Queueing Systems
https://doi.org/10.1007/s11134-021-09695-x

Heavy traffic queue length scaling in switches with reconfiguration delay

Chang-Heng Wang[1] · **Siva Theja Maguluri**[2] · **Tara Javidi**[1]

Received: 9 October 2019 / Revised: 30 January 2021 / Accepted: 9 February 2021

Abstract

The Adaptive MaxWeight policy achieves optimal throughput for switches with nonzero reconfiguration delay and has been shown to have good delay performance in simulation. In this paper, we analyze the queue length behavior of a switch with nonzero reconfiguration delay operating under the Adaptive MaxWeight. We first show that the Adaptive MaxWeight policy exhibits a weak state space collapse behavior in steady state, which can be viewed as an inheritance of a similar property of the MaxWeight policy in a switch with zero reconfiguration delay. The weak state space collapse result is then utilized to obtain an asymptotically tight bound on an expression involving the steady-state queue length and the probability of reconfiguration for the Adaptive MaxWeight policy in the heavy traffic regime. We then derive the relation between the expected schedule duration and the steady-state queue length through Lyapunov drift analysis and characterize bounds for the expected steady-state queue length. While the resulting queue length bounds are not asymptotically tight, they suggest an approximate queue length scaling behavior, which approaches the optimal scaling with respect to the traffic load and the reconfiguration delay when the hysteresis function of the Adaptive MaxWeight policy approaches a linear function.

Keywords Switch scheduling · Reconfiguration delay · MaxWeight · Heavy traffic · State space collapse

Mathematics Subject Classification 60K25 · 90B15

✉ Chang-Heng Wang
chw009@ucsd.edu

Siva Theja Maguluri
siva.theja@gatech.edu

Tara Javidi
tjavidi@ucsd.edu

1 University of California, San Diego, La Jolla, CA, USA

2 Georgia Institute of Technology, Atlanta, GA, USA

Published online: 05 March 2021

1 Introduction

Modern data centers aggregate huge amount of computing and storage resource to support high demand applications such as cloud computing, large-scale web applications, and big data analytics. With the ever increasing number of resources and the communication demand between these resources, the interconnecting networks face stringent performance challenges. Optical switches emerge as a promising candidate to address this challenge since they can support higher data bandwidth relatively easier than traditional electronic switches, and also have lower power consumption. However, optical switches pose another challenge different from traditional electronic switches that makes it difficult to directly substitute electronic switches: optical switches typically exhibit a delay following each schedule reconfiguration, during which no packet transmission can occur [6,9]. This delay is referred as the **reconfiguration delay**, and it makes the switch scheduling problem more difficult. For example, it is known that with the reconfiguration delay, the well-known MaxWeight policy [16,17] is not even throughput optimal.

For the scheduling of switches with reconfiguration delay, many works in the literature consider the problem in a "quasi-static" sense: decomposing the traffic demand into a sequence of efficient schedules (in the sense of minimizing service time and number of schedule reconfigurations) for a predetermined time horizon, for example, [10,12,13,18]. The performance of such solutions is usually limited by the duration of the time horizon and may require some prior knowledge of the traffic arrival rate to achieve good performance. In contrast, [4,5,20] consider the dynamic scheduling problem for switches with reconfiguration delay. The Adaptive MaxWeight policy [20] (or a similar variant, Switching Curve-Based policy [4]) makes a schedule decision at every time slot, as opposed to scheduling over a time horizon. The key idea of Adaptive MaxWeight is to reconfigure the schedule only when the current schedule is not good enough, which will be described with more detail in Sect. 3. The idea has also been generalized in [19] to introduce a large class of scheduling policies for switches with reconfiguration delay. These policies have been shown to guarantee throughput optimality under mild assumptions on arrival traffic, which means that the policies can ensure finite expected queue lengths whenever there exists any policy that provides such guarantee.

In this work, we focus on the delay analysis of the Adaptive MaxWeight policy. Beyond throughput optimality, it is desirable to further consider the delay behavior (or equivalently, queue length behavior) in order to better evaluate the performance. However, similar to the MaxWeight policy in switches without reconfiguration delay, an exact expression for the steady-state queue length behavior remains an open problem. Therefore, one usually approaches the delay performance analysis by studying the queue length scaling with respect to either the number of queues in the system or the traffic load. In this paper, we consider the queue length scaling with the traffic load for switches with reconfiguration delay (operated under Adaptive MaxWeight) in the heavy traffic regime. In particular, the arrival rates considered in this paper approach the boundary of the capacity region, with a limiting arrival traffic where all input ports and all output ports are saturated. The contributions of this paper include the following:

(i) We prove that in the considered heavy traffic regime, the steady-state queue length under Adaptive MaxWeight exhibits a behavior similar to the state space collapse (SSC) as introduced in [15], which is referred as weak state space collapse (WSSC) here.

(ii) We derive an upper bound on the expected sum of queue lengths in the heavy traffic regime. The derivation utilizes a drift analysis technique introduced in [7], and the Lyapunov function proposed in [15]. Combined with the WSSC result, the analysis provides an asymptotically tight bound on an expression involving the steady-state queue length and the probability of reconfiguration for the Adaptive MaxWeight policy in the heavy traffic regime.

(iii) We derive the relation between the expected schedule duration with the expected steady-state queue length and combine this relation with the aforementioned asymptotically tight bound to derive bounds on the steady-state queue length. While the resulting bounds are not asymptotically tight, they suggest an approximate scaling behavior of the steady-state queue length in the heavy traffic regime.

(iv) We derive a universal lower bound on the expected sum of the queue lengths for switches with reconfiguration under any scheduling policy. While the well-known lower bound for switches without reconfiguration delay trivially applies, it does not provide an insight on how reconfiguration delay limits the performance of the scheduling policies. The lower bound derived in this work refines the previous one and identifies the effect of the reconfiguration delay. Comparing with this lower bound, we show that Adaptive MaxWeight approaches the optimal scaling with respect to the traffic load and to the reconfiguration delay.

The rest of the paper is organized as follows: The switch with reconfiguration delay model and the notion of throughput optimality is introduced in Sect. 2. In Sect. 3, we briefly introduce the Adaptive MaxWeight policy and some of its properties. We then present our main result regarding heavy traffic queue length behavior in Sect. 4. We first establish the WSSC property of the Adaptive MaxWeight policy. With the WSSC result, we then establish a queue length upper bound for the steady-state queue length in the heavy traffic regime, which is dependent on the expected schedule duration. Using the schedule weight as a Lyapunov function, we further derive the relation between the expected schedule duration and the expected steady-state queue length through drift analysis, and then derive the scaling of steady-state queue length in heavy traffic. In Sect. 5, we derive some benchmark performance for AdaptiveMaxWeight to compare with, including a universal queue length lower bound for any scheduling policy as well as a queue length upper bound of a benchmark policy known as the Fixed Frame MaxWeight. Section 6 presents some simulation results for AdaptiveMaxWeight, in an effort to characterize the scaling of the expected queue length with respect to some system parameters, and in comparison with the scaling derived in this paper. Finally, we conclude with a summary and some future directions in Sect. 7.

2 System model

2.1 Switch model and arrival traffic

The model considered in this paper is an $n \times n$ input-queued switch, which has n input ports and n output ports. Each input port maintains n separate queues (either physically or virtually), each storing packets destined to one output port. We denote the queue storing packets at input i and destined for output j with the pair (i, j). This model is also known as the input-queued switch.

The system considered is assumed to be time-slotted, with the time indexed as $t \in \mathbb{N}_+ = \{0, 1, 2, \ldots\}$. Each slot duration is the transmission time of a single packet, which is assumed to be a fixed value. Let $a_{ij}(t)$ be the number of packets arriving at queue (i, j) at time t. Let $q_{ij}(t)$ be the number of packets in the queue (i, j) at the beginning of the time slot t. Write $\mathbf{a}(t) = [a_{ij}(t)]$, $\mathbf{q}(t) = [q_{ij}(t)]$, where $\mathbf{a}(t), \mathbf{q}(t) \in \mathbb{N}_+^{n\times n}$.

