I} \left(\langle \mathbf{c}^{(k)}, \mathbf{S}(\overline{\mathbf{Q}}) \rangle \neq b^{(k)} \right) \right] \\ &= (b^{(k)} - \gamma^{(k)}) (1 - \pi^{(k)}), \end{aligned}$$

where the inequality follows from the definition of $\gamma^{(k)}$ in (48), and from the equality from the definition of $\pi^{(k)}$ in (47). Using this bound together with (50) yields (49). \diamond

Claim 1 implies the following additional fact:

$$\begin{aligned} &\mathbb{E} \left[\left(b^{(k)} - \langle \mathbf{c}^{(k)}, \mathbf{S}(\overline{\mathbf{Q}}) \rangle \right)^2 \right] \\ &= (1 - \pi^{(k)}) \mathbb{E} \left[\left(b^{(k)} - \langle \mathbf{c}^{(k)}, \mathbf{S}(\overline{\mathbf{Q}}) \rangle \right)^2 \mid \left(\langle \mathbf{c}^{(k)}, \mathbf{S}(\overline{\mathbf{Q}}) \rangle \neq b^{(k)} \right) \right] \\ &\leq \frac{\epsilon^{(k)}}{\gamma^{(k)}} \left((b^{(k)})^2 + \langle \mathbf{c}^{(k)}, S_{max} \mathbf{1} \rangle^2 \right) \\ &= O(\epsilon^{(k)}) \end{aligned} \quad (51)$$

where the inequality follows from (49) and the fact that $S_l \leq S_{max}$ for all l . This result establishes in a certain probabilistic sense that $\langle \mathbf{c}^{(k)}, \mathbf{S}(\overline{\mathbf{Q}}) \rangle$ is close $b^{(k)}$ if $\epsilon^{(k)}$ is small. This result confirms the intuition that when $\epsilon^{(k)}$ is small, i.e., when $\boldsymbol{\lambda}$ is close to the face $\mathcal{F}^{(k)}$, the MW Scheduler must mostly select service rates on $\mathcal{F}^{(k)}$ so that the average service rate vector exceeds the given arrival rate vector componentwise to ensure stability.

Our final remark before studying (33) and (34) concerns the geometry of the scheduling capacity region \mathcal{R} . Since the number of possible rate vectors is finite, the number of faces in the rate region is finite. Therefore, for each face $\mathcal{F}^{(k)}$ of the region \mathcal{R} , there exists an angle $\theta^{(k)} \in (0, \pi/2]$ such that

$$\langle \mathbf{c}^{(k)}, \mathbf{S}(\mathbf{Q}) \rangle = b^{(k)}, \quad \text{for all } \mathbf{Q} \text{ such that } \frac{\|\mathbf{Q}^{(k)}\|}{\|\mathbf{Q}\|} \geq \cos(\theta^{(k)}), \quad (52)$$

where $\mathbf{S}(\mathbf{Q})$ is the service rate vector selected by the MW Scheduler for the given \mathbf{Q} as in Definition 2. Note that $\theta^{(k)}$ identifies a *cone* around the line $\mathbf{c}^{(k)}$ such that any \mathbf{Q} in the cone leads to a rate allocation on the face $\mathcal{F}^{(k)}$.

We are now ready to study each term in (33) and (34) with $\mathbf{c} := \mathbf{c}^{(k)}$ to establish (44).

$$\begin{aligned}
& \mathbb{E} \left[\langle \mathbf{c}^{(k)}, \overline{\mathbf{Q}} \rangle \langle \mathbf{c}^{(k)}, \mathbf{S}(\overline{\mathbf{Q}}) - \mathbf{A}(\overline{\mathbf{Q}}) \rangle \right] \\
&= \mathbb{E} \left[\|\overline{\mathbf{Q}}_{\parallel}\| (b^{(k)} - \langle \mathbf{c}^{(k)}, \boldsymbol{\lambda} \rangle) - \mathbb{E} \left[\|\overline{\mathbf{Q}}_{\parallel}\| (b^{(k)} - \langle \mathbf{c}^{(k)}, \mathbf{S}(\overline{\mathbf{Q}}) \rangle) \right] \right] \\
&\stackrel{(a)}{=} \epsilon^{(k)} \mathbb{E} \left[\|\overline{\mathbf{Q}}_{\parallel}\| \right] - \mathbb{E} \left[\|\overline{\mathbf{Q}}_{\parallel}\| \cos(\theta_{\overline{\mathbf{Q}}, \overline{\mathbf{Q}}_{\parallel}^{(k)}}) (b^{(k)} - \langle \mathbf{c}^{(k)}, \mathbf{S}(\overline{\mathbf{Q}}) \rangle) \right] \\
&\stackrel{(b)}{=} \epsilon^{(k)} \mathbb{E} \left[\|\overline{\mathbf{Q}}_{\parallel}\| \right] - \mathbb{E} \left[\|\overline{\mathbf{Q}}_{\parallel}\| \cos(\theta_{\overline{\mathbf{Q}}, \overline{\mathbf{Q}}_{\parallel}^{(k)}}) \mathcal{I}(\theta_{\overline{\mathbf{Q}}, \overline{\mathbf{Q}}_{\parallel}^{(k)}} > \theta^{(k)}) (b^{(k)} - \langle \mathbf{c}^{(k)}, \mathbf{S}(\overline{\mathbf{Q}}) \rangle) \right] \\
&\stackrel{(c)}{=} \epsilon^{(k)} \mathbb{E} \left[\|\overline{\mathbf{Q}}_{\parallel}\| \right] - \mathbb{E} \left[\|\overline{\mathbf{Q}}_{\perp}^{(k)}\| \cot(\theta_{\overline{\mathbf{Q}}, \overline{\mathbf{Q}}_{\parallel}^{(k)}}) \mathcal{I}(\theta_{\overline{\mathbf{Q}}, \overline{\mathbf{Q}}_{\parallel}^{(k)}} > \theta^{(k)}) (b^{(k)} - \langle \mathbf{c}^{(k)}, \mathbf{S}(\overline{\mathbf{Q}}) \rangle) \right] \\
&\stackrel{(d)}{\geq} \epsilon^{(k)} \mathbb{E} \left[\|\overline{\mathbf{Q}}_{\parallel}\| \right] - \mathbb{E} \left[\|\overline{\mathbf{Q}}_{\perp}^{(k)}\| \mathcal{I}(\theta_{\overline{\mathbf{Q}}, \overline{\mathbf{Q}}_{\parallel}^{(k)}} > \theta^{(k)}) (b^{(k)} - \langle \mathbf{c}^{(k)}, \mathbf{S}(\overline{\mathbf{Q}}) \rangle) \right] \cot(\theta^{(k)}) \\
&\geq \epsilon^{(k)} \mathbb{E} \left[\|\overline{\mathbf{Q}}_{\parallel}\| \right] - \mathbb{E} \left[\|\overline{\mathbf{Q}}_{\perp}^{(k)}\| (b^{(k)} - \langle \mathbf{c}^{(k)}, \mathbf{S}(\overline{\mathbf{Q}}) \rangle) \right] \cot(\theta^{(k)}) \\
&\stackrel{(e)}{\geq} \epsilon^{(k)} \mathbb{E} \left[\|\overline{\mathbf{Q}}_{\parallel}\| \right] - \cot(\theta^{(k)}) \sqrt{\mathbb{E} \left[\|\overline{\mathbf{Q}}_{\perp}^{(k)}\|^2 \right] \mathbb{E} \left[(b^{(k)} - \langle \mathbf{c}^{(k)}, \mathbf{S}(\overline{\mathbf{Q}}) \rangle)^2 \right]} \\
&\stackrel{(f)}{\geq} \epsilon^{(k)} \mathbb{E} \left[\|\overline{\mathbf{Q}}_{\parallel}\| \right] - \cot(\theta^{(k)}) \sqrt{\frac{\epsilon^{(k)} N_2^{(k)}}{\gamma^{(k)}} \left((b^{(k)})^2 + \langle \mathbf{c}^{(k)}, S_{max} \mathbf{1} \rangle^2 \right)}, \tag{53}
\end{aligned}$$

where the step (a) follows from the definitions of $\epsilon^{(k)}$ in (11) and of the angle $\theta_{\mathbf{x}, \mathbf{y}}$ between two vectors \mathbf{x} and \mathbf{y} given in (2); (b) is true from the definition of $\theta^{(k)}$; (c) is true since $\|\overline{\mathbf{Q}}_{\perp}^{(k)}\| = \|\overline{\mathbf{Q}}_{\parallel}\| \sin(\theta_{\overline{\mathbf{Q}}, \overline{\mathbf{Q}}_{\parallel}^{(k)}})$; (d) is true since cotangent function is a decreasing nonnegative-valued function in $(0, \pi/2]$; (e) follows from Cauchy-Schwartz Inequality; and (f) follows from Proposition 2 and (51). We note that, in (53), the first term is $O(\epsilon^{(k)})$ while the second term is $O(\sqrt{\epsilon^{(k)}})$.

Next, we bound the first right-hand-side term in (33): $\mathbb{E} \left[\langle \mathbf{c}^{(k)}, \mathbf{A}(\overline{\mathbf{Q}}) - \mathbf{S}(\overline{\mathbf{Q}}) \rangle^2 \right]$

$$\begin{aligned}
&\stackrel{(a)}{=} \mathbb{E} \left[\left(\langle \mathbf{c}^{(k)}, \mathbf{A} \rangle - b^{(k)} \right)^2 \right] + \mathbb{E} \left[\left(b^{(k)} - \langle \mathbf{c}^{(k)}, \mathbf{S}(\overline{\mathbf{Q}}) \rangle \right)^2 \right] \\
&\quad + 2 \left(\langle \mathbf{c}^{(k)}, \boldsymbol{\lambda} \rangle - b^{(k)} \right) \mathbb{E} \left[b^{(k)} - \langle \mathbf{c}^{(k)}, \mathbf{S}(\overline{\mathbf{Q}}) \rangle \right] \\
&\stackrel{(b)}{=} \mathbb{E} \left[\left(\langle \mathbf{c}^{(k)}, \mathbf{A} \rangle - b^{(k)} \right)^2 \right] + \mathbb{E} \left[\left(b^{(k)} - \langle \mathbf{c}^{(k)}, \mathbf{S}(\overline{\mathbf{Q}}) \rangle \right)^2 \right] - 2 \epsilon^{(k)} \mathbb{E} \left[b^{(k)} - \langle \mathbf{c}^{(k)}, \mathbf{S}(\overline{\mathbf{Q}}) \rangle \right] \\
&\stackrel{(c)}{\leq} \mathbb{E} \left[\left(\langle \mathbf{c}^{(k)}, \mathbf{A} - \boldsymbol{\lambda} \rangle + \langle \mathbf{c}^{(k)}, \boldsymbol{\lambda} \rangle - b^{(k)} \right)^2 \right] + \mathbb{E} \left[\left(b^{(k)} - \langle \mathbf{c}^{(k)}, \mathbf{S}(\overline{\mathbf{Q}}) \rangle \right)^2 \right] \\
&\stackrel{(d)}{=} \mathbb{E} \left[\langle \mathbf{c}^{(k)}, \mathbf{A} - \boldsymbol{\lambda} \rangle^2 \right] + 2 \epsilon^{(k)} \langle \mathbf{c}^{(k)}, \mathbb{E}[\mathbf{A}] - \boldsymbol{\lambda} \rangle + (\epsilon^{(k)})^2 + \mathbb{E} \left[\left(b^{(k)} - \langle \mathbf{c}^{(k)}, \mathbf{S}(\overline{\mathbf{Q}}) \rangle \right)^2 \right] \\
&\stackrel{(e)}{\leq} \langle \mathbf{c}^{(k)} \rangle^2 + \boldsymbol{\sigma}^2 + (\epsilon^{(k)})^2 + \frac{\epsilon^{(k)}}{\gamma^{(k)}} \left((b^{(k)})^2 + \langle \mathbf{c}^{(k)}, S_{max} \mathbf{1} \rangle^2 \right) \\
&\stackrel{(f)}{=} \zeta^{(\boldsymbol{\epsilon}, k)} + \frac{\epsilon^{(k)}}{\gamma^{(k)}} \left((b^{(k)})^2 + \langle \mathbf{c}^{(k)}, S_{max} \mathbf{1} \rangle^2 \right) \tag{54}
\end{aligned}$$

where the step (a) follows simply from expanding the square after adding and subtracting $b^{(k)}$, and noting that the exogenous arrival rate vector \mathbf{A} has mean $\boldsymbol{\lambda}$ and is independent of the service rate vector $\mathbf{S}(\overline{\mathbf{Q}})$; (b) follows from the definition

of $\epsilon^{(k)}$ in (11); (c) follows from adding and subtracting $\boldsymbol{\lambda}$ in the first expression, and from noting that $b^{(k)} - \langle \mathbf{c}^{(k)}, \mathbf{r} \rangle \geq 0$ for any $\mathbf{r} \in \mathcal{R}$ and hence for $\mathbf{S}(\overline{\mathbf{Q}})$; (d) follows, again, from the definition of $\epsilon^{(k)}$; (e) follows from (51) and uses the notation $\boldsymbol{\sigma}^2$ for the variance vector for the arrival process; and (f) uses the definition of the parameter $\zeta^{(\boldsymbol{\epsilon}, k)}$. Notice that all the terms except the first term vanishes in (54) as $\epsilon^{(k)} \downarrow 0$.

Next, we bound the last term in (33) using (46) and the fact that $U_l \leq S_{max}$ for all l :

$$\mathbb{E} \left[\langle \mathbf{c}^{(k)}, \mathbf{U}(\overline{\mathbf{Q}}) \rangle^2 \right] \leq \langle \mathbf{c}, S_{max} \mathbf{1} \rangle \mathbb{E} \left[\langle \mathbf{c}^{(k)}, \mathbf{U}(\overline{\mathbf{Q}}) \rangle \right] \leq \epsilon^{(k)} \langle \mathbf{c}^{(k)}, S_{max} \mathbf{1} \rangle, \quad (55)$$

which is also vanishing as $\epsilon^{(k)} \downarrow 0$.

Finally, we consider the term (34). While the argument essentially follows that of (36), we need to pay more attention to the zero entries of $\mathbf{c}^{(k)}$, which did not exist in the JSQ case. To that end, we define $\mathcal{L}_{++}^{(k)} \triangleq \{l \in \{1, \dots, L\} : c_l^{(k)} > 0\}$ to denote the strictly positive entries of $\mathbf{c}^{(k)}$. Then, we focus on only these components by defining the following restricted vectors living in the $|\mathcal{L}_{++}^{(k)}|$ -dimensional real space:

$$\tilde{\mathbf{c}}^{(k)} \triangleq (c_l^{(k)})_{l \in \mathcal{L}_{++}^{(k)}}, \quad \tilde{\mathbf{Q}} \triangleq (Q_l)_{l \in \mathcal{L}_{++}^{(k)}}, \quad \tilde{\mathbf{U}} \triangleq (U_l)_{l \in \mathcal{L}_{++}^{(k)}}.$$

For convenience, we will denote $\tilde{\mathbf{Q}}[t+1] = (Q_l[t+1])_{l \in \mathcal{L}_{++}^{(k)}}$ as $\tilde{\mathbf{Q}}^+$. In this reduced space, we further define the projection and perpendicular of $\tilde{\mathbf{Q}}^+$ and $\tilde{\mathbf{U}}$ with respect to $\tilde{\mathbf{c}}^{(k)}$ as:

$$\tilde{\mathbf{Q}}_{\parallel}^+ \triangleq \langle \tilde{\mathbf{c}}^{(k)}, \tilde{\mathbf{Q}}^+ \rangle \tilde{\mathbf{c}}^{(k)}, \quad \tilde{\mathbf{Q}}_{\perp}^+ \triangleq \tilde{\mathbf{Q}}^+ - \tilde{\mathbf{Q}}_{\parallel}^+, \quad \tilde{\mathbf{U}}_{\parallel} \triangleq \langle \tilde{\mathbf{c}}^{(k)}, \tilde{\mathbf{U}} \rangle \tilde{\mathbf{c}}^{(k)}, \quad \tilde{\mathbf{U}}_{\perp} \triangleq \tilde{\mathbf{U}} - \tilde{\mathbf{U}}_{\parallel}.$$

Now, since $\tilde{\mathbf{c}}^{(k)} \succ \mathbf{0}$ satisfies $\|\tilde{\mathbf{c}}^{(k)}\| = 1$, the statement of Lemma 9 applies to $\tilde{\mathbf{c}}^{(k)}$, $\tilde{\mathbf{Q}}^+$, and $\tilde{\mathbf{U}}$ to yield

$$\langle \tilde{\mathbf{c}}^{(k)}, \tilde{\mathbf{Q}}^+ \rangle \langle \tilde{\mathbf{c}}^{(k)}, \tilde{\mathbf{U}} \rangle = \langle -\tilde{\mathbf{Q}}_{\perp}^+, \tilde{\mathbf{U}} \rangle.$$

This result, together with the fact that $\langle \mathbf{c}^{(k)}, \mathbf{Q}^+ \rangle \langle \mathbf{c}^{(k)}, \mathbf{U} \rangle = \langle \tilde{\mathbf{c}}^{(k)}, \tilde{\mathbf{Q}}^+ \rangle \langle \tilde{\mathbf{c}}^{(k)}, \tilde{\mathbf{U}} \rangle$, where we used the notation $\mathbf{Q}^+ \triangleq \mathbf{Q}[t+1]$, allows us to bound (34) as follows (in the following we temporarily use $\mathbb{E}_{\overline{\mathbf{Q}}}$ to imply that \mathbf{Q} is distributed as $\overline{\mathbf{Q}}$ in the expectation):

$$\begin{aligned} (34) &\leq \mathbb{E}_{\overline{\mathbf{Q}}} \left[\langle \mathbf{c}^{(k)}, \mathbf{Q}^+ \rangle \langle \mathbf{c}^{(k)}, \mathbf{U} \rangle \right] \\ &= \mathbb{E}_{\overline{\mathbf{Q}}} \left[\langle \tilde{\mathbf{c}}^{(k)}, \tilde{\mathbf{Q}}^+ \rangle \langle \tilde{\mathbf{c}}^{(k)}, \tilde{\mathbf{U}} \rangle \right] \\ &= \mathbb{E}_{\overline{\mathbf{Q}}} \left[\langle -\tilde{\mathbf{Q}}_{\perp}^+, \tilde{\mathbf{U}} \rangle \right] \\ &\stackrel{(a)}{\leq} \sqrt{\mathbb{E}_{\overline{\mathbf{Q}}} \left[\|\tilde{\mathbf{Q}}_{\perp}^+\|^2 \right]} \sqrt{\mathbb{E}_{\overline{\mathbf{Q}}} \left[\|\tilde{\mathbf{U}}\|^2 \right]} \\ &\stackrel{(b)}{\leq} \sqrt{\mathbb{E}_{\overline{\mathbf{Q}}} \left[\|\tilde{\mathbf{Q}}_{\perp}^+\|^2 \right]} \sqrt{\mathbb{E}_{\overline{\mathbf{Q}}} \left[\|\tilde{\mathbf{U}}\|^2 \right]}, \end{aligned} \quad (56)$$

where (a) follows from Cauchy-Schwartz inequality, and (b) is true since the distribution of \mathbf{Q}^+ is the same as \mathbf{Q} under steady-state. Next, we bound the expectations in (56).

The first expectation of (56) satisfies:

$$\mathbb{E}_{\overline{\mathbf{Q}}} \left[\|\tilde{\mathbf{Q}}_{\perp}\|^2 \right] \leq \mathbb{E} \left[\|\overline{\mathbf{Q}}_{\perp}\|^2 \right] \leq N_2^{(k)}, \quad (57)$$

where the first inequality follows from the fact that \mathbf{Q}_{\perp} and $\tilde{\mathbf{Q}}_{\perp}$ are equal for all positions $\mathcal{L}_{++}^{(k)}$, while \mathbf{Q}_{\perp} potentially contains additional non-zero components, and hence cannot be smaller in magnitude. The second inequality follows from Proposition 2 that establishes the state-space collapse of MWS.

The second expectation of (56) satisfies:

$$\mathbb{E}_{\overline{\mathbf{Q}}} \left[\|\tilde{\mathbf{U}}\|^2 \right] = \mathbb{E}_{\overline{\mathbf{Q}}} \left[\sum_{l \in \mathcal{L}_{++}^{(k)}} \tilde{U}_l^2 \right] \leq \frac{S_{max}}{c_{min}^{(k)}} \mathbb{E} \left[\langle \mathbf{c}^{(k)}, \mathbf{U}(\overline{\mathbf{Q}}) \rangle \right], \quad (58)$$

where $c_{min}^{(k)} \triangleq \min_{m \in \mathcal{L}_{++}^{(k)}} c_m^{(k)} > 0$. Here, the last inequality is true since $c_l^{(k)} \geq c_{min}^{(k)}$

for all $l \in \mathcal{L}_{++}^{(k)}$, and $\tilde{U}_l \leq S_{max}$ for all l .

Substituting the bounds (57) and (58) back in (56) together with the fact (46) yields:

$$(34) \leq \sqrt{\epsilon^{(k)} N_2^{(k)} \frac{S_{max}}{c_{min}^{(k)}}} \quad (59)$$

To complete, we reintroduce the $(\epsilon), (\epsilon, k)$ superscript to emphasize the dependence on ϵ and substitute the derived bounds (53), (54), (55), (59) in (33) and (34) to get

$$\mathbb{E} \left[\langle \mathbf{c}^{(k)}, \overline{\mathbf{Q}}^{(\epsilon)} \rangle \right] \leq \frac{\zeta^{(\epsilon, k)}}{2\epsilon^{(k)}} + \overline{B}_1^{(\epsilon, k)},$$

where

$$\begin{aligned} \overline{B}_1^{(\epsilon, k)} \triangleq & \cot(\theta^{(k)}) \sqrt{\frac{N_2^{(k)} \left((b^{(k)})^2 + \langle \mathbf{c}^{(k)}, S_{max} \mathbf{1} \rangle^2 \right)}{\epsilon^{(k)} \gamma^{(k)}}} \\ & + \frac{\left((b^{(k)})^2 + \langle \mathbf{c}^{(k)}, S_{max} \mathbf{1} \rangle^2 \right)}{2\gamma^{(k)}} + \frac{\langle \mathbf{c}^{(k)}, S_{max} \mathbf{1} \rangle}{2} + \sqrt{\frac{N_2^{(k)} S_{max}}{\epsilon^{(k)} c_{min}^{(k)}}}, \end{aligned}$$

which is $o(1/\epsilon^{(k)})$ as claimed. Then, (45) immediately follows by taking the limit as $\epsilon^{(k)}$. \diamond

6 Some Extensions of the Results on the Scheduling Problem

In this section, we obtain bounds on the n^{th} moment of the steady-state queue lengths for the scheduling problem, and we also discuss how to handle channel fading in the derivations of the bounds. Both of these extensions introduce some challenges, but can be essentially addressed by the methodology presented in the previous sections. The n^{th} moment analysis can also be performed for JSQ routing, but is omitted here since it is quite similar to the scheduling case.

6.1 n^{th} Moment Analysis

We follow the steps outlined in Section 1.1 to first develop lower bounds on the n^{th} moment of steady-state queue-lengths, and utilize the state-space collapse result of Section 4 to find corresponding upper bounds. Then, we will establish, as before, the heavy-traffic optimality of these policies by showing that the appropriate dominant terms of the lower and upper bounds match as the arrival rate vector approaches one of the faces of the capacity region \mathcal{R} .

1. Lower Bounds on the n^{th} moment: We first extend the approach applied in Section 3 to the n^{th} moment of the lower-bounding system in Figure 5.

Lemma 10 *For the system of Figure 5 with a given service process $\{\beta[t]\}_t$, consider the arrival process $\{\alpha^{(\epsilon)}[t]\}_t$, parameterized by $\epsilon > 0$, with mean $\alpha^{(\epsilon)}$ satisfying $\epsilon = \beta - \alpha^{(\epsilon)}$, and with variance denoted as $\sigma_{\alpha^{(\epsilon)}}^2$. Let the queue-length process, denoted by $\{\Phi^{(\epsilon)}[t]\}_t$, evolve as in (5) with $\alpha[t] := \alpha^{(\epsilon)}[t]$.*

Then, $\{\Phi^{(\epsilon)}[t]\}_t$ is a positive Harris recurrent Markov Chain ([30]) for any $\epsilon > 0$, and converges in distribution to a random variable $\bar{\Phi}^{(\epsilon)}$ with all bounded moments.