We assume the arrival processes $a_{ij}(t)$ to be independent from each other, where $i, j \in \{1, 2, \ldots, N\}, i \neq j$. For each queue (i, j), the arrival process $a_{ij}(t)$ is assumed to be i.i.d. across time slots, with mean $\mathbb{E}[a_{ij}(t)] = \lambda_{ij}$ and variance $\mathrm{Var}(a_{ij}(t)) = \sigma_{ij}^2$. We also assume that $a_{ij}(t)$ has a finite support, i.e., $\exists\, a_{\max} < \infty$ such that $a_{ij}(t) \leq a_{\max}$.

2.2 Schedules and reconfiguration delay

Let $\mathbf{s}(t) \in \{0, 1\}^{N\times N}$ denote the schedule at time slot t, which indicates the queues that are being scheduled by the switch. We set $s_{ij}(t) = 1$ if queue (i, j) is scheduled at time t, and $s_{ij}(t) = 0$ otherwise.

The feasible schedules for the network are determined by the network topology and physical constraints on simultaneous data transmissions. We let $\mathcal{S} \subset \mathbb{R}^{n\times n}$ denote the set of all feasible schedules, i.e., $\mathbf{s}(t) \in \mathcal{S}$ for all t. We assume at any t each input port can only transmit to at most one output port, and each output port can only receive from at most one input port, i.e., $\sum_i s_{ij}(t) \leq 1$, $\sum_j s_{ij}(t) \leq 1$. This schedule constraint determines the set of feasible schedules $\mathcal{S}$. Among the feasible schedules, for schedules that satisfy $\sum_i s_{ij}(t) = 1$, $\sum_j s_{ij}(t) = 1$ for all i, j, we referred to these schedules as maximal schedules.

Upon reconfiguring a schedule, the network incurs a reconfiguration delay, during which no packet can be transmitted. We make this notion formal through the following two definitions:

Definition 1 Let $\{t_k^S\}_{k=0}^{\infty}$ denote the time instances when the schedule is reconfigured. The schedule between two schedule reconfiguration time instances remains the same, i.e.,

$$\mathbf{s}(\tau) = \mathbf{s}(t_k^S), \quad \forall \tau \in [t_k^S, t_{k+1}^S - 1].$$

Definition 2 Let Δ_r be the reconfiguration delay associated with reconfiguring the schedule of the network. During the period of schedule reconfiguration, i.e., $\forall t \in$

$\cup_{k=0}^{\infty}[t_k^S, t_k^S + \Delta_r]$, the switch does not serve any of its queues. We assume the reconfiguration delay to be an integer multiple of a time slot.

Let $r(t)$ denote the time remaining in the reconfiguration delay, with $r(t) = 0$ indicating that the switch is not in reconfiguration at time t. Therefore, if $t \in [t_k^S, t_k^S + \Delta_r]$ for some $k \in \mathbb{N}_+$, we have $r(t) = \Delta_r - (t - t_k^S)$; and $r(t) = 0$ for all other t.

With the above definitions, we may then write the queue dynamics for any queue (i, j) as

$$\begin{aligned} q_{ij}(t+1) &= \left[q_{ij}(t) + a_{ij}(t) - s_{ij}(t)\mathbb{1}_{\{r(t)=0\}}\right]^+ \\ &= q_{ij}(t) + a_{ij}(t) - s_{ij}(t)\mathbb{1}_{\{r(t)=0\}} + u_{ij}(t), \end{aligned} \tag{1}$$

where $\mathbb{1}_E$ is the indicator function of the event E, and $[x]^+ = \max\{x, 0\}$. Note that $u_{ij}(t) \in \{0, 1\}$ is the unused service of queue (i, j) when the queue is empty, which is defined as

$$u_{ij}(t) := \begin{cases} 1, & \text{if } q_{ij}(t) + a_{ij}(t) - s_{ij}(t)\mathbb{1}_{\{r(t)=0\}} = -1, \\ 0, & \text{otherwise.} \end{cases}$$

A useful property of the unused service is that $u_{ij}(t) = 1$ only when $q_{ij}(t+1) = 0$, or in other words, $u_{ij}(t)q_{ij}(t+1) = 0$.

The schedule at each time slot $\mathbf{s}(t)$ is determined by a scheduling policy. In this paper, we consider scheduling policies that determine the schedule at time t based on the queue length $\mathbf{q}(t)$ and the previous schedule $\mathbf{s}(t-1)$. Under this type of policy, the process $\{\mathbf{X}(t)\}_{t=0}^{\infty}$ with $\mathbf{X}(t) = (\mathbf{q}(t), \mathbf{s}(t), r(t)) \in \mathbb{N}_+^{n\times n} \times \{0, 1\}^{n\times n} \times \{0, 1, \ldots, \Delta_r\} \triangleq \mathcal{X}$ that describes the switch model is then a discrete time Markov chain.

2.3 Stability and capacity region

A queue (i, j) is strongly stable if its queue length $q_{ij}(t)$ satisfies

$$\limsup_{t\to\infty} \frac{1}{t}\sum_{\tau=1}^{t} \mathbb{E}[q_{ij}(\tau)] < \infty,$$

and we say the system of queues is stable if queue (i, j) is strongly stable for all $i, j \in \{1, 2, \ldots, N\}$. A scheduling policy is said to stabilize the system if the system is stable under that scheduling policy for a given traffic rate matrix. With this notion of stability, we define the capacity region $\mathcal{C}$ of the network as the set of all traffic rate matrices such that there exists a scheduling policy which stabilizes the system.

The capacity region is given by the convex hull of the feasible schedules $\mathcal{S}$ [17], that is,

$$\mathcal{C} = \left\{\sum_{\mathbf{s}\in\mathcal{S}} \alpha_{\mathbf{s}}\mathbf{s} : \sum_{\mathbf{s}\in\mathcal{S}} \alpha_{\mathbf{s}} < 1,\ \alpha_{\mathbf{s}} \geq 0,\ \forall \mathbf{s} \in \mathcal{S}\right\}.$$

We also define an outer boundary of $\mathcal{C}$ where all input ports and output ports are saturated, as

$$\mathcal{F} = \left\{ \boldsymbol{\lambda} \in \mathbb{R}^{n \times n} : \sum_i \lambda_{ij} = 1, \sum_j \lambda_{ij} = 1, \forall i, j \in \{1, 2, \ldots, n\} \right\}.$$

For any traffic rate matrix $\boldsymbol{\lambda} \in \mathcal{C}$, we say that $\boldsymbol{\lambda}$ is admissible, and define the load of the traffic as $\rho(\boldsymbol{\lambda}) = \max\{r : \boldsymbol{\lambda} \in r\mathcal{C},\ 0 < r < 1\}$. We say that a scheduling policy is throughput optimal if it stabilizes the system for any traffic rate matrix $\boldsymbol{\lambda} \in \mathcal{C}$.

3 Adaptive MaxWeight policy

It is known that for switches without reconfiguration delay, the MaxWeight policy is throughput optimal [16] and has optimal delay scaling in the heavy traffic regime [14, 15]. However, with the presence of reconfiguration delay, the MaxWeight policy is not even throughput optimal since it does not account for the overhead of frequent schedule reconfiguration [3].

The Adaptive MaxWeight scheduling policy is presented in Algorithm 1. The main idea behind Adaptive MaxWeight is to reconfigure the schedule when the current schedule is not "good" enough. Using the schedule weight as the measure of a schedule, Adaptive MaxWeight computes the schedule weight difference between the current schedule and the MaxWeight schedule, W^* (which is the "best" schedule under this measure) and compares this weight difference to a threshold which is a function of the maximum weight, $g(W^*)$. When the schedule weight difference exceeds the threshold, we reconfigure the schedule to the MaxWeight schedule, otherwise we keep the current schedule.