Moreover, the n^{th} moment of $\bar{\Phi}^{(\epsilon)}$ can be lower-bounded as

$$\epsilon^n \mathbb{E} \left[\left(\bar{\Phi}^{(\epsilon)} \right)^n \right] \geq n! \left(\frac{\zeta^{(\epsilon)}}{2} \right)^n - B_n^{(\epsilon)}, \quad n \geq 1, \quad (60)$$

where $B_n^{(\epsilon)}$ vanishes with $\epsilon \downarrow 0$, i.e., $\lim_{\epsilon \downarrow 0} B_n^{(\epsilon)} = 0$

Therefore, in the heavy-traffic limit as the mean arrival rate approaches the mean service rate from below, i.e., as $\epsilon \downarrow 0$, and assuming the variance $\sigma_{\alpha^{(\epsilon)}}^2$ converges to a constant σ_α^2 , the lower bounds become

$$\liminf_{\epsilon^n \downarrow 0} \epsilon^n \mathbb{E} \left[\left(\bar{\Phi}^{(\epsilon)} \right)^n \right] \geq n! \left(\frac{\zeta}{2} \right)^n, \quad n \geq 1, \quad (61)$$

where $\zeta \triangleq \sigma_\alpha^2 + \nu_\beta^2$.

Proof See Appendix C. ◇

Lemma 10 reveals an interesting fact that, in the heavy-traffic limit, the dominant terms of the n^{th} moment of $\bar{\Phi}^{(\epsilon)}$ only depends on ζ , which in turn depends only on the variances of the arrival and service processes. This is consistent with Brownian approximations, which utilize central limit theorem to approximate the system behavior using the first two moments of the arrival and service processes.

We can now apply this generic result to the scheduling problem using the same construction as in Section 3.2.

Lemma 11 *For the scheduling problem of Section 2.3 with a given set of feasible schedules \mathcal{S} , consider the exogenous arrival vector process $\{\mathbf{A}^{(\epsilon)}[t]\}_t$ with mean vector $\boldsymbol{\lambda}^{(\epsilon)} \in \text{Int}(\mathcal{R})$ and with variance vector denoted as $(\boldsymbol{\sigma}^{(\epsilon)})^2 \triangleq \left((\sigma_l^{(\epsilon)})^2 \right)_{l=1}^L$. Here, the superscript vector is given by $\boldsymbol{\epsilon} = (\epsilon^{(k)})_{k=1}^K \succ 0$, where $\epsilon^{(k)}$ and $\lambda^{(k)}$ gives the distance and the projection of $\boldsymbol{\lambda}^{(\epsilon)}$ to the k^{th} hyperplane as defined in (11) and (12), respectively. Accordingly, let the queue-length process under MW Scheduling with this*

arrival process be denoted as $\{\mathbf{Q}^{(\epsilon)}[t]\}_t$, evolving as in (1). Moreover, let $\overline{\mathbf{Q}}^{(\epsilon)}$ denote a random vector with the same distribution as the steady-state distribution of $\{\mathbf{Q}^{(\epsilon)}[t]\}_t$ (by Lemma 3).

Then, for each $k \in \{1, \dots, K\}$, and with $\zeta^{(\epsilon,k)} \triangleq \langle (\mathbf{c}^{(k)})^2, (\boldsymbol{\sigma}^{(\epsilon)})^2 \rangle$,

$$(\epsilon^{(k)})^n \mathbb{E} \left[\langle \mathbf{c}^{(k)}, \overline{\mathbf{Q}}^{(\epsilon)} \rangle^n \right] \geq n! \left(\frac{\zeta^{(\epsilon,k)}}{2} \right)^n - B_n^{(\epsilon,k)}, \quad n \geq 1, \quad (62)$$

where $B_n^{(\epsilon,k)}$ vanishes as $\epsilon^{(k)} \downarrow 0$.

Further, for each $k \in \mathcal{K}_{\boldsymbol{\lambda}^{(\epsilon)}}$, consider the **heavy-traffic limit** $\epsilon^{(k)} \downarrow 0$; and suppose that the variance vector $(\boldsymbol{\sigma}^{(\epsilon)})^2$ approaches a constant vector $\boldsymbol{\sigma}^2$. Then, defining $\zeta^{(k)} \triangleq \langle (\mathbf{c}^{(k)})^2, \boldsymbol{\sigma}^2 \rangle$, we have

$$\liminf_{\epsilon^{(k)} \downarrow 0} (\epsilon^{(k)})^n \mathbb{E} \left[\langle \mathbf{c}^{(k)}, \overline{\mathbf{Q}}^{(\epsilon)} \rangle^n \right] \geq n! \left(\frac{\zeta^{(k)}}{2} \right)^n, \quad n \geq 1. \quad (63)$$

2. State-Space Collapse: The state space collapse result for MW Scheduling provided in Proposition 2 applies directly. Next, we utilize this result to develop upper bounds on the n^{th} moment of the steady-state queue-lengths.

3. Upper Bounds and n^{th} Moment Heavy-Traffic-Optimality of MW Scheduling: The main idea behind the analysis is to utilize the state-space collapse result, which implies that the total unused service under the MW Scheduler is small unless *all* queue-lengths are small. Thus, as the system gets heavily loaded, the total unused service vanishes, making the system act similarly to the lower bounding system investigated above. The following result builds on this to establish the n^{th} moment heavy-traffic optimality of MW Scheduler.

Proposition 5 Consider the scheduling problem under the exogenous arrival vector process $\{\mathbf{A}^{(\epsilon)}[t]\}_t$ with mean vector $\boldsymbol{\lambda}^{(\epsilon)} \in \text{Int}(\mathcal{R})$ and with variance vector $(\boldsymbol{\sigma}^{(\epsilon)})^2 \triangleq \left((\sigma_l^{(\epsilon)})^2 \right)_{l=1}^L$. The parameter $\epsilon = (\epsilon^{(k)})_{k=1}^K \succ 0$, where $\epsilon^{(k)}$ gives the distance and the projection of $\boldsymbol{\lambda}^{(\epsilon)}$ to the k^{th} hyperplane as defined in (11). Then, under the MW Scheduling, the limiting steady-state queue-vector $\overline{\mathbf{Q}}^{(\epsilon)}$ satisfies, for each $k \in \mathcal{K}_{\boldsymbol{\lambda}^{(\epsilon)}}$,

$$(\epsilon^{(k)})^n \mathbb{E} \left[\langle \mathbf{c}^{(k)}, \overline{\mathbf{Q}}^{(\epsilon)} \rangle \right] \leq n! \left(\frac{\zeta^{(\epsilon,k)}}{2} \right)^n + \overline{B}_n^{(\epsilon,k)}, \quad n \geq 1, \quad (64)$$

where we recall that $\zeta^{(\epsilon,k)} \triangleq \langle (\mathbf{c}^{(k)})^2, (\boldsymbol{\sigma}^{(\epsilon)})^2 \rangle$ as it is defined in Lemma 6, and $\overline{B}_n^{(\epsilon,k)}$ is vanishing as $\epsilon^{(k)} \downarrow 0$.

Also, in the **heavy traffic limit**, where we consider a sequence of exogenous arrival processes $\{\mathbf{A}^{(\epsilon,k)}[t]\}_t$ with their mean vector $\boldsymbol{\lambda}_{\mathbf{A}}^{(\epsilon)}$ approaching the k^{th} dominant face in $\mathcal{K}_{\boldsymbol{\lambda}^{(\epsilon)}}$ along its normal $\mathbf{c}^{(k)}$, and variance vectors $(\boldsymbol{\sigma}_{\mathbf{A}}^{(\epsilon)})^2$ approaching a constant vector $\boldsymbol{\sigma}^2$, we have

$$\limsup_{\epsilon^{(k)} \downarrow 0} (\epsilon^{(k)})^n \mathbb{E} \left[\langle \mathbf{c}^{(k)}, \overline{\mathbf{Q}}^{(\epsilon)} \rangle \right] \leq n! \left(\frac{\zeta^{(k)}}{2} \right)^n, \quad n \geq 1, \quad (65)$$

where $\zeta^{(k)} \triangleq \langle (\mathbf{c}^{(k)})^2, \boldsymbol{\sigma}^2 \rangle$.

Hence, comparing the heavy-traffic lower-bound (63) for any feasible policy to the heavy-traffic upper-bound (65) for MW Scheduler establishes the n^{th} -moment heavy-traffic optimality of MW Scheduling Policy.

Proof See Appendix D. \diamond

We note that the n^{th} moment argument concerns optimality of the norm of the projection $\|\overline{\mathbf{Q}}_{\parallel}^{(k)}\|$ onto the vector $\mathbf{c}^{(k)}$. It is interesting to relate this norm to the norm of the queue-length vector $\|\overline{\mathbf{Q}}\|$ when $n = 2$ to show the 2^{nd} moment optimality of MW Scheduling in minimizing $\lim_{\epsilon^{(k)} \downarrow 0} (\epsilon^{(k)})^2 \mathbb{E}[\|\overline{\mathbf{Q}}\|^2]$. This consequence is provided in the following corollary.

Corollary 1 *Under the same conditions as in Proposition 5, the MW Scheduler achieves, in the heavy-traffic limit,*

$$\limsup_{\epsilon^{(k)} \downarrow 0} (\epsilon^{(k)})^2 \mathbb{E}[\|\overline{\mathbf{Q}}\|^2] \leq \frac{(\zeta^{(k)})^2}{2}, \quad (66)$$

where $\zeta^{(k)} \triangleq \langle (\mathbf{c}^{(k)})^2, \boldsymbol{\sigma}^2 \rangle$. Furthermore, since the right-hand-side is the smallest heavy-traffic limit achievable by any policy, this establishes 2^{nd} -moment heavy-traffic optimality of MW Scheduling.

Proof The proof simply follows from noting that

$$\|\mathbf{Q}_{\parallel}^{(k)}\|^2 \leq \|\mathbf{Q}\|^2 = \|\mathbf{Q}_{\parallel}^{(k)}\|^2 + \|\mathbf{Q}_{\perp}^{(k)}\|^2,$$

from Pythagorean Theorem, and from combining the state-space collapse result $\mathbb{E}[\|\mathbf{Q}_{\perp}^{(k)}\|^2] \leq N_2^{(k)}$ where $N_2^{(k)}$ is independent of $\epsilon^{(k)}$ with the tight heavy-traffic lower and upper bounds (61) and (65) on $\mathbb{E}[\|\mathbf{Q}_{\parallel}^{(k)}\|^2]$. \diamond

6.2 Channel Fading

We consider the same setup as in Section 2.3, depicted in Figure 2, receiving nonnegative-integer-valued vector of arrivals $\{\mathbf{A}[t]\}_{t \geq 0}$ with $A_l[t] \leq A_{max}, \forall l, t$, distributed independently over links and also identically over time. However, instead of a *fixed* set of feasible rate vectors \mathcal{S} , we allow the feasible set to evolve randomly in time over a *finite* state space. In particular, we let $\{J[t]\}_{t \geq 0}$ be an i.i.d. sequence of random variables capturing the global state of the channel states of all links in the network. We assume that $J[t] \in \mathcal{J}$ for some set \mathcal{J} with finite cardinality, and let $\psi_j \triangleq \mathbb{P}(J[t] = j)$. Then, each global state $j \in \mathcal{J}$ yields a set of feasible rate vectors $\mathcal{S}^{(j)}$ that can be provided under that state. We assume that $\mathcal{S}^{(j)}$ has finite cardinality for each j with $S_l \leq S_{max}$, for all l and each $\mathbf{S} \triangleq (S_l)_l \in \mathcal{S}^{(j)}$. For simplicity, we also assume that each feasible rate vector in $\mathcal{S}^{(j)}$ is composed of non-negative integers, although the arguments hold for any discrete set of choices.