The selection of the threshold determines the performance of the policy. In [20], it has been shown that if $g(x) = (1-\gamma)x^{1-\delta}$, then Adaptive MaxWeight is throughput optimal. This result has been generalized in [19] to any strictly increasing and sublinear function $g(x)$, where a sublinear function $g(x)$ is defined as any function that satisfies $\lim_{x\to\infty} \frac{g(x)}{x} = 0$. The strictly increasing and sublinear function $g(x)$ is referred to as the hysteresis function, and the throughput optimality result is stated as the following fact.

Fact 1 *Given any reconfiguration delay $\Delta_r > 0$, and given any sublinear and strictly increasing hysteresis function g, the Markov chain $\mathbf{X}(t)$ is positive recurrent for any admissible traffic rate matrix under Adaptive MaxWeight. Therefore, Adaptive MaxWeight is throughput optimal.*

While the throughput optimality is a desirable property, it may be considered as only a first-order performance metric, in the sense that it only guarantees bounded expected queue length (and thus bounded expected delay), but the queue length could still be very large. One more step forward is to characterize its expected queue length, which is the main theme in the rest of this paper.

Algorithm 1 Adaptive MaxWeight Scheduling Policy

Require: Sublinear and strictly increasing function $g(\cdot)$
 for each $t = 0, 1, \ldots$ **do**
  $\mathbf{s}^*(t) \leftarrow \arg\max_{\mathbf{s}\in\mathcal{S}} \sum_{ij} q_{ij}(t)s_{ij}$
  $W^*(t) \leftarrow \max_{\mathbf{s}\in\mathcal{S}} \sum_{ij} q_{ij}(t)s_{ij}$
  $W(t) \leftarrow \sum_{ij} q_{ij}(t)s_{ij}(t-1)$
  $\Delta W(t) \leftarrow W^*(t) - W(t)$
  if $\Delta W(t) > g(W^*(t))$ **then**
   $\mathbf{s}(t) \leftarrow \mathbf{s}^*(t)$
  else
   $\mathbf{s}(t) \leftarrow \mathbf{s}(t-1)$
  end if
 end for

4 Heavy traffic analysis of Adaptive MaxWeight

Studying queue length or delay performance for a queueing system such as a switch in general is challenging. Therefore, analyses of such systems are mostly considered within certain asymptotic regimes. In this paper, we focus on the heavy traffic regime and make use of a drift technique developed in [7]. The outline of the heavy traffic analysis for the Adaptive MaxWeight is sketched as follows: In Sect. 4.1, we first introduce and establish a weak state space collapse property for Adaptive MaxWeight. The weak state space collapse property states that the queue length matrix is "almost" concentrated in a cone $\mathcal{K}$ (to be defined later) in the heavy traffic regime. The cone $\mathcal{K}$ has the geometric property that for any queue length matrix in the cone $\mathcal{K}$, all maximal schedules have the same weight for this queue length matrix, and this geometric property plays a key role in the following analysis toward a queue length upper bound. In Sect. 4.2, we then apply the drift analysis to Adaptive MaxWeight, which utilizes a Lyapunov function proposed in [15] and was developed to utilize the geometric property of the cone $\mathcal{K}$. The result of the drift analysis is a steady-state queue length upper bound that is dependent on the expected schedule duration. Lastly, in Sect. 4.3, we characterize the relation between the expected schedule duration and the queue length and use this relation to derive asymptotic bounds on the expected queue length at reconfiguration times.

In this paper, we are interested in the queue length behavior of switches with reconfiguration delay in the heavy traffic regime. In particular, we consider a sequence of switch systems indexed by ϵ, where each switch system has i.i.d. arrival traffic $\mathbf{a}^{(\epsilon)}(t)$ with mean and variance given by

$$\mathbb{E}[\mathbf{a}^{(\epsilon)}(t)] = \boldsymbol{\lambda}^{(\epsilon)} = \boldsymbol{\nu}(1-\epsilon), \quad \mathrm{Var}[\mathbf{a}^{(\epsilon)}(t)] = \left(\boldsymbol{\sigma}^{(\epsilon)}\right)^2,$$

where $\boldsymbol{\nu} \in \mathcal{F}$ and $\left(\boldsymbol{\sigma}^{(\epsilon)}\right)^2 \to \boldsymbol{\sigma}^2$ as $\epsilon \to 0$. The traffic load of each switch is $\rho = 1-\epsilon$. Recall that $\mathcal{F}$ is the set of critically loaded rate matrix with all ports saturated. The sequence of switches considered here have arrival rate matrices that approach $\boldsymbol{\nu}$ as we take $\epsilon \to 0$.

4.1 Weak state space collapse

It was shown in [15] that for a switch with no reconfiguration delay, the MaxWeight scheduling exhibits a multi-dimensional state space collapse. To be specific, let $\mathbf{e}^{(i)}$ denote the matrix with the ith row being all ones and zeros everywhere else, and $\tilde{\mathbf{e}}^{(j)}$ denote the matrix with the jth column being all ones and zeros everywhere else. As $\epsilon \to 0$, the steady-state queue length $\bar{\mathbf{q}}^{(\epsilon)}$ "concentrates" in the cone spanned by the matrices $\{\mathbf{e}^{(i)}\}_{i=1}^{n} \cup \{\tilde{\mathbf{e}}^{(j)}\}_{j=1}^{n}$, i.e.,

$$\mathcal{K} = \left\{\mathbf{x} \in \mathbb{R}^{n\times n} : \mathbf{x} = \sum_{i} w_i \mathbf{e}^{(i)} + \sum_{j} \tilde{w}_j \tilde{\mathbf{e}}^{(j)}, \text{ where } w_i, \tilde{w}_j \in \mathbb{R}_+ \text{ for all } i, j\right\},$$

in the sense that the projection of $\bar{\mathbf{q}}^{(\epsilon)}$ onto $\mathcal{K}$ is the dominant component in $\bar{\mathbf{q}}^{(\epsilon)}$. More specifically, for any $\mathbf{x} \in \mathbb{R}^{n\times n}$, define the projection of $\mathbf{x}$ on to $\mathcal{K}$ as

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

where $\|\cdot\|$ denotes the l_2-norm, and define $\mathbf{x}_{\perp} = \mathbf{x} - \mathbf{x}_{\|}$. In the heavy traffic limit ($\epsilon \to 0$), all moments of $\bar{\mathbf{q}}_{\perp}^{(\epsilon)}$ are bounded by a constant, and hence this is a negligible component in $\bar{\mathbf{q}}^{(\epsilon)}$ since it can be shown that $\|\bar{\mathbf{q}}^{(\epsilon)}\|$ is $\Omega(1/\epsilon)$. This is referred as state space collapse (SSC) in [15].

The cone $\mathcal{K}$ has the property that for any given queue length matrix $\mathbf{q} \in \mathcal{K}$, all maximal schedules have the same weight $\sum_{ij} q_{ij} s_{ij}$ [15]. In other words, the SSC property implies that the weights of all maximal schedules are equalized.

In this paper, we consider a weaker notion of the SSC property for switches with reconfiguration delay operated under the Adaptive MaxWeight policy.