Capacity Region Under Channel Fading: With channel fading, the maximum achievable rate region becomes

$$\begin{aligned} \mathcal{R} &\triangleq \sum_{j \in \mathcal{J}} \psi_j \text{Convex Hull}(\mathcal{S}^{(j)}) \\ &= \text{Convex Hull} \left\{ \sum_{j \in \mathcal{J}} \psi_j \mathbf{s}^{(j)} : \mathbf{s}^{(j)} \in \mathcal{S}^{(j)}, \text{ for each } j \in \mathcal{J} \right\} \end{aligned}$$

which is henceforth called the *Fading Capacity Region*. Notice that the set $\text{Convex Hull}(\mathcal{S}^{(j)})$ for each $j \in \mathcal{J}$ simply yields a convex polyhedral set in \mathbb{R}_+^L as in (4). Thus, their finite weighted-sum also yields a polyhedral set that can be equivalently described, with a convenient abuse of notation, as:

$$\mathcal{R} = \{\mathbf{r} \geq \mathbf{0} : \langle \mathbf{c}^{(k)}, \mathbf{r} \rangle \leq b^{(k)}, k = 1, \dots, K\}, \quad (67)$$

where K denotes the finite (and minimal) number of hyperplanes that fully describe the polyhedron. We refer the reader to Definition 3 for the notions of the hyperplane $\mathcal{H}^{(k)}$, the pair $(\mathbf{c}^{(k)}, b^{(k)})$, and the face $\mathcal{F}^{(k)}$, which identically apply to \mathcal{R} .

Maximum Weight (MW) Scheduler Under Fading: In each slot, the purpose of the scheduler is to select a feasible rate vector from the feasible set to achieve stability of the queueing network in the long run. A well-known generalization (introduced in [38]) of the earlier MW Scheduler (cf. Definition 2) to this fading case is the following: given the queue-length vector $\mathbf{Q}[t]$ and the global channel state $J[t]$ at the beginning of slot t , the service rate vector is selected as

$$\mathbf{S}[t] \triangleq \mathbf{S}(\mathbf{Q}[t], J[t]) = \text{RAND} \left\{ \arg \max_{\mathbf{S} \in \mathcal{S}_{J[t]}} \langle \mathbf{Q}[t], \mathbf{S} \rangle \right\}. \quad (68)$$

Note that the rate vectors in \mathcal{R} are typically not instantaneously realizable, but only in the mean sense. In particular, *mean* service rate vector over the channel variations provided by the above MW Scheduler conditioned over a queue-length vector \mathbf{Q} satisfies:

$$\mathbf{R}(\mathbf{Q}) \triangleq \mathbb{E}[\mathbf{S}[t] | \mathbf{Q}[t] = \mathbf{Q}] = \text{RAND} \left\{ \arg \max_{\mathbf{R} \in \mathcal{R}} \langle \mathbf{Q}, \mathbf{R} \rangle \right\}. \quad (69)$$

We next comment on the application of the three steps of our methodology outlined in Section 1.1 to the steady-state performance of this MW Scheduler in the fading scenario.

1. Moment Bounds and Lower Bounds under Fading: The statement and proof of Lemma 3 apply once we replace the non-fading capacity region with the above fading capacity region. However, we note that the specific values of $\mathbf{c}^{(k)}$ and $b^{(k)}$ are different in the non-fading and fading capacity regions, as the latter also incorporates the channel fading distribution $\boldsymbol{\psi} \triangleq (\psi_j)_j$.

Next, we construct the lower-bounding system associated with the k^{th} face of \mathcal{R} , for any given $k \in \{1, \dots, K\}$, as in the non-fading case, except that the

service process is no longer at a constant rate but must incorporate the channel fading distribution ψ . Recalling that the pair $(\mathbf{c}^{(k)}, b^{(k)})$ describes the associated hyperplane, we first define

$$b^{(j,k)} \triangleq \max_{\mathbf{s} \in \mathcal{S}^{(j)}} \langle \mathbf{c}^{(k)}, \mathbf{s} \rangle, \quad \text{for each } j \in \mathcal{J},$$

which yields the maximum $\mathbf{c}^{(k)}$ -weighted service rate available in channel state j .

We are now ready to describe the governing arrival and service statistics of the lower bounding system (cf. Figure 5): the arrival process $\{\alpha^{(k)}[t]\}_{t \geq 0}$ of the lower bounding system associated with hyperplane $\mathcal{H}^{(k)}$ is set to $\alpha^{(k)}[t] = \langle \mathbf{c}^{(k)}, \mathbf{A}[t] \rangle$, while the service process $\{\beta^{(k)}[t]\}_{t \geq 0}$ is distributed as:

$$\mathbb{P}(\beta^{(k)}[t] = b^{(j,k)}) = \psi_j, \quad \text{for each } j \in \mathcal{J}.$$

We note that $\{b^{(j,k)}\}_{j \in \mathcal{J}}$ may be identical for different j , in which case their probabilities are aggregated. It is also true that $b^{(k)} = \mathbb{E}[\beta^{(k)}[t]]$ by construction of \mathcal{R} .

We note that the queue-length process $\{\Phi^{(k)}[t]\}_{t \geq 0}$ driven by the above arrival $\{\alpha^{(k)}[t]\}_{t \geq 0}$ and service $\{\beta^{(k)}[t]\}_{t \geq 0}$ processes as in (5) is stochastically smaller than $\{\langle \mathbf{c}^{(k)}, \mathbf{Q}[t] \rangle\}_{t \geq 0}$, where $\{\mathbf{Q}[t]\}_{t \geq 0}$ is the queue-length vector process under *any feasible* scheduling strategy. This follows from a coupling argument that utilizes the facts that: the service process of both the lower bounding and the actual queueing systems are governed by the same fading distribution ψ ; and that $b^{(j,k)}$, by definition, is the largest $\mathbf{c}^{(k)}$ -weighted service that can be provided to the queueing system under channel state j . Hence, any lower bound on $\{\Phi^{(k)}[t]\}_t$ is also a lower bound on $\{\langle \mathbf{c}^{(k)}, \mathbf{Q}[t] \rangle\}_t$.

The setup and definitions of Section 4.2 apply directly to the fading case given that \mathcal{R} is the fading capacity region defined above rather than the non-fading capacity region of Definition 3. Then, we have the following equivalent of Lemma 6.

Lemma 12 *For the above scheduling problem with a given set of channel fading distribution $\psi = (\psi_j)_j$ and their associated set of feasible schedules $\{\mathcal{S}^{(j)}\}_j$, consider the exogenous arrival vector process $\{\mathbf{A}^{(\epsilon)}[t]\}_t$ with mean vector $\boldsymbol{\lambda}^{(\epsilon)} \in \text{Int}(\mathcal{R})$ and with variance vector denoted as $(\boldsymbol{\sigma}^{(\epsilon)})^2 \triangleq ((\sigma_l^{(\epsilon)})^2)_{l=1}^L$. Here, the superscript vector is given by $\boldsymbol{\epsilon} = (\epsilon^{(k)})_{k=1}^K \succ 0$, where $\epsilon^{(k)}$ and $\lambda^{(k)}$ gives the distance and the projection of $\boldsymbol{\lambda}^{(\epsilon)}$ to the k^{th} hyperplane of (67) as defined in (11) and (12), respectively. Accordingly, let the queue-length process under MW Scheduling (see (68)) with this arrival process be denoted as $\{\mathbf{Q}^{(\epsilon)}[t]\}_t$, evolving as in (1). Moreover, let $\bar{\mathbf{Q}}^{(\epsilon)}$ denote a random random vector with the same distribution as the steady-state distribution of $\{\mathbf{Q}^{(\epsilon)}[t]\}_t$.*

Then, for each $k \in \{1, \dots, K\}$, we have

$$\mathbb{E} \left[\langle \mathbf{c}^{(k)}, \bar{\mathbf{Q}}^{(\epsilon)} \rangle \right] \geq \frac{\zeta^{(\epsilon,k)}}{2\epsilon^{(k)}} - B_1^{(k)} \quad (70)$$

where $\zeta^{(\epsilon,k)} \triangleq \mathbb{E} \left[\left(\langle \mathbf{c}^{(k)}, \mathbf{A}^{(\epsilon)} \rangle - \beta^{(k)} \right)^2 \right] = \langle (\mathbf{c}^{(k)})^2, (\boldsymbol{\sigma}^{(\epsilon)})^2 \rangle + \text{Var}(\beta^{(k)}) + (\epsilon^{(k)})^2$, and $B_1^{(k)} \triangleq \frac{S_{\max}}{2}$.

Further, for each $k \in \mathcal{K}_{\lambda(\epsilon)}$, consider the **heavy-traffic limit** $\epsilon^{(k)} \downarrow 0$; and suppose that the variance vector $(\sigma^{(\epsilon)})^2$ approaches a constant vector σ^2 . Then, defining $\zeta^{(k)} \triangleq \langle (\mathbf{c}^{(k)})^2, \sigma^2 \rangle + \text{Var}(\beta^{(k)})$, we have

$$\lim_{\epsilon^{(k)} \downarrow 0} \epsilon^{(k)} \mathbb{E} \left[\langle \mathbf{c}^{(k)}, \overline{\mathbf{Q}}^{(\epsilon)} \rangle \right] \geq \frac{\zeta^{(k)}}{2} \quad (71)$$

Comparing these lower bounds with their non-fading counterparts (14) and (15), we note that they include $\text{Var}(\beta^{(k)})$ which captures the impact of the fading distribution ψ on the steady-state mean queue-length levels.

2. State-Space Collapse of MW Scheduling under Fading: The state-space collapse argument under fading follows the same line of argument as in Section 4. First, the statement and proof of Lemma 7 applies without modification since the maximum service rate is uniformly bounded by S_{max} in every channel state. Next, we follow the development of Section 4.2 to consider the performance of a sequence of systems associated with a sequence of arrival processes $\{\mathbf{A}^{(\epsilon)}[t]\}_{t \geq 0}$ parameterized by $\epsilon = (\epsilon^{(k)})_k \succ \mathbf{0}$ as defined in (11).

Then, the corresponding queue-length vector process under MW Scheduling $\{\mathbf{Q}^{(\epsilon)}[t]\}_t$ converge in distribution to $\overline{\mathbf{Q}}^{(\epsilon)}$. Also, the definitions of $\mathbf{Q}_{\parallel}^{(\epsilon,k)}$, $\mathbf{Q}_{\perp}^{(\epsilon,k)}$, and $\mathcal{K}_{\lambda(\epsilon)}^{\circ}$ remain the same as in Section 4.2 with \mathcal{R} representing the fading capacity region (67). Then, the following state-space collapse result follows.

Proposition 6 *Assume $\lambda^{(\epsilon)} \in \text{Int}(\mathcal{R})$ in the scheduling problem with the corresponding set of $\{\epsilon^{(k)}\}_k$ as defined in (11). Then, under the MW Scheduling Policy, for each $k \in \mathcal{K}_{\lambda(\epsilon)}^{\circ}$, there exist finite constants $\{N_r^{(k)}\}_{r=1,2,\dots}$ such that $\mathbb{E} \left[\|\overline{\mathbf{Q}}_{\perp}^{(\epsilon,k)}\|^r \right] \leq N_r^{(k)}$, for all $\epsilon \succ \mathbf{0}$, each $k \in \mathcal{K}_{\lambda(\epsilon)}^{\circ}$, and each $r = 1, 2, \dots$.*

Proof (Outline) The proof of Proposition 2 directly applies to this statement with \mathcal{R} defined as in (67), and by replacing $\mathbb{E}[\mathbf{S} | \mathbf{Q}]$ with $\mathbf{R}(\mathbf{Q})$ as defined in (69). Most importantly, with these substitutions, the key property (30) continues to hold for the MW Scheduler under fading, which allows the rest of the argument to apply without modification. \diamond

3. Upper Bounds and Heavy-Traffic-Optimality of MW Scheduling under Fading: Similarly to the first two steps, the upper bound arguments of Section 5 extend to the fading scenario with minor modifications. First, Lemma 8 applies once we replace $\mathbf{S}(\overline{\mathbf{Q}})$ with $\mathbf{S}(\overline{\mathbf{Q}}, J)$ (see (68)) to capture the channel randomness, and let the expectation be over the channel fading distribution ψ as well. Then, the statement of Proposition 4 applies almost without modification to the fading case, which is repeated here for convenience.

Proposition 7 *Consider the scheduling problem under fading with capacity region (67) serving the exogenous arrival vector process $\{\mathbf{A}^{(\epsilon)}[t]\}_t$ with mean vector $\lambda^{(\epsilon)} \in \text{Int}(\mathcal{R})$ and with variance vector $(\sigma^{(\epsilon)})^2 \triangleq \left((\sigma_l^{(\epsilon)})^2 \right)_{l=1}^L$. The parameter $\epsilon = (\epsilon^{(k)})_{k=1}^K \succ \mathbf{0}$, where $\epsilon^{(k)}$ gives the distance and the projection of $\lambda^{(\epsilon)}$ to the k^{th} hyperplane of \mathcal{R}*

as defined in (11). Then, under MW Scheduling, the limiting steady-state queue-vector $\overline{\mathbf{Q}}^{(\epsilon)}$ satisfies, for each $k \in \mathcal{K}_{\lambda^{(\epsilon)}}^{\circ}$,

$$\mathbb{E} \left[\langle \mathbf{c}^{(k)}, \overline{\mathbf{Q}}^{(\epsilon)} \rangle \right] \leq \frac{\zeta^{(\epsilon,k)}}{2\epsilon^{(k)}} + \overline{B}_1^{(\epsilon,k)}, \quad (72)$$

where we recall that $\zeta^{(\epsilon,k)} \triangleq \langle (\mathbf{c}^{(k)})^2, (\boldsymbol{\sigma}^{(\epsilon)})^2 \rangle + \text{Var}(\beta^{(k)}) + (\epsilon^{(k)})^2$ as it is defined in Lemma 12, and $\overline{B}_1^{(\epsilon,k)}$ is $o\left(\frac{1}{\epsilon^{(k)}}\right)$, i.e., $\lim_{\epsilon^{(k)} \downarrow 0} \epsilon^{(k)} \overline{B}_1^{(\epsilon,k)} = 0$.

Also, in the **heavy traffic limit**, where we consider a sequence of exogenous arrival processes $\{\mathbf{A}^{(\epsilon,k)}[t]\}_t$ with their mean vector $\lambda_{\mathbf{A}}^{(\epsilon)}$ approaching the k^{th} dominant face in $\mathcal{K}_{\lambda^{(\epsilon)}}^{\circ}$ along its normal $\mathbf{c}^{(k)}$, and variance vectors $(\boldsymbol{\sigma}_{\mathbf{A}}^{(\epsilon)})^2$ approaching a constant vector $\boldsymbol{\sigma}^2$, we have

$$\lim_{\epsilon^{(k)} \downarrow 0} \epsilon^{(k)} \mathbb{E} \left[\langle \mathbf{c}^{(k)}, \overline{\mathbf{Q}}^{(\epsilon)} \rangle \right] \leq \frac{\zeta^{(k)}}{2}, \quad (73)$$

where $\zeta^{(k)} \triangleq \langle (\mathbf{c}^{(k)})^2, (\boldsymbol{\sigma}^{(\epsilon)})^2 \rangle + \text{Var}(\beta^{(k)})$.

Hence, comparing the heavy-traffic lower-bound (71) for any feasible policy to the heavy-traffic upper-bound (73) for MW Scheduler establishes the first moment heavy-traffic optimality of MW Scheduling Policy under fading.

Proof (Outline) We point to a few modifications in the proof of Proposition 4 that yields the proof of this statement. The need for these modifications arise from the fact that the MW scheduler does not directly select its allocation from the fading capacity region \mathcal{R} , but from the instantaneous feasible set of schedules $\{\mathcal{S}^{J[t]}\}_t$ available at the time. This subtlety can be handled partly by working with conditional expectation $\mathbf{R}(\mathbf{Q})$ defined in (69) instead of $\mathbf{S}(\mathbf{Q})$, and partly by introducing conditional probabilities in the analysis.

Omitting the details, we point to the similarities and differences from the non-fading case: the derivation of (46) follows with $\mathbf{S}(\overline{\mathbf{Q}})$ replaced by $\mathbf{R}(\overline{\mathbf{Q}})$; the definition of $\pi^{(k)}$ in (47) is modified to per channel state j as

$$\pi^{(j,k)} \triangleq \mathbb{P} \left(\langle \mathbf{c}^{(k)}, \mathbf{S}(\overline{\mathbf{Q}}, J) \rangle = b^{(j,k)} \mid J = j \right),$$

for each j, k ; the definition of $\gamma^{(k)}$ in (48) is unmodified; the statement of Claim 1 is true for the above conditional probability $\pi_j^{(k)}$, such that $(1 - \pi_j^{(k)}) = O(\epsilon^{(k)})$ for each channel state j . This establishes that under heavy-traffic conditions, the MW Scheduler will operate on the dominant face with high probability in every channel state. The rest of the argument follows the non-fading case since the previous fact allows us to approach the fading case as a time-average of non-fading cases, each operating on rate vectors contributing to the dominant face $\mathcal{F}^{(k)}$. \diamond

7 Conclusions

The main contribution of the paper is to show that drift conditions in steady-state can be used to obtain bounds on the moments of queue lengths that are tight in heavy traffic. The key new idea here is to derive an appropriate notion

of state-space collapse in steady-state which sharpen the bounds obtained using drift conditions. The results presented in this paper apply to the case where the state collapses to a single dimension. An interesting topic for further research is to understand whether the ideas presented here apply more generally.

Acknowledgements

We thank the anonymous reviewers for their useful suggestions, and Bin Li and Siva Theja Maguluri for their careful proof-reading of an earlier version of the paper.

Research supported in part by an AFOSR MURI FA 9550-10-10573, Army MURIs W911NF-07-1-0287 and W911NF-08-1-0233, AFOSR Grant FA-9550-08-1-0432, DTRA Grant HDTRA1-08-1-0016, and NSF Grants CAREER-CNS-0953515 and CCF-0916664.

References

1. S. Asmussen. *Applied Probability and Queues*. Springer-Verlag, New York, 2003.
2. S. L. Bell and R. J. Williams. Dynamic scheduling of a parallel server system in heavy traffic with complete resource pooling: asymptotic optimality of a threshold policy. *Electronic J. of Probability*, pages 1044–1115, 2005.
3. D. Bertsimas, D. Gamarnik, and J. N. Tsitsiklis. Performance of multiclass Markovian queueing networks via piecewise linear Lyapunov functions. *Annals of Applied Probability*, 2001.
4. D. Bertsimas, I. Ch. Paschalidis, and J. N. Tsitsiklis. Optimization of multiclass queueing processing networks. *Annals of Applied Probability*, 4:43–75, 1994.
5. M. Bramson. State space collapse with application to heavy-traffic limits for multiclass queueing networks. *Queueing Systems Theory and Applications*, pages 89–148, 1998.
6. J. G. Dai and W. Lin. Asymptotic optimality of maximum pressure policies in stochastic processing networks. *Annals of Applied Probability*, 2008.
7. A. Eryilmaz, R. Srikant, and J. Perkins. Stable scheduling policies for fading wireless channels. *IEEE/ACM Transactions on Networking*, 13(2):411–424, 2005.
8. G. J. Foschini and J. Salz. A basic dynamic routing problem and diffusion. *IEEE Transactions on Communications*, 26(3):320–327, March 1978.
9. D. Gamarnik and A. Zeevi. Validity of heavy-traffic steady-state approximations in generalized Jackson networks. *Ann. Appl. Probab.*, page 5690, 2006.
10. N. Gans and G. van Ryzin. Optimal control of a multiclass, flexible queueing system. *Operations Research*, 1997.
11. N. Gans and G. van Ryzin. Optimal dynamic scheduling of a general class of parallel-processing queueing systems. *Advances in Applied Probability*, 1998.
12. P. W. Glynn and A. Zeevi. Bounding stationary expectations of markov processes. In *Markov Processes and Related Topics: A Festschrift for Thomas G. Kurtz, IMS Collections*, 2008.
13. G. Gupta and N. B. Shroff. Delay analysis and optimality of scheduling policies for multi-hop wireless networks. *IEEE/ACM Trans. on Networking*, 2010.
14. B. Hajek. Hitting-time and occupation-time bounds implied by drift analysis with applications. *Ann. Appl. Probab.*, pages 502–525, 1982.
15. J. M. Harrison. Brownian models of queueing networks with heterogeneous customer populations. In *Stochastic Differential Systems, Stochastic Control Theory and Applications, IMA Vol. 10*, eds. W. Fleming and P.-L. Lions (Springer, New York), pages 147–186, 1988.
16. J. M. Harrison. Heavy traffic analysis of a system with parallel servers: Asymptotic optimality of discrete review policies. *Ann. Appl. Probab.*, pages 822–848, 1998.
17. J. M. Harrison and M. J. Lopez. Heavy traffic resource pooling in parallel-server systems. *Queueing Systems*, 33:339–368, 1999.

18. T. Ji, E. Athanasopoulou, and R. Srikant. On optimal scheduling algorithms for small generalized switches. *IEEE/ACM Transactions on Networking*, 2010.
19. 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. *Annals of Applied Probability*, 2009.
20. F. P. Kelly and C. N. Laws. Dynamic routing in open queueing networks: Brownian models, cut constraints and resource pooling. *Queueing Systems Theory and Applications*, pages 47–86, 1993.
21. J. F. C. Kingman. On queues in heavy traffic. *J. Roy. Statist. Soc., ser. B*, pages 383–392, 1962.
22. J. F. C. Kingman. Some inequalities for the queue GI/G/1. *Biometrika*, pages 315–324, 1962.
23. P. R. Kumar and S. P. Meyn. Stability of queueing networks and scheduling policies. *IEEE Transactions on Automatic Control*, 40:251–260, February 1995.
24. S. Kumar and P. R. Kumar. Performance bounds for queueing networks and scheduling policies. *IEEE Transactions on Automatic Control*, 39:1600–1611, 1994.
25. S. Kumar, R. Srikant, and P. R. Kumar. Bounding blocking probabilities and throughput in queueing networks with buffer constraints. *Queueing Systems Theory and Applications*, 28:55–77, 1998.
26. C. N. Laws. Resource pooling in queueing networks with dynamic routing. *Advances in Applied Probability*, pages 699–726, 1992.
27. R. Leelahakriengkrai and R. Agrawal. Scheduling in multimedia wireless networks. In *Proc. of the International Teletraffic Congress*, 2001.
28. S. P. Meyn. *Control techniques for complex networks*. Cambridge University Press, 2007.
29. S. P. Meyn. Stability and asymptotic optimality of generalized maxweight policies. *SIAM Journal on Control and Optimization*, 2007.
30. S. P. Meyn and R. L. Tweedie. *Markov Chains and Stochastic Stability*. Springer-Verlag, 1993.
31. M. Neely, E. Modiano, and C. Rohrs. Dynamic power allocation and routing for time varying wireless networks. *IEEE Selected Areas in Communications*, 2005.
32. M. I. Reiman. Some diffusion approximations with state space collapse. In *Proceedings of International Seminar on Modelling and Performance Evaluation Methodology, Lecture Notes in Control and Information Sciences*, pages 209–240. Springer, Berlin, 1983.
33. D. Shah and D. J. Wischik. The teleology of scheduling algorithms for switched networks under light load, critical load, and overload. <http://www.cs.ucl.ac.uk/staff/D.Wischik/Research/netsched.html>. Earlier version in *Proc. IEEE Infocom 2006* under the title “Optimal scheduling algorithms for input-queued switches”.
34. D. Shah and D. J. Wischik. Lower bound and optimality in switched networks. In *Proceedings of Allerton Conference on Communication, Control and Computation*, 2008.
35. S. Shakkottai, R. Srikant, and A. L. Stolyar. Pathwise optimality of the exponential scheduling rule for wireless channels. *Advances in Applied Probability*, pages 1021–1045, 2004.
36. A. L. Stolyar. Maxweight scheduling in a generalized switch: State space collapse and workload minimization in heavy traffic. *Annals of Applied Probability*, pages 1–53, 2004.
37. A. L. Stolyar and E. Yudovina. Systems with large flexible server pools: Instability of “natural” load balancing, 2011. Available at <http://arxiv.org/abs/1012.4140>.
38. 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*, pages 1936–1948, December 1992.
39. L. Tassiulas and A. Ephremides. Dynamic server allocation to parallel queues with randomly varying connectivity. *IEEE Transactions on Information Theory*, 39:466–478, March 1993.
40. L. Tassiulas and A. Ephremides. Dynamic scheduling for minimum delay in tandem and parallel constrained queueing models. *Annals of Operation Research*, 18:333–355, 1994.
41. V. J. Venkataramanan and X. Lin. On the queue-overflow probability of wireless systems: A new approach combining large deviations with Lyapunov functions, 2009. Preprint.
42. W. Whitt. Weak convergence theorems for priority queues: Preemptive resume discipline. *J. Appl. Probab.*, pages 74–94, 1971.
43. R. J. Williams. Diffusion approximations for open multiclass queueing networks: Sufficient conditions involving state space collapse. *Queueing Systems Theory and Applications*, 30:27–88, 1998.