Definition 3 (*Weak state space collapse*) Given a sequence of switch systems $\mathbf{X}^{(\epsilon)}(t) = (\mathbf{q}^{(\epsilon)}(t), \mathbf{s}^{(\epsilon)}(t), r^{(\epsilon)}(t))$, parametrized by $0 < \epsilon < 1$, suppose each switch system is positive recurrent and converges in distribution to a steady-state random vector $\bar{\mathbf{X}}^{(\epsilon)} = (\bar{\mathbf{q}}^{(\epsilon)}, \bar{\mathbf{s}}^{(\epsilon)}, \bar{r}^{(\epsilon)})$. We say that the sequence of switch systems exhibit a weak state space collapse if

$$\lim_{\epsilon\to 0} \frac{\mathbb{E}\left[\|\bar{\mathbf{q}}_{\perp}^{(\epsilon)}\|\right]}{\mathbb{E}\left[\|\bar{\mathbf{q}}^{(\epsilon)}\|\right]} = 0.$$

In contrast to the SSC in [15], where the moments of $\|\bar{\mathbf{q}}_{\perp}\|$ are bounded by finite constants, the weak state space collapse (WSSC) only requires the ratio $\|\bar{\mathbf{q}}_{\perp}\|/\|\bar{\mathbf{q}}\|$ converging to 0 as $\epsilon \to 0$. This could be considered as a multiplicative type of SSC, which has the similar flavor to the multiplicative state space collapse in the diffusion approximation literature [2,8]. It is worth noting that a recent work [21] also considers the multiplicative type of SSC in the context of bandwidth sharing using a different approach, and the growth rate of the moment bound with respect to ϵ could also be identified.

In the rest of this section, we use Lemma 1 to derive the WSSC for switches with reconfiguration delay operated under the Adaptive MaxWeight policy. Lemma 1 is a T-step version of [1, Theorem 1] where T could be any fixed integer. The proof of this lemma can be completed by simply replacing the transition probability by a T-step transition probability in [1, Theorem 1], hence we omit the proof here.

Lemma 1 *Consider an irreducible and aperiodic Markov Chain $\{X(t)\}_{t\geq 0}$ over a countable state space $\mathcal{X}$, and suppose $Z : \mathcal{X} \to \mathbb{R}_+$ is a nonnegative Lyapunov function. For any fixed integer $T > 0$, we define the T-step drift $\Delta^{\mathrm{T}} Z(X)$ of Z at state X as*

$$\Delta^{\mathrm{T}} Z(X) = [Z(X(t+T)) - Z(X(t))] \, \mathbb{1}_{\{X(t)=X\}}.$$

Suppose that, for some $T > 0$, the T-step drift satisfies the following conditions C.1 and C.2:

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

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

C.2 *There exists a $D < \infty$ such that, for all $X \in \mathcal{X}$,*

$$\Pr\left\{ |\Delta^{\mathrm{T}} Z(X)| \leq D \right\} = 1.$$

If the Markov chain $\{X(t)\}_{t\geq 0}$ converges in distribution to a random variable $\bar{X}$, then

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

With Lemma 1, we are now able to show the following proposition, which is essential to establish the WSSC result for the Adaptive MaxWeight policy.

Proposition 1 *Consider a set of switch systems with a fixed reconfiguration delay $\Delta_r > 0$, parametrized by $0 < \epsilon < 1$, all operated under the Adaptive MaxWeight policy with hysteresis function $g(\cdot)$, where $g(\cdot)$ is sublinear and strictly increasing. Each system has arrival process $\mathbf{a}^{(\epsilon)}(t)$ as described in Sect. 2. The mean arrival rate vector $\boldsymbol{\lambda}^{(\epsilon)} = (1-\epsilon)\boldsymbol{\nu}$ for some fixed $\boldsymbol{\nu} \in \mathcal{F}$ is such that $\nu_{\min} \triangleq \min_{ij} \nu_{ij} > 0$. Let the variance $\left(\boldsymbol{\sigma}^{(\epsilon)}\right)^2$ of the arrival process satisfy that $\|\boldsymbol{\sigma}^{(\epsilon)}\|^2 \leq \tilde{\sigma}^2$ for some $\tilde{\sigma}^2$ not dependent on ϵ.*

Let $\mathbf{X}^{(\epsilon)}(t) \in \mathcal{X}$ denote the process that determines each system, which is positive recurrent and hence converges to a steady-state random vector in distribution, denoted as $\bar{\mathbf{X}}^{(\epsilon)} = (\bar{\mathbf{q}}^{(\epsilon)}, \bar{\mathbf{s}}^{(\epsilon)}, \bar{r}^{(\epsilon)})$. Then, for any fixed θ with $0 < \theta < 1/2$, and for each system with $0 < \epsilon \leq \nu_{\min}/4\|\boldsymbol{\nu}\|$, the steady-state queue length satisfies

$$\mathbb{E}\Big[\|\bar{\mathbf{q}}_{\perp}^{(\epsilon)}\| - \theta \|\bar{\mathbf{q}}_{\|}^{(\epsilon)}\| \Big] \leq M_\theta, \tag{2}$$

where M_θ is a function of θ, $\tilde{\sigma}$, $a_{\max}$, $\nu_{\min}$ and n, but is independent of ϵ.

The proof of Proposition 1 is given in Appendix A. Comparing Proposition 1 with [15, Proposition 2], we may see that we no longer have the guarantee that all moments of $\|\bar{\mathbf{q}}_\perp^{(\epsilon)}\|$ are bounded here. However, we can still show that $\mathbb{E}[\|\bar{\mathbf{q}}_\perp^{(\epsilon)}\|]$ is negligible compared to $\mathbb{E}[\|\bar{\mathbf{q}}^{(\epsilon)}\|]$ as $\epsilon \to 0$, hence we consider this as a weak version of SSC. In particular, notice that the constant M_θ is independent of ϵ, and that $\mathbb{E}[\|\bar{\mathbf{q}}^{(\epsilon)}\|] \to \infty$ as $\epsilon \to 0$. Then, since $\mathbb{E}[\|\bar{\mathbf{q}}_\perp^{(\epsilon)}\|] \leq \theta\mathbb{E}[\|\bar{\mathbf{q}}_\parallel^{(\epsilon)}\|] + M_\theta \leq \theta\mathbb{E}[\|\bar{\mathbf{q}}^{(\epsilon)}\|] + M_\theta$ for any $\epsilon > 0$, we have $\lim_{\epsilon\to 0} \frac{\mathbb{E}\left[\|\bar{\mathbf{q}}_\perp^{(\epsilon)}\|\right]}{\mathbb{E}\left[\|\bar{\mathbf{q}}^{(\epsilon)}\|\right]} \leq \theta$ for any $\theta > 0$. Therefore, we may conclude that

$$\lim_{\epsilon\to 0} \frac{\mathbb{E}\left[\|\bar{\mathbf{q}}_\perp^{(\epsilon)}\|\right]}{\mathbb{E}\left[\|\bar{\mathbf{q}}^{(\epsilon)}\|\right]} = 0. \tag{3}$$

4.2 Drift analysis

With the WSSC result from the previous subsection, we now utilize Lyapunov drift analysis similar to [15] to derive an asymptotically tight bound on an expression involving the steady-state queue length $\mathbb{E}[\sum_{ij} \bar{\mathbf{q}}_{ij}]$ and the probability of reconfiguration $\Pr\{\bar{r}^{(\epsilon)} > 0\}$, as shown in Theorem 1.