A Proof of Lemma 2

The proof follows from the application of Lemma 1 to the Markov Chain $X[t] := \mathbf{Q}[t]$ using the Lyapunov function $Z(X) := V(\mathbf{Q}) \triangleq \|\mathbf{Q}\|$. First, we check that both conditions (C1) and (C2) are satisfied. We start with (C1):

$$\begin{aligned} \mathbb{E}[\Delta V(\mathbf{Q}) \mid \mathbf{Q}[t] = \mathbf{Q}] &= \mathbb{E}[\|\mathbf{Q}[t+1]\| - \|\mathbf{Q}[t]\| \mid \mathbf{Q}[t] = \mathbf{Q}] \\ &= \mathbb{E}\left[\sqrt{\|\mathbf{Q}[t+1]\|^2} - \sqrt{\|\mathbf{Q}[t]\|^2} \mid \mathbf{Q}[t] = \mathbf{Q}\right] \\ &\leq \frac{1}{2\|\mathbf{Q}\|} \mathbb{E}[\|\mathbf{Q}[t+1]\|^2 - \|\mathbf{Q}[t]\|^2 \mid \mathbf{Q}[t] = \mathbf{Q}], \end{aligned} \quad (74)$$

where the inequality follows from the fact that $f(x) = \sqrt{x}$ is concave for $x \geq 0$ so that $f(y) - f(x) \leq (y-x)f'(x) = \frac{(y-x)}{2\sqrt{x}}$ with $y := \|\mathbf{Q}[t+1]\|^2$ and $x := \|\mathbf{Q}[t]\|^2$. Next, we study the difference in (74), which is simply the mean drift of the quadratic Lyapunov function $W(\mathbf{Q}) \triangleq \|\mathbf{Q}\|^2$. We shall omit the time reference $[t]$ after the first step for brevity.

$$\begin{aligned} \mathbb{E}[\Delta W(\mathbf{Q}) \mid \mathbf{Q}] &= \mathbb{E}[\|\mathbf{Q}[t+1]\|^2 - \|\mathbf{Q}\|^2 \mid \mathbf{Q}] \\ &= \mathbb{E}[\|\mathbf{Q} + \mathbf{A} - \mathbf{S} + \mathbf{U}\|^2 - \|\mathbf{Q}\|^2 \mid \mathbf{Q}] \\ &= \mathbb{E}[\|\mathbf{Q} + \mathbf{A} - \mathbf{S}\|^2 + 2\langle \mathbf{Q} + \mathbf{A} - \mathbf{S}, \mathbf{U} \rangle + \|\mathbf{U}\|^2 - \|\mathbf{Q}\|^2 \mid \mathbf{Q}] \\ &\stackrel{(a)}{\leq} \mathbb{E}[\|\mathbf{Q} + \mathbf{A} - \mathbf{S}\|^2 - \|\mathbf{Q}\|^2 \mid \mathbf{Q}] \\ &= \mathbb{E}[2\langle \mathbf{Q}, \mathbf{A} - \mathbf{S} \rangle + \|\mathbf{A} - \mathbf{S}\|^2 \mid \mathbf{Q}] \\ &\leq 2\mathbb{E}[\langle \mathbf{Q}, \mathbf{A} - \mathbf{S} \rangle \mid \mathbf{Q}] + K_1, \end{aligned} \quad (75)$$

where the inequality (a) follows from the fact that $U_l(Q_l + A_l - S_l) = -U_l^2 \leq 0$, for each l , and $K_1 \triangleq L \max(A_{max}, S_{max})^2$ is finite since both the arrival and service processes are bounded.

Next, we bound the first term in (75) by first defining $\epsilon \triangleq \mu_\Sigma - \lambda_\Sigma$, and then defining a hypothetical arrival rate vector $\boldsymbol{\lambda} = (\lambda_l)_l$ with respect to the given service rate vector $\boldsymbol{\mu}$ and ϵ such that $\lambda_l \triangleq \mu_l - \frac{\epsilon}{L}$, for each l . Note that $\sum_{l=1}^L \lambda_l = \mu_\Sigma - \epsilon = \lambda_\Sigma$. Then, we can massage the first term in (75) as

$$\begin{aligned} \mathbb{E}[\langle \mathbf{Q}, \mathbf{A} - \mathbf{S} \rangle \mid \mathbf{Q}] &= \langle \mathbf{Q}, \mathbb{E}[\mathbf{A} \mid \mathbf{Q}] - \boldsymbol{\lambda} \rangle - \langle \mathbf{Q}, \boldsymbol{\mu} - \boldsymbol{\lambda} \rangle \\ &\stackrel{(a)}{=} \langle \mathbf{Q}, \mathbb{E}[\mathbf{A} \mid \mathbf{Q}] - \boldsymbol{\lambda} \rangle - \frac{\epsilon}{L} \langle \mathbf{Q}, \mathbf{1} \rangle \\ &\stackrel{(b)}{=} \mathbb{E}[A_\Sigma \mid \mathbf{Q}] Q_{min} - \langle \mathbf{Q}, \boldsymbol{\lambda} \rangle - \frac{\epsilon}{L} \|\mathbf{Q}\|_1 \quad \left\{ Q_{min} \triangleq \min_{1 \leq m \leq L} Q_m \geq 0 \right\} \\ &= \lambda_\Sigma Q_{min} - \sum_{l=1}^L \lambda_l Q_l - \frac{\epsilon}{L} \|\mathbf{Q}\| \\ &= -\sum_{l=1}^L \lambda_l (Q_l - Q_{min}) - \frac{\epsilon}{L} \|\mathbf{Q}\| \\ &\stackrel{(c)}{\leq} -\frac{\epsilon}{L} \|\mathbf{Q}\| \end{aligned}$$

where step (a) follows from the definition of $\boldsymbol{\lambda}$; (b) follows from the definitions of the JSQ policy (see Definition 1) and the l_1 norm $\|\mathbf{Q}\|_1 \triangleq \sum_{l=1}^L |Q_l|$; (c) is trivially true since $Q_{min} \leq Q_l$ for all l . Using this bound in (75) and back in (74) yields

$$\mathbb{E}[\Delta V(\mathbf{Q}) \mid \mathbf{Q}[t] = \mathbf{Q}] \leq -\frac{\epsilon}{L} + \frac{K_1}{2\|\mathbf{Q}\|},$$

which verifies Condition (C1). Moving on to Condition (C2), we have

$$\begin{aligned}
|\Delta V(\mathbf{Q})| &= \left| \|\mathbf{Q}[t+1]\| - \|\mathbf{Q}[t]\| \right| \mathcal{I}(\mathbf{Q}[t] = \mathbf{Q}) \\
&= \|\mathbf{Q}[t+1] - \mathbf{Q}[t]\| \mathcal{I}(\mathbf{Q}[t] = \mathbf{Q}) \\
&\stackrel{(a)}{\leq} \|\mathbf{Q}[t+1] - \mathbf{Q}[t]\|_1 \mathcal{I}(\mathbf{Q}[t] = \mathbf{Q}) \\
&\leq L \max_{1 \leq l \leq L} |Q_l[t+1] - Q_l[t]| \mathcal{I}(\mathbf{Q}[t] = \mathbf{Q}) \\
&\stackrel{(b)}{\leq} 2L \max(A_{max}, S_{max})
\end{aligned} \tag{76}$$

where (a) follows from the fact that $\|\mathbf{x}\|_1 \geq \|\mathbf{x}\|$ for any $\mathbf{x} \in \mathbb{R}^L$; and (b) is true since $A_l[t]$ and $S_l[t]$ are bounded by A_{max} and S_{max} , respectively, for all $l = 1, \dots, L$ and all $t \geq 0$. This verifies Condition (C2) and completes the proof.

B Proof of Lemma 7

We first prove (20):

$$\begin{aligned}
\Delta V_{\perp}(\mathbf{Q}) &= \left[\|\mathbf{Q}_{\perp}[t+1]\| - \|\mathbf{Q}_{\perp}[t]\| \right] \mathcal{I}(\mathbf{Q}[t] = \mathbf{Q}) \\
&= \left[\sqrt{\|\mathbf{Q}_{\perp}[t+1]\|^2} - \sqrt{\|\mathbf{Q}_{\perp}[t]\|^2} \right] \mathcal{I}(\mathbf{Q}[t] = \mathbf{Q}) \\
&\leq \frac{1}{2\|\mathbf{Q}_{\perp}[t]\|} \left[\|\mathbf{Q}_{\perp}[t+1]\|^2 - \|\mathbf{Q}_{\perp}[t]\|^2 \right] \mathcal{I}(\mathbf{Q}[t] = \mathbf{Q}) \\
&= \frac{1}{2\|\mathbf{Q}_{\perp}[t]\|} \left[\underbrace{\left(\|\mathbf{Q}[t+1]\|^2 - \|\mathbf{Q}[t]\|^2 \right)}_{=\Delta W(\mathbf{Q})} \mathcal{I}(\mathbf{Q}[t] = \mathbf{Q}) - \underbrace{\left(\|\mathbf{Q}_{\parallel}[t+1]\|^2 - \|\mathbf{Q}_{\parallel}[t]\|^2 \right)}_{=\Delta W_{\parallel}(\mathbf{Q})} \mathcal{I}(\mathbf{Q}[t] = \mathbf{Q}) \right],
\end{aligned}$$

where the inequality follows from the fact that $f(x) = \sqrt{x}$ is concave for $x \geq 0$ so that $f(y) - f(x) \leq (y-x)f'(x) = \frac{(y-x)}{2\sqrt{x}}$ with $y := \|\mathbf{Q}_{\perp}[t+1]\|^2$ and $x := \|\mathbf{Q}_{\perp}[t]\|^2$. Also, the last step follows from Pythagoras Theorem (3) with $\mathbf{x} := \mathbf{Q}_{\parallel}[\cdot]$ and $\mathbf{y} := \mathbf{Q}_{\perp}[\cdot]$.

Next, we prove (21):

$$\begin{aligned}
|\Delta V_{\perp}(\mathbf{Q})| &= \left| \|\mathbf{Q}_{\perp}[t+1]\| - \|\mathbf{Q}_{\perp}[t]\| \right| \mathcal{I}(\mathbf{Q}[t] = \mathbf{Q}) \\
&\stackrel{(a)}{\leq} \|\mathbf{Q}_{\perp}[t+1] - \mathbf{Q}_{\perp}[t]\| \mathcal{I}(\mathbf{Q}[t] = \mathbf{Q}) \\
&\stackrel{(b)}{=} \|\mathbf{Q}[t+1] - \mathbf{Q}[t] - \mathbf{Q}_{\parallel}[t] + \mathbf{Q}_{\parallel}[t+1]\| \mathcal{I}(\mathbf{Q}[t] = \mathbf{Q}) \\
&\stackrel{(c)}{\leq} (\|\mathbf{Q}[t+1] - \mathbf{Q}[t]\| + \|\mathbf{Q}_{\parallel}[t+1] - \mathbf{Q}_{\parallel}[t]\|) \mathcal{I}(\mathbf{Q}[t] = \mathbf{Q}) \\
&\stackrel{(d)}{\leq} 2\|\mathbf{Q}[t+1] - \mathbf{Q}[t]\| \mathcal{I}(\mathbf{Q}[t] = \mathbf{Q}) \\
&\stackrel{(e)}{\leq} 2\sqrt{L} \max_{1 \leq l \leq L} |Q_l[t+1] - Q_l[t]| \mathcal{I}(\mathbf{Q}[t] = \mathbf{Q}) \\
&\stackrel{(f)}{\leq} 2\sqrt{L} \max(A_{max}, S_{max})
\end{aligned}$$

where: (a) follows from the fact that $\|\|\mathbf{x}\| - \|\mathbf{y}\|\| \leq \|\mathbf{x} - \mathbf{y}\|$ for each $\mathbf{x}, \mathbf{y} \in \mathbb{R}^L$; (b) follows from the definition of $\mathbf{Q} = \mathbf{Q}_{\perp} + \mathbf{Q}_{\parallel}$; (c) follows from triangle inequality; (d) follows from the non-expansive nature of the projection onto a convex set once we note that $\mathbf{Q}_{\parallel}[\cdot]$ is the projection of $\mathbf{Q}[\cdot]$ onto the line along \mathbf{c} , which implies that $\|\mathbf{Q}_{\parallel}[t+1] - \mathbf{Q}_{\parallel}[t]\| \leq \|\mathbf{Q}[t+1] - \mathbf{Q}[t]\|$; (e) trivially follows from the definition of $\|\cdot\|$; and (f) is true since $A_l[t]$ and $S_l[t]$ are respectively assumed to be bounded by A_{max} and S_{max} for all $l = 1, \dots, L$ and all $t \geq 0$.