Theorem 1 *Consider a set of switch systems with a fixed reconfiguration delay $\Delta_r > 0$, parametrized by $0 < \epsilon < 1$, all operated under the Adaptive MaxWeight policy with hysteresis function $g(\cdot)$, where $g(\cdot)$ is a sublinear and strictly increasing function. Each system has arrival process $\mathbf{a}^{(\epsilon)}(t)$ as described in Sect. 2. The mean arrival rate vector $\boldsymbol{\lambda}^{(\epsilon)} = (1-\epsilon)\boldsymbol{\nu}$ for some fixed $\boldsymbol{\nu} \in \mathcal{F}$ is such that $\nu_{\min} \triangleq \min_{ij} \nu_{ij} > 0$, and the variance of the arrival process is denoted as $\left(\boldsymbol{\sigma}^{(\epsilon)}\right)^2$. Then, for each system with $0 < \epsilon \leq \nu_{\min}/4\|\boldsymbol{\nu}\|$, and any fixed number θ satisfying $0 < \theta < \frac{1}{2}$, the steady-state distribution $\bar{\mathbf{X}}^{(\epsilon)} = (\bar{\mathbf{q}}^{(\epsilon)}, \bar{\mathbf{s}}^{(\epsilon)}, \bar{r}^{(\epsilon)})$ of the Markov chain $\mathbf{X}^{(\epsilon)} = (\mathbf{q}^{(\epsilon)}, \mathbf{s}^{(\epsilon)}, r^{(\epsilon)})$ satisfies*

$$\left(\epsilon - \Pr\{\bar{r}^{(\epsilon)} > 0\}\right)\left(\mathbb{E}\left[\sum_{ij} \bar{q}_{ij}^{(\epsilon)}\right]\right) \leq \frac{1-\frac{1}{2n}}{1-2n^3\theta}\|\boldsymbol{\sigma}^{(\epsilon)}\|^2 + B_1(\theta, \epsilon, n) \tag{4}$$

and

$$\left(\epsilon - \Pr\{\bar{r}^{(\epsilon)} > 0\}\right)\left(\mathbb{E}\left[\sum_{ij} \bar{q}_{ij}^{(\epsilon)}\right]\right) \geq \frac{1-\frac{1}{2n}}{1+3n^3\theta}\|\boldsymbol{\sigma}^{(\epsilon)}\|^2 + B_2(\theta, \epsilon, n), \tag{5}$$

where $\lim_{\epsilon\downarrow 0} B_1(\theta,\epsilon,n) = 0$ *and* $\lim_{\epsilon\downarrow 0} B_2(\theta,\epsilon,n) = 0$. *Since we may take* θ *arbitrarily close to* 0*, then in the heavy traffic limit as* $\epsilon \downarrow 0$*, if* $\left(\boldsymbol{\sigma}^{(\epsilon)}\right)^2 \to \tilde{\boldsymbol{\sigma}}^2$*, we have*

$$\lim_{\epsilon\downarrow 0}\left(\epsilon - \Pr\{\bar{r}^{(\epsilon)} > 0\}\right)\left(\mathbb{E}\Big[\sum_{ij}\bar{q}_{ij}^{(\epsilon)}\Big]\right) = \left(1-\frac{1}{2n}\right)\|\tilde{\boldsymbol{\sigma}}\|^2. \tag{6}$$

The asymptotically tight bounds in Theorem 1 take a form similar to the expected queue length bound of the MaxWeight policy in switches without reconfiguration delay [15, Theorem 1], except that the probability of reconfiguration $\Pr\{\bar{r}^{(\epsilon)} > 0\}$ is deducted from ϵ here.

To gain a clearer insight into the impact of the reconfiguration delay on the expected queue length, we take a further look at the upper bound (4). Let $\beta_1 = \frac{1-\frac{1}{2n}}{1-2n^3\theta}\|\boldsymbol{\sigma}^{(\epsilon)}\|^2 + B_1(\theta,\epsilon,n)$. We may rearrange the terms in (4) and obtain

$$\mathbb{E}\Big[\sum_{ij}\bar{q}_{ij}^{(\epsilon)}\Big] \le \frac{\beta_1}{\epsilon} + \frac{\beta_1}{\epsilon}\frac{\Pr\{\bar{r}^{(\epsilon)} > 0\}}{\epsilon - \Pr\{\bar{r}^{(\epsilon)} > 0\}}. \tag{7}$$

Comparing with the expected queue length upper bound of the MaxWeight policy in switches without reconfiguration delay [15, Theorem 1], we note that the first term in (7) has the same scaling with respect to ϵ in the heavy traffic limit. Therefore, the second term can be viewed as the overhead incurred by the reconfiguration delay and is determined by the probability of reconfiguration $\Pr\{\bar{r}^{(\epsilon)} > 0\}$. We continue the analysis of the probability of reconfiguration $\Pr\{\bar{r}^{(\epsilon)} > 0\}$ through the expected schedule duration in the next subsection.

We now proceed with the proof of Theorem 1.

Proof of Theorem 1 For simplicity of notation, we drop the superscript (ϵ) in the following proof. We also use $\mathbb{E}_{\bar{\mathbf{X}}}[\,\cdot\,]$ to denote the expectation given that $X(t)$ follows the steady-state distribution $\bar{\mathbf{X}}$.

In the proof of Theorem 1, we consider a drift analysis technique from [15], which devises a Lyapunov function that is catered to the geometric properties of the cone $\mathcal{K}$. Consider the following Lyapunov function from [15]:

$$V(\mathbf{X}) = \sum_i\Big(\sum_j q_{ij}\Big)^2 + \sum_j\Big(\sum_i q_{ij}\Big)^2 - \frac{1}{n}\Big(\sum_{ij} q_{ij}\Big)^2.$$

It may be shown using Lemma 1 that for steady-state $\bar{\mathbf{X}}$, the expectation $\mathbb{E}[V(\bar{\mathbf{X}})]$ is finite. We thus have zero drift for $V(\bar{\mathbf{X}})$ at steady state:

$$\mathbb{E}_{\bar{\mathbf{X}}}\Big[V(\mathbf{X}(t+1)) - V(\mathbf{X}(t))\Big] = \mathbb{E}_{\bar{\mathbf{X}}}\Big[V(\mathbf{X}(t+1))\Big] - \mathbb{E}_{\bar{\mathbf{X}}}\Big[V(\mathbf{X}(t))\Big] = 0.$$

We now evaluate the above drift terms with the queue length dynamics (1) and rewrite the expression as

$$\mathcal{T}_1 + \mathcal{T}_2 + \mathcal{T}_3 + \mathcal{T}_4 = 0, \tag{8}$$

where

$$\begin{aligned}
\mathcal{T}_1 = \mathbb{E}_{\bar{\mathbf{X}}}\Big[& 2\sum_i \Big(\sum_j q_{ij}(t)\Big)\Big(\sum_j \big(a_{ij}(t) - s_{ij}(t)\mathbb{1}_{\{r(t)=0\}}\big)\Big) \\
& + 2\sum_j \Big(\sum_i q_{ij}(t)\Big)\Big(\sum_i \big(a_{ij}(t) - s_{ij}(t)\mathbb{1}_{\{r(t)=0\}}\big)\Big) \\
& - \frac{2}{n}\Big(\sum_{ij} q_{ij}(t)\Big)\Big(\sum_{ij} \big(a_{ij}(t) - s_{ij}(t)\mathbb{1}_{\{r(t)=0\}}\big)\Big)\Big], \\
\mathcal{T}_2 = \mathbb{E}_{\bar{\mathbf{X}}}\Big[& \sum_i \Big(\sum_j \big(a_{ij}(t) - s_{ij}(t)\mathbb{1}_{\{r(t)=0\}}\big)\Big)^2 \\
& + \sum_j \Big(\sum_i \big(a_{ij}(t) - s_{ij}(t)\mathbb{1}_{\{r(t)=0\}}\big)\Big)^2 \\
& - \frac{1}{n}\Big(\sum_{ij} \big(a_{ij}(t) - s_{ij}(t)\mathbb{1}_{\{r(t)=0\}}\big)\Big)^2\Big], \\
\mathcal{T}_3 = \mathbb{E}_{\bar{\mathbf{X}}}\Big[& -\sum_i \Big(\sum_j u_{ij}(t)\Big)^2 - \sum_j \Big(\sum_i u_{ij}(t)\Big)^2 + \frac{1}{n}\Big(\sum_{ij} u_{ij}(t)\Big)^2\Big], \\
\mathcal{T}_4 = \mathbb{E}_{\bar{\mathbf{X}}}\Big[& 2\sum_i \Big(\sum_j q_{ij}(t+1)\Big)\Big(\sum_j u_{ij}(t)\Big) \\
& + 2\sum_j \Big(\sum_i q_{ij}(t+1)\Big)\Big(\sum_i u_{ij}(t)\Big) - \frac{2}{n}\Big(\sum_{ij} q_{ij}(t+1)\Big)\Big(\sum_{ij} u_{ij}(t)\Big)\Big].
\end{aligned}$$

The choice of the Lyapunov function $V(\cdot)$ is to make the $\mathcal{T}_4$ term zero when the queue length matrix $\mathbf{q}$ is in the cone $\mathcal{K}$ (see [15] for more detailed discussion). This allows us to combine with the WSSC result to obtain a tight bound.