C Proof of Lemma 10

We have already argued the weak convergence and boundedness of all moments of the limiting random variable $\bar{\Phi}^{(\epsilon)}$ for each $\epsilon > 0$ in the proof of Lemma 4. Recalling the evolution (8), we study the mean drift of the Lyapunov function $W_n(\Phi) \triangleq \|\Phi\|^n$. In the following, we temporarily omit the time reference $[t]$ and the superscript (ϵ) for ease of exposition:

$$\begin{aligned}
\mathbb{E}[\Delta W_n(\Phi[t]) \mid \Phi[t] = \Phi] &= \mathbb{E}[(\Phi + \alpha - \beta + \chi)^n - \Phi^n \mid \Phi] \\
&= \mathbb{E}\left[(\Phi + \alpha - \beta)^n - \Phi^n + \sum_{i=0}^{n-1} \binom{n}{i} (\Phi + \alpha - \beta)^i \chi^{n-i} \mid \Phi\right] \\
&\stackrel{(a)}{=} \mathbb{E}\left[(\Phi + \alpha - \beta)^n - \Phi^n + \chi^n \sum_{i=0}^{n-1} \binom{n}{i} (-1)^i \mid \Phi\right] \\
&= \mathbb{E}[(\Phi + \alpha - \beta)^n - \Phi^n - (-\chi)^n \mid \Phi] \\
&= \mathbb{E}\left[\sum_{i=0}^{n-1} \binom{n}{i} \Phi^i (\alpha - \beta)^{n-i} - (-\chi)^n \mid \Phi\right] \\
&= \mathbb{E}\left[\sum_{i=0}^{n-2} \binom{n}{i} \Phi^i (\alpha - \beta)^{n-i} + n\Phi^{n-1}(\alpha - \beta) - (-\chi)^n \mid \Phi\right],
\end{aligned}$$

where (a) uses the fact that $\chi = -(\Phi + \alpha - \beta)\mathcal{I}(\Phi + \alpha - \beta < 0)$ by definition of χ . Taking expectations of both sides with respect to the steady-state distribution, i.e., setting $\Phi[t] = \bar{\Phi}$, and noting that $\mathbb{E}[\Delta W_n(\bar{\Phi})] = 0$ yields:

$$0 = \mathbb{E}\left[\sum_{i=0}^{n-2} \binom{n}{i} \bar{\Phi}^i (\alpha - \beta)^{n-i} + n\bar{\Phi}^{n-1}(\alpha - \beta) - (-\chi)^n\right]$$

Re-arranging terms, noting the independence of the arrival and service processes from each other and Φ , and recalling that $\mathbb{E}[\beta - \alpha] = \epsilon$ by construction allows us to write:

$$n\epsilon\mathbb{E}[\bar{\Phi}^{n-1}] = \sum_{i=0}^{n-2} \binom{n}{i} \mathbb{E}[\bar{\Phi}^i] \mathbb{E}[(\alpha - \beta)^{n-i}] - \mathbb{E}[(-\chi)^n]$$

We separate the final term in the summation and multiply both sides with ϵ^{n-2}/n to get:

$$\epsilon^{n-1}\mathbb{E}[\bar{\Phi}^{n-1}] = \frac{(n-1)}{2}\epsilon^{(n-2)}\mathbb{E}[\bar{\Phi}^{n-2}]\mathbb{E}[(\alpha - \beta)^2] \quad (77)$$

$$+ \sum_{i=0}^{n-3} \binom{n}{i} \frac{\epsilon^{(n-2)}\mathbb{E}[\bar{\Phi}^i]}{n} \mathbb{E}[(\alpha - \beta)^{n-i}] - \frac{\epsilon^{(n-2)}\mathbb{E}[(-\chi)^n]}{n}. \quad (78)$$

Next, we investigate terms in (78) to show that they vanish with $\epsilon \downarrow 0$. To that end, we first note that

$$\left|\frac{\epsilon^{(n-2)}\mathbb{E}[(-\chi)^n]}{n}\right| \leq \left(\frac{\epsilon^{(n-2)}S_{max}^{n-1}}{n}\right)\mathbb{E}[\chi] = \left(\frac{\epsilon^{(n-1)}S_{max}^{n-1}}{n}\right), \quad (79)$$

where the final equality follows from the fact that $\mathbb{E}[\chi] = \epsilon$ under steady-state operation. Clearly, the final expression vanishes with $\epsilon \downarrow 0$ for any $n \geq 2$.

Next, we turn the summation in (78) to argue inductively that $\lim_{\epsilon \downarrow 0} \epsilon^{n-2}\mathbb{E}[\bar{\Phi}^i] = 0$ for all $n \geq 3$ and for each $i = 0, 1, \dots, n-3$. This result, when proven, implies that the whole sum is vanishing with ϵ . The first step of the induction holds trivially since $\mathbb{E}[\bar{\Phi}^0] = \epsilon$. Suppose the claim is true for some $n \geq 3$ and all $i = 0, 1, \dots, n-3$. We would like to confirm it for

$(n+1)$ and $i = 0, 1, \dots, n-2$ as well. This is straight-forward for all $i = 0, 1, \dots, n-1$. The case when $i = n-2$ requires us to investigate $\epsilon^{n-1}\mathbb{E}[\overline{\Phi}^{n-2}]$ using the expansion (77)-(78):

$$\epsilon^{n-1}\mathbb{E}[\overline{\Phi}^{n-2}] = \sum_{i=0}^{n-2} \binom{n}{i} \left(\frac{\epsilon^{n-3}\mathbb{E}[\overline{\Phi}^i]}{n} \right) \mathbb{E}[(\alpha - \beta)^{n-i}] - \frac{\epsilon^{n-2}\mathbb{E}[(-\chi)^n]}{n}.$$

The right-hand-side of this expression vanishes with ϵ since $\epsilon^{n-2}\mathbb{E}[\overline{\Phi}^i]$ vanishes for each $i = 0, 1, \dots, n-3$ by the induction hypothesis, and $\epsilon^{n-2}\mathbb{E}[(-\chi)^n]$ vanishes according to (79). We note that the moments of $(\alpha - \beta)$ are bounded since the arrival and service processes are both bounded. This completes the induction proof, and hence establish that $\lim_{\epsilon \downarrow 0} \epsilon^{n-2}\mathbb{E}[\overline{\Phi}^i] = 0$ for all $n \geq 3$ and for each $i = 0, 1, \dots, n-3$.

Returning to (77)-(78) and revoking the (ϵ) notation to highlight the dependence on ϵ , we have thus proven that there exist $\{C_n^{(\epsilon)}\}_{n \geq 1}$ that vanish with ϵ such that

$$\begin{aligned} \epsilon^n \mathbb{E}[(\overline{\Phi}^{(\epsilon)})^n] &\geq \frac{n}{2} \epsilon^{(n-1)} \mathbb{E}[(\overline{\Phi}^{(\epsilon)})^{n-1}] \mathbb{E}[(\alpha^{(\epsilon)} - \beta)^2] - C_n^{(\epsilon)} \\ &= n \left(\frac{\zeta^{(\epsilon)}}{2} \right) \epsilon^{(n-1)} \mathbb{E}[(\overline{\Phi}^{(\epsilon)})^{n-1}] - C_n^{(\epsilon)} \\ &\geq n! \left(\frac{\zeta^{(\epsilon)}}{2} \right)^n - B_n^{(\epsilon)}, \end{aligned} \tag{80}$$

where $\zeta^{(\epsilon)} \triangleq (\sigma_{\Sigma}^{(\epsilon)})^2 + \nu_{\Sigma}^2 + \epsilon^2$, and $B_n^{(\epsilon)} \triangleq \sum_{k=1}^n \frac{n!}{(n-k)!} \left(\frac{\zeta^{(\epsilon)}}{2} \right)^k C_{n-k}^{(\epsilon)}$ also vanishes with ϵ , which proves (60). The heavy-traffic result (61) then follows immediately by taking the limit of both sides as $\epsilon \downarrow 0$.

D Proof of Proposition 5

We temporarily omit the superscript (ϵ) associated with the arrival and queue-length processes for ease of exposition, and consider any $k \in \mathcal{K}_{\lambda}^{\circ}$. We, then, study mean drift of the Lyapunov

function $W_n^{(k)}(\mathbf{Q}) \triangleq \langle \mathbf{c}^{(k)}, \mathbf{Q} \rangle^n$ associate with the n^{th} moment.

$$\begin{aligned}
\mathbb{E}[\Delta W_n^{(k)}(\mathbf{Q}[t]) \mid \mathbf{Q}[t] = \mathbf{Q}] &= \mathbb{E} \left[\langle \mathbf{c}^{(k)}, \mathbf{Q} + \mathbf{A} - \mathbf{S} + \mathbf{U} \rangle^n - \langle \mathbf{c}^{(k)}, \mathbf{Q} \rangle^n \mid \mathbf{Q} \right] \\
&= \mathbb{E} \left[\langle \mathbf{c}^{(k)}, \mathbf{Q} + \mathbf{A} - \mathbf{S} \rangle^n - \langle \mathbf{c}^{(k)}, \mathbf{Q} \rangle^n \mid \mathbf{Q} \right] \\
&\quad + \sum_{i=0}^{n-1} \binom{n}{i} \mathbb{E} \left[\langle \mathbf{c}^{(k)}, \mathbf{Q} + \mathbf{A} - \mathbf{S} \rangle^i \langle \mathbf{c}^{(k)}, \mathbf{U} \rangle^{n-i} \mid \mathbf{Q} \right] \\
&= \mathbb{E} \left[(\langle \mathbf{c}^{(k)}, \mathbf{Q} + \mathbf{A} \rangle - b^{(k)} + b^{(k)} - \langle \mathbf{c}^{(k)}, \mathbf{S} \rangle)^n - \langle \mathbf{c}^{(k)}, \mathbf{Q} \rangle^n \mid \mathbf{Q} \right] \\
&\quad + \sum_{i=0}^{n-1} \binom{n}{i} \mathbb{E} \left[\langle \mathbf{c}^{(k)}, \mathbf{Q} + \mathbf{A} - \mathbf{S} \rangle^i \langle \mathbf{c}^{(k)}, \mathbf{U} \rangle^{n-i} \mid \mathbf{Q} \right] \\
&= \mathbb{E} \left[(\langle \mathbf{c}^{(k)}, \mathbf{Q} + \mathbf{A} \rangle - b^{(k)})^n - \langle \mathbf{c}^{(k)}, \mathbf{Q} \rangle^n \mid \mathbf{Q} \right] \\
&\quad + \sum_{i=0}^{n-1} \binom{n}{i} \mathbb{E} \left[(\langle \mathbf{c}^{(k)}, \mathbf{Q} + \mathbf{A} \rangle - b^{(k)})^i (b^{(k)} - \langle \mathbf{c}^{(k)}, \mathbf{S} \rangle)^{n-i} \mid \mathbf{Q} \right] \\
&\quad + \sum_{i=0}^{n-1} \binom{n}{i} \mathbb{E} \left[\langle \mathbf{c}^{(k)}, \mathbf{Q} + \mathbf{A} - \mathbf{S} \rangle^i \langle \mathbf{c}^{(k)}, \mathbf{U} \rangle^{n-i} \mid \mathbf{Q} \right] \\
&= n \langle \mathbf{c}^{(k)}, \mathbf{Q} \rangle^{n-1} \mathbb{E}[\langle \mathbf{c}^{(k)}, \mathbf{A} \rangle - b^{(k)} \mid \mathbf{Q}] \tag{81} \\
&\quad + \sum_{i=0}^{n-2} \binom{n}{i} \mathbb{E} \left[\langle \mathbf{c}^{(k)}, \mathbf{Q} \rangle^i (\langle \mathbf{c}^{(k)}, \mathbf{A} \rangle - b^{(k)})^{n-i} \mid \mathbf{Q} \right] \\
&\quad + \sum_{i=0}^{n-1} \binom{n}{i} \mathbb{E} \left[(\langle \mathbf{c}^{(k)}, \mathbf{Q} + \mathbf{A} \rangle - b^{(k)})^i (b^{(k)} - \langle \mathbf{c}^{(k)}, \mathbf{S} \rangle)^{n-i} \mid \mathbf{Q} \right] \\
&\quad + \sum_{i=0}^{n-1} \binom{n}{i} \mathbb{E} \left[\langle \mathbf{c}^{(k)}, \mathbf{Q} + \mathbf{A} - \mathbf{S} \rangle^i \langle \mathbf{c}^{(k)}, \mathbf{U} \rangle^{n-i} \mid \mathbf{Q} \right]
\end{aligned}$$