We now simplify each term in the following: For the term T_1, since the schedule generated by the Adaptive MaxWeight policy is a maximal schedule,[1] we have $\sum_i s_{ij}(t) = \sum_j s_{ij}(t) = 1$, and $\sum_{ij} s_{ij}(t) = n$. On the other hand, since the packet arrival $\mathbf{a}(t)$ is independent of the queue length $\mathbf{q}(t)$, we may then apply $\mathbb{E}[\sum_i a_{ij}(t)] = \mathbb{E}[\sum_i a_{ij}(t)] = 1 - \epsilon$ and $\mathbb{E}[\sum_{ij} a_{ij}(t)] = n(1 - \epsilon)$ from the

[1] From the definition of the MaxWeight schedule, we know that for any maximal schedule $\mathbf{s}'$ that is not a maximal schedule, there always exists a maximal and MaxWeight schedule $\mathbf{s}^*$ which covers all the queues served by $\mathbf{s}'$, hence we can assume that the MaxWeight policy always generates a maximal schedule without loss of generality. It then follows that the Adaptive MaxWeight always generates a maximal schedule since its schedule is generated by the MaxWeight policy, either at the current or a previous time slot.

assumption of the packet arrival process, and simplify T_1 as

$$
\begin{aligned}
\mathcal{T}_1 &= \mathbb{E}_{\bar{\mathbf{X}}}\Big[2\Big(\sum_{ij} q_{ij}(t)\Big)\Big((1-\epsilon) - \mathbb{1}_{\{r(t)=0\}}\Big)\Big] \\
&= -2\epsilon \mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} q_{ij}(t)\Big] + 2\mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} q_{ij}(t)\mathbb{1}_{\{r(t)>0\}}\Big] \\
&= -2\epsilon \mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} q_{ij}(t)\Big] + 2\mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} q_{ij}(t) - \sum_{ij} q_{ij}(t)\mathbb{1}_{\{r(t)=0\}}\Big] \\
&= -2\epsilon \mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} q_{ij}(t)\Big] + 2\mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} q_{ij}(t)\Big]\Big(\Pr{}_{\bar{\mathbf{X}}}\{r(t)>0\} + \Pr{}_{\bar{\mathbf{X}}}\{r(t)=0\}\Big) \\
&\quad - 2\mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} q_{ij}(t)\Big| r(t)=0\Big]\Pr{}_{\bar{\mathbf{X}}}\{r(t)=0\} \\
&= -2\Big(\epsilon - \Pr{}_{\bar{\mathbf{X}}}\{r(t)>0\}\Big)\mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} q_{ij}(t)\Big] \\
&\quad + 2\Big(\mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} q_{ij}(t)\Big] - \mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} q_{ij}(t)\Big| r(t)=0\Big]\Big)\Pr{}_{\bar{\mathbf{X}}}\{r(t)=0\}. \qquad (9)
\end{aligned}
$$

For the term $\mathcal{T}_2$, we again use the fact that $s(t)$ is a maximal schedule, and from the assumption of the packet arrival process, we can derive $\mathbb{E}\Big[\sum_i(\sum_j a_{ij}(t)-1)^2\Big] = \mathbb{E}\Big[\sum_j(\sum_i a_{ij}(t)-1)^2\Big] = \|\boldsymbol{\sigma}\|^2 + n\epsilon^2$ and $\mathbb{E}\Big[(\sum_{ij} a_{ij}(t)-n)^2\Big] = \|\boldsymbol{\sigma}\|^2 + n^2\epsilon^2$, and thus

$$
\begin{aligned}
\mathcal{T}_2 &= \mathbb{E}_{\bar{\mathbf{X}}}\Big[\Big(\sum_i\big(\sum_j a_{ij}(t)-1\big)^2 + \sum_j\big(\sum_i a_{ij}(t)-1\big)^2 - \frac{1}{n}\big(\sum_{ij} a_{ij}(t)-n\big)^2\Big) \\
&\quad + \Big(\sum_i\big(2\sum_j a_{ij}(t)-1\big) + \sum_j\big(2\sum_i a_{ij}(t)-1\big) - \frac{1}{n}\big(2n\sum_{ij} a_{ij}(t)-n^2\big)\Big) \\
&\quad \times \mathbb{1}_{\{r(t)>0\}}\Big] \\
&= \mathbb{E}_{\bar{\mathbf{X}}}\Big[\Big(2(\|\boldsymbol{\sigma}\|^2+n\epsilon^2) - \frac{1}{n}(\|\boldsymbol{\sigma}\|^2+n^2\epsilon^2)\Big) + \Big(2n(1-\epsilon)-n\Big)\mathbb{1}_{\{r(t)>0\}}\Big] \\
&= \Big((2-\frac{1}{n})\|\boldsymbol{\sigma}\|^2 + n\epsilon^2\Big) + n(1-2\epsilon)\Pr{}_{\bar{\mathbf{X}}}\{r(t)>0\}. \qquad (10)
\end{aligned}
$$

For the term $\mathcal{T}_3$, since $u_{ij}(t) \le s_{ij}(t)$, we have $\sum_i u_{ij} \le 1$, $\sum_j u_{ij} \le 1$ and $\sum_{ij} u_{ij} \le n$. Therefore,

$$
\mathcal{T}_3 \le \mathbb{E}_{\bar{\mathbf{X}}}\Big[\frac{1}{n}\Big(\sum_{ij} u_{ij}(t)\Big)^2\Big] \le \mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} u_{ij}(t)\Big]
$$

and

$$\mathcal{T}_3 \geq -\mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_i \Big(\sum_j u_{ij}(t)\Big)^2\Big] - \mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_j \Big(\sum_i u_{ij}(t)\Big)^2\Big]$$
$$\geq -\mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_i \Big(\sum_j u_{ij}(t)\Big)\Big] - \mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_j \Big(\sum_i u_{ij}(t)\Big)\Big] = -2\mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} u_{ij}(t)\Big].$$

The above bounds involve the expected sum of unused services between schedule reconfiguration time instances $\mathbb{E}_{\bar{\mathbf{X}}}[\sum_{ij} u_{ij}(t)]$. One way to determine this value is to set the drift of $\sum_{ij} \bar{\mathbf{q}}_{ij}$ to zero: Since we have $\mathbb{E}_{\bar{\mathbf{X}}}[\sum_{ij} q_{ij}(t+1)] = \mathbb{E}_{\bar{\mathbf{X}}}[\sum_{ij} q_{ij}(t+1)]$ at steady state, we have

$$\begin{aligned}
&\mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} q_{ij}(t+1)\Big] - \mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} q_{ij}(t)\Big] \\
&= \mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij}\Big(a_{ij}(t) - s_{ij}(t)\mathbb{1}_{\{r(t)=0\}} + u_{ij}(t)\Big)\Big] \\
&= n(1-\epsilon) - n\Pr{}_{\bar{\mathbf{X}}}\{r(t)=0\}) + \mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} u_{ij}(t)\Big] \\
&= n(1-\epsilon) - n(1-\Pr{}_{\bar{\mathbf{X}}}\{r(t)>0\}) + \mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} u_{ij}(t)\Big] = 0
\end{aligned}$$

which implies

$$\mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} u_{ij}(t)\Big] = n\Big(\epsilon - \Pr{}_{\bar{\mathbf{X}}}\{r(t)>0\}\Big). \tag{11}$$