Note in (81) that $\mathbb{E}[\langle \mathbf{c}^{(k)}, \mathbf{A} \rangle - b^{(k)} \mid \mathbf{Q}] = -\epsilon^{(k)}$ by construction. Thus, taking the expectation of both sides over the steady-state distribution, i.e. setting $\mathbf{Q} = \overline{\mathbf{Q}}$, noting $\mathbb{E}[\Delta W_n^{(k)}(\overline{\mathbf{Q}})] = 0$, and multiplying both sides with $(\epsilon^{(k)})^{n-2}$ yields:

$$n(\epsilon^{(k)})^{n-1} \mathbb{E}[\langle \mathbf{c}^{(k)}, \overline{\mathbf{Q}} \rangle^{n-1}] = \sum_{i=0}^{n-2} \binom{n}{i} (\epsilon^{(k)})^{n-2} \mathbb{E} \left[\langle \mathbf{c}^{(k)}, \overline{\mathbf{Q}} \rangle^i (\langle \mathbf{c}^{(k)}, \mathbf{A} \rangle - b^{(k)})^{n-i} \right] \tag{82}$$

$$+ \sum_{i=0}^{n-1} \binom{n}{i} (\epsilon^{(k)})^{n-2} \mathbb{E} \left[(\langle \mathbf{c}^{(k)}, \overline{\mathbf{Q}} + \mathbf{A} \rangle - b^{(k)})^i (b^{(k)} - \langle \mathbf{c}^{(k)}, \mathbf{S}(\overline{\mathbf{Q}}) \rangle)^{n-i} \right] \tag{83}$$

$$+ \sum_{i=0}^{n-1} \binom{n}{i} (\epsilon^{(k)})^{n-2} \mathbb{E} \left[\langle \mathbf{c}^{(k)}, \overline{\mathbf{Q}} + \mathbf{A} - \mathbf{S}(\overline{\mathbf{Q}}) \rangle^i \langle \mathbf{c}^{(k)}, \mathbf{U}(\overline{\mathbf{Q}}) \rangle^{n-i} \right] \tag{84}$$

We note the resemblance of this expression to (77)-(78) of the lower-bounding system with $\overline{\Phi}$ replaced with $\langle \mathbf{c}^{(k)}, \overline{\mathbf{Q}} \rangle$. Accordingly, the claimed upper bound (64) follows from induction as in the argument (80) once we establish that (83) and (84) both vanish as $\epsilon^{(k)} \downarrow 0$. Thus, we next study the behavior of each of these expressions with $\epsilon^{(k)}$ for all $n \geq 2$.

We first study the expectation in (83) to show that it is of order $\sqrt{\epsilon^{(k)}}$ and hence (83) must vanish as $\epsilon^{(k)} \downarrow 0$ for any $n \geq 2$. For any $i = 0, 1, \dots, n-1$, we have

$$\begin{aligned}
& \mathbb{E} \left[\langle \mathbf{c}^{(k)}, \bar{\mathbf{Q}} + \mathbf{A} \rangle - b^{(k)} \rangle^i (b^{(k)} - \langle \mathbf{c}^{(k)}, \mathbf{S}(\bar{\mathbf{Q}}) \rangle)^{n-i} \right] \\
&= \sum_{j=0}^i \binom{i}{j} \mathbb{E} \left[\langle \mathbf{c}^{(k)}, \bar{\mathbf{Q}} \rangle^j \langle \mathbf{c}^{(k)}, \mathbf{A} \rangle - b^{(k)} \rangle^{j-i} (b^{(k)} - \langle \mathbf{c}^{(k)}, \mathbf{S}(\bar{\mathbf{Q}}) \rangle)^{n-i} \right] \\
&= \sum_{j=0}^i \binom{i}{j} \mathbb{E} \left[\|\bar{\mathbf{Q}}_{\parallel}^{(k)}\|^j (b^{(k)} - \langle \mathbf{c}^{(k)}, \mathbf{S}(\bar{\mathbf{Q}}) \rangle)^{n-i} \right] \mathbb{E} \left[\langle \mathbf{c}^{(k)}, \mathbf{A} \rangle - b^{(k)} \rangle^{j-i} \right] \quad (85)
\end{aligned}$$

where in the last step we used the independence of arrival processes, and used the definition of $\bar{\mathbf{Q}}_{\parallel}^{(k)}$. Following the same argument as in the derivation of (53) to bound the first expectation as:

$$\begin{aligned}
\mathbb{E} \left[\|\bar{\mathbf{Q}}_{\parallel}^{(k)}\|^j (b^{(k)} - \langle \mathbf{c}^{(k)}, \mathbf{S}(\bar{\mathbf{Q}}) \rangle)^{n-i} \right] &\leq (\cot(\theta^{(k)}))^j \mathbb{E} \left[\|\bar{\mathbf{Q}}_{\perp}^{(k)}\|^j (b^{(k)} - \langle \mathbf{c}^{(k)}, \mathbf{S}(\bar{\mathbf{Q}}) \rangle)^{n-i} \right] \\
&\leq (\cot(\theta^{(k)}))^j \sqrt{\mathbb{E} \left[\|\bar{\mathbf{Q}}_{\perp}^{(k)}\|^{2j} \right] \mathbb{E} \left[(b^{(k)} - \langle \mathbf{c}^{(k)}, \mathbf{S}(\bar{\mathbf{Q}}) \rangle)^{2(n-i)} \right]}
\end{aligned}$$

We know from Proposition 2 that $\mathbb{E} \left[\|\bar{\mathbf{Q}}_{\perp}^{(k)}\|^{2j} \right] \leq N_{2j}^{(k)}$ for some finite $N_{2j}^{(k)}$ independent of $\epsilon^{(k)}$. Also, since $i \in \{0, \dots, n-1\}$, we can show that

$$\mathbb{E} \left[(b^{(k)} - \langle \mathbf{c}^{(k)}, \mathbf{S}(\bar{\mathbf{Q}}) \rangle)^{2(n-i)} \right] \leq \frac{\epsilon^{(k)}}{\gamma^{(k)}} \left((b^{(k)})^{2(n-i)} + \langle \mathbf{c}^{(k)}, S_{max} \mathbf{1} \rangle^{2(n-i)} \right) = O(\epsilon^{(k)})$$

exactly as argued in (51). Thus, these two bounds together with the fact that $A_l \leq A_{max}$ for all l establishes that (85) = $O\left(\sqrt{\epsilon^{(k)}}\right)$, which, in turn, proves that (83) = $O\left((\epsilon^{(k)})^{n-\frac{3}{2}}\right)$, i.e., (83) vanishes as $\epsilon^{(k)} \downarrow 0$ for all $n \geq 2$.

Next, we study the expectations in (84) to show that they also vanish as $\epsilon^{(k)} \downarrow 0$ for all $n \geq 2$. We recall the $\tilde{\mathbf{c}}, \tilde{\mathbf{Q}}, \tilde{\mathbf{U}}, \tilde{\mathbf{Q}}^+$, and $\mathbb{E}_{\tilde{\mathbf{Q}}[\cdot]}$ notation introduced for the derivations (56)-(59) and follow the same line of reasoning. In particular, we note that, for each $i = 0, 1, \dots, n-1$, the expectation in (84) can be bounded as

$$\begin{aligned}
\mathbb{E}_{\tilde{\mathbf{Q}}} \left[\langle \mathbf{c}^{(k)}, \mathbf{Q} + \mathbf{A} - \mathbf{S} \rangle^i \langle \mathbf{c}^{(k)}, \mathbf{U} \rangle^{n-i} \right] &\leq \mathbb{E}_{\tilde{\mathbf{Q}}} \left[\langle \mathbf{c}^{(k)}, \mathbf{Q}^+ \rangle^i \langle \mathbf{c}^{(k)}, \mathbf{U} \rangle^{n-i} \right] \\
&= \mathbb{E}_{\tilde{\mathbf{Q}}} \left[\langle \tilde{\mathbf{c}}^{(k)}, \tilde{\mathbf{Q}}^+ \rangle^i \langle \tilde{\mathbf{c}}^{(k)}, \tilde{\mathbf{U}} \rangle^{n-i} \right] \\
&= \mathbb{E}_{\tilde{\mathbf{Q}}} \left[\|\tilde{\mathbf{Q}}_{\parallel}^+\|^i \|\tilde{\mathbf{U}}_{\parallel}\|^{n-i} \right] \quad (86)
\end{aligned}$$

To bound (86), we consider two cases separately: $i \leq n/2$, and $i > n/2$. When $i \leq n/2$, the argument closely follows (56)-(59) to show that (86) = $O\left(\sqrt{\epsilon^{(k)}}\right)$. When $i > n/2$, we multiply and divide (assuming the non-trivial case of $\|\tilde{\mathbf{U}}_{\parallel}\| \neq 0$) by $\|\tilde{\mathbf{U}}_{\parallel}\|^{2i-n}$ to get

$$\begin{aligned}
\mathbb{E}_{\tilde{\mathbf{Q}}} \left[\|\tilde{\mathbf{Q}}_{\parallel}^+\|^i \|\tilde{\mathbf{U}}_{\parallel}\|^{n-i} \right] &= \mathbb{E}_{\tilde{\mathbf{Q}}} \left[\frac{\|\tilde{\mathbf{Q}}_{\parallel}^+\|^i \|\tilde{\mathbf{U}}_{\parallel}\|^i}{\|\tilde{\mathbf{U}}_{\parallel}\|^{2i-n}} \right] \\
&\stackrel{(a)}{=} \mathbb{E}_{\tilde{\mathbf{Q}}} \left[\frac{\langle \tilde{\mathbf{Q}}_{\perp}^+, \tilde{\mathbf{U}} \rangle^i}{\|\tilde{\mathbf{U}}_{\parallel}\|^{2i-n}} \right] \\
&= \mathbb{E}_{\tilde{\mathbf{Q}}} \left[\left\langle \tilde{\mathbf{Q}}_{\perp}^+, \frac{\tilde{\mathbf{U}}}{\|\tilde{\mathbf{U}}_{\parallel}\|^{2-\frac{n}{i}}} \right\rangle^i \right] \\
&\stackrel{(b)}{\leq} \sqrt{\mathbb{E}_{\tilde{\mathbf{Q}}} \left[\|\tilde{\mathbf{Q}}_{\perp}^+\|^{2i} \right] \mathbb{E}_{\tilde{\mathbf{Q}}} \left[\|\tilde{\mathbf{U}}_{\parallel}\|^{2(n-i)} \right]}
\end{aligned}$$

where (a) follows from Lemma 9 similarly as in the derivation of (56); and (b) follows from Cauchy-Schwartz inequality after minor modifications. It is then easy to show, as in (57) and

(58), that $\mathbb{E}_{\mathbf{Q}} \left[\|\tilde{\mathbf{Q}}_{\perp}^+\|^{2i} \right] \leq N_{2i}^{(k)}$ with $N_{2i}^{(k)}$ defined in Proposition 2 and $\mathbb{E}_{\mathbf{Q}} \left[\|\tilde{\mathbf{U}}\|^{2(n-i)} \right] = O(\epsilon^{(k)})$ since $i \leq n - 1$. Hence, we have shown that (86) = $O\left(\sqrt{\epsilon^{(k)}}\right)$ in this case as well. Substituting this result back in (84) proves the claimed result that (84) = $O\left(\left(\epsilon^{(k)}\right)^{n-\frac{3}{2}}\right)$, i.e., (84) vanishes as $\epsilon^{(k)} \downarrow 0$ for all $n \geq 2$.

As we noted before, the fact that (83) + (84) vanishes with $\epsilon^{(k)}$ is sufficient (exactly as in the lower-bound argument of (80)) to inductively derive the upper-bound (64) via the recursive relationship in (82). Consequently, the heavy-traffic result (65) follows immediately, completing the proof.