Hence, we have

$$-2n\Big(\epsilon - \Pr{}_{\bar{\mathbf{X}}}\{r(t)>0\}\Big) \leq \mathcal{T}_3 \leq n\Big(\epsilon - \Pr{}_{\bar{\mathbf{X}}}\{r(t)>0\}\Big). \tag{12}$$

For the term $\mathcal{T}_4$, we first follow the same line in [15] to simplify the expression. Rewrite each term and use $\mathbf{q}(t+1) = \mathbf{q}_{\|} + \mathbf{q}_{\perp}$ to obtain

$$\begin{aligned}
\mathcal{T}_4 &= 2\mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} u_{ij}(t)\Big(\sum_{j'} q_{ij'}(t+1) + \sum_{i'} q_{i'j}(t+1) - \frac{1}{n}\sum_{i'j'} q_{i'j'}(t+1)\Big)\Big] \\
&= 2\mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} u_{ij}(t)\Big(\sum_{j'} q_{\|ij'}(t+1) + \sum_{i'} q_{\|i'j}(t+1) - \frac{1}{n}\sum_{i'j'} q_{\|i'j'}(t+1)\Big) \\
&\quad + \sum_{ij} u_{ij}(t)\Big(\sum_{j'} q_{\perp ij'}(t+1) + \sum_{i'} q_{\perp i'j}(t+1) - \frac{1}{n}\sum_{i'j'} q_{\perp i'j'}(t+1)\Big)\Big].
\end{aligned}$$

Since $u_{ij}(t)q_{ij}(t+1) = 0$, when $u_{ij}(t) = 0$, we have $q_{ij}(t+1) = 0$ and thus $q_{\parallel ij}(t+1) = -q_{\perp ij}(t+1)$. Also, since $\mathbf{q}_{\parallel}$ is in the cone $\mathcal{K}$ by definition, we may use a property of the cone $\mathcal{K}$ [15, Lemma 1] to obtain

$$
\begin{aligned}
q_{\parallel ij}(t+1) &= \frac{1}{n}\sum_{j'} q_{\parallel ij'}(t+1) + \frac{1}{n}\sum_{i'} q_{\parallel i'j}(t+1) + \frac{1}{n^2}\sum_{i'j'} q_{\parallel i'j'}(t+1) \\
&= -q_{\perp ij}(t+1)
\end{aligned}
$$

and further simplify $\mathcal{T}_4$ as

$$
\begin{aligned}
\mathcal{T}_4 = 2\mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} u_{ij}(t)\Big(&-nq_{\perp ij}(t+1) + \sum_{j'} q_{\perp ij'}(t+1) + \sum_{i'} q_{\perp i'j}(t+1) \\
&-\frac{1}{n}\sum_{i'j'} q_{\perp i'j'}(t+1)\Big)\Big]
\end{aligned}
$$

from which we can note that $\mathcal{T}_4$ is zero when the residual component $\mathbf{q}_{\perp}$ is zero (which means that the queue length matrix $\mathbf{q} \in \mathcal{K}$). It then remains to express $\mathcal{T}_4$ in terms of $\mathbf{q}_{\perp}$ in order to apply the WSSC result.

We now simplify each term as an inner product using $\mathbf{e}^{(i)}$, $\tilde{\mathbf{e}}^{(j)}$, or $\mathbf{1}$, the all-ones matrix. In particular, with the following simplification:

$$
\begin{aligned}
\sum_{ij} u_{ij}(t)q_{\perp ij}(t+1), &= \Big\langle \mathbf{u}(t), \mathbf{q}_{\perp}(t+1)\Big\rangle, \\
\sum_{ij} u_{ij}(t)\sum_{j'} q_{\perp ij'}(t+1) &= \sum_{i}\Big(\sum_{j} u_{ij}(t)\sum_{j'} q_{\perp ij'}(t+1)\Big) \\
&= \sum_{i}\Big\langle \mathbf{u}(t), \mathbf{e}^{(i)}\Big\rangle\Big\langle \mathbf{q}_{\perp}(t+1), \mathbf{e}^{(i)}\Big\rangle, \\
\sum_{ij} u_{ij}(t)\sum_{i'} q_{\perp i'j}(t+1) &= \sum_{j}\Big(\sum_{i} u_{ij}(t)\sum_{i'} q_{\perp i'j}(t+1)\Big) \\
&= \sum_{j}\Big\langle \mathbf{u}(t), \tilde{\mathbf{e}}^{(j)}\Big\rangle\Big\langle \mathbf{q}_{\perp}(t+1), \tilde{\mathbf{e}}^{(j)}\Big\rangle, \\
\sum_{ij} u_{ij}(t)\sum_{i'j'} q_{\perp i'j'}(t+1) &= \Big\langle \mathbf{u}(t), \mathbf{1}\Big\rangle\Big\langle \mathbf{q}_{\perp}(t+1), \mathbf{1}\Big\rangle,
\end{aligned}
$$

we simplify $\mathcal{T}_4$ as

$$\mathcal{T}_4 = 2\mathbb{E}_{\bar{\mathbf{X}}}\Big[\Big\langle \mathbf{u}(t), -n\mathbf{q}_\perp(t+1) + \sum_i \langle \mathbf{q}_\perp(t+1), \mathbf{e}^{(i)}\rangle \mathbf{e}^{(i)} + \sum_j \langle \mathbf{q}_\perp(t+1), \tilde{\mathbf{e}}^{(j)}\rangle \tilde{\mathbf{e}}^{(j)} - \frac{1}{n}\langle \mathbf{q}_\perp(t+1), \mathbf{1}\rangle \mathbf{1}\Big\rangle\Big].$$

We now derive bounds for the simplified $\mathcal{T}_4$ term. From the definition of the cone $\mathcal{K}$, we can derive $\langle \mathbf{q}_\perp, \mathbf{e}^{(i)}\rangle \leq 0$ for all i, $\langle \mathbf{q}_\perp, \tilde{\mathbf{e}}^{(j)}\rangle \leq 0$ for all j, and $\langle \mathbf{q}_\perp, \mathbf{1}\rangle \leq 0$. We then obtain

$$\begin{aligned}\mathcal{T}_4 &\leq 2\mathbb{E}_{\bar{\mathbf{X}}}\Big[\Big\langle \mathbf{u}(t), -n\mathbf{q}_\perp(t+1) - \frac{1}{n}\langle \mathbf{q}_\perp(t+1), \mathbf{1}\rangle \mathbf{1}\Big\rangle\Big] \\ &\leq 2\mathbb{E}_{\bar{\mathbf{X}}}\Big[\big\|\mathbf{u}(t)\big\| \big\| - n\mathbf{q}_\perp(t+1) - \frac{1}{n}\langle \mathbf{q}_\perp(t+1), \mathbf{1}\rangle \mathbf{1}\big\|\Big],\end{aligned}$$

where the last inequality follows from the Cauchy–Schwartz inequality. Note that the second term can further be bounded as

$$\begin{aligned}\Big\| - n\mathbf{q}_\perp(t+1) - \frac{1}{n}\langle \mathbf{q}_\perp(t+1), \mathbf{1}\rangle \mathbf{1}\Big\| &\overset{(a)}{\leq} n\|\mathbf{q}_\perp(t+1)\| + \frac{1}{n}|\langle \mathbf{q}_\perp(t+1), \mathbf{1}\rangle|\|\mathbf{1}\| \\ &\overset{(b)}{\leq} n\|\mathbf{q}_\perp(t+1)\| + \frac{\|\mathbf{1}\|\|\mathbf{1}\|}{n}\|\mathbf{q}_\perp(t+1)\| \\ &= 2n\|\mathbf{q}_\perp(t+1)\|,\end{aligned}$$

where (a) follows from the triangle inequality, and (b) follows from the Cauchy–Schwartz inequality. We then obtain

$$\mathcal{T}_4 \leq 4n\mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{u}(t)\|\|\mathbf{q}_\perp(t+1)\|\Big].$$

Similarly, we can derive a lower bound for $\mathcal{T}_4$:

$$\begin{aligned}\mathcal{T}_4 &\geq 2\mathbb{E}_{\bar{\mathbf{X}}}\Big[\Big\langle \mathbf{u}(t), -n\mathbf{q}_\perp(t+1) + \sum_i \langle \mathbf{q}_\perp(t+1), \mathbf{e}^{(i)}\rangle \mathbf{e}^{(i)} \\ &\quad + \sum_j \langle \mathbf{q}_\perp(t+1), \tilde{\mathbf{e}}^{(j)}\rangle \tilde{\mathbf{e}}^{(j)}\Big\rangle\Big] \\ &\geq -2\mathbb{E}_{\bar{\mathbf{X}}}\Big[\big\|\mathbf{u}(t)\big\| \big\| - n\mathbf{q}_\perp(t+1) + \sum_i \langle \mathbf{q}_\perp(t+1), \mathbf{e}^{(i)}\rangle \mathbf{e}^{(i)} \\ &\quad + \sum_j \langle \mathbf{q}_\perp(t+1), \tilde{\mathbf{e}}^{(j)}\rangle \tilde{\mathbf{e}}^{(j)}\big\|\Big]\end{aligned}$$

$$\geq -2\mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{u}(t)\|\Big(\|-n\mathbf{q}_\perp(t+1)\|+\|\sum_i\langle\mathbf{q}_\perp(t+1),\mathbf{e}^{(i)}\rangle\mathbf{e}^{(i)}\|$$

$$+\|\sum_j\langle\mathbf{q}_\perp(t+1),\tilde{\mathbf{e}}^{(j)}\rangle\tilde{\mathbf{e}}^{(j)}\|\Big)\Big].$$

Since $\langle\mathbf{e}^{(i)},\mathbf{e}^{(i)}\rangle = n$ and $\langle\mathbf{e}^{(i)},\mathbf{e}^{(i')}\rangle = 0$ for $i \neq i'$, we have

$$\begin{aligned}
&\|\sum_i\langle\mathbf{q}_\perp(t+1),\mathbf{e}^{(i)}\rangle\mathbf{e}^{(i)}\|\\
&=\Big\langle\sum_i\langle\mathbf{q}_\perp(t+1),\mathbf{e}^{(i)}\rangle\mathbf{e}^{(i)},\sum_i\langle\mathbf{q}_\perp(t+1),\mathbf{e}^{(i)}\rangle\mathbf{e}^{(i)}\Big\rangle^{\frac{1}{2}}\\
&=\Big(\sum_i n\langle\mathbf{q}_\perp(t+1),\mathbf{e}^{(i)}\rangle^2\Big)^{\frac{1}{2}}\\
&=\Big(\sum_i n\Big(\sum_j q_{\perp ij}(t+1)\Big)^2\Big)^{\frac{1}{2}}\\
&\leq\Big(\sum_i n^2\Big(\sum_j q_{\perp ij}^2(t+1)\Big)\Big)^{\frac{1}{2}}\\
&= n\|\mathbf{q}_\perp(t+1)\|.
\end{aligned}$$

Similarly, we can derive

$$\|\sum_j\langle\mathbf{q}_\perp(t+1),\tilde{\mathbf{e}}^{(j)}\rangle\tilde{\mathbf{e}}^{(j)}\| \leq n\|\mathbf{q}_\perp(t+1)\|$$

and thus

$$\mathcal{T}_4 \geq -6n\mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{u}(t)\|\|\mathbf{q}_\perp(t+1)\|\Big].$$

We now bound the common term $\mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{u}(t)\|\|\mathbf{q}_\perp(t+1)\|\Big]$ in the upper and lower bounds. Since each component of $\mathbf{u}(t)$ is either 0 or 1, we have $\|\mathbf{u}(t)\| = \sqrt{\sum_{ij}u_{ij}(t)} \leq \sum_{ij}u_{ij}(t)$. Also, with the boundedness of the packet arrival, we have $\|\mathbf{q}(t+1)\| \leq \|\mathbf{q}(t)\| + na_{\max}$, which implies $\|\mathbf{q}_\perp(t+1)\| \leq \|\mathbf{q}_\perp(t)\| + na_{\max}$. Therefore,

$$\begin{aligned}
\mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{u}(t)\|\|\mathbf{q}_\perp(t+1)\|\Big] &\leq \mathbb{E}_{\bar{\mathbf{X}}}\Big[\Big(\sum_{ij}u_{ij}(t)\Big)\Big(\|\mathbf{q}_\perp(t)\|+na_{\max}\Big)\Big]\\
&= \mathbb{E}_{\bar{\mathbf{X}}}\Big[\Big(\sum_{ij}u_{ij}(t)\Big)\|\mathbf{q}_\perp(t)\|\Big]+\mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij}u_{ij}(t)\Big]na_{\max}.
\end{aligned}$$

Since $\sum_{ij} u_{ij}(t) \leq n$, we may write $\sum_{ij} u_{ij}(t) \leq n\mathbb{1}_{\{\sum_{ij} u_{ij}(t)>0\}}$ and thus

$$\begin{aligned}
&\mathbb{E}_{\bar{\mathbf{X}}}\Big[\Big(\sum_{ij} u_{ij}(t)\Big)\|\mathbf{q}_{\perp}(t)\|\Big] \\
&\leq n\mathbb{E}_{\bar{\mathbf{X}}}\Big[\mathbb{1}_{\{\sum_{ij} u_{ij}(t)>0\}}\|\mathbf{q}_{\perp}(t)\|\Big] \\
&= n\Big(\mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{q}_{\perp}(t)\|\Big] - \mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{q}_{\perp}(t)\|\mathbb{1}_{\{\sum_{ij} u_{ij}(t)=0\}}\Big]\Big) \\
&= n\Big(\mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{q}_{\perp}(t)\|\Big]\mathbb{E}_{\bar{\mathbf{X}}}\Big[\mathbb{1}_{\{\sum_{ij} u_{ij}(t)>0\}}\Big] + \mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{q}_{\perp}(t)\|\Big]\mathbb{E}_{\bar{\mathbf{X}}}\Big[\mathbb{1}_{\{\sum_{ij} u_{ij}(t)=0\}}\Big] \\
&\quad - \mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{q}_{\perp}(t)\|\mathbb{1}_{\{\sum_{ij} u_{ij}(t)=0\}}\Big]\Big).
\end{aligned}$$

Since $\mathbb{1}_{\{\sum_{ij} u_{ij}(t)>0\}} \leq \sum_{ij} u_{ij}(t)$, and $\mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{q}_{\perp}(t)\|\mathbb{1}_{\{\sum_{ij} u_{ij}(t)=0\}}\Big] = \mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{q}_{\perp}(t)\|\Big| \sum_{ij} u_{ij}(t) = 0\Big]\Pr_{\bar{\mathbf{X}}}\Big\{ \sum_{ij} u_{ij}(t) = 0\Big\}$, we further obtain

$$\begin{aligned}
&\mathbb{E}_{\bar{\mathbf{X}}}\Big[\Big(\sum_{ij} u_{ij}(t)\Big)\|\mathbf{q}_{\perp}(t)\|\Big] \\
&\leq n\mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{q}_{\perp}(t)\|\Big]\mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} u_{ij}(t)\Big] \\
&\quad + n\Big(\mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{q}_{\perp}(t)\|\Big] - \mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{q}_{\perp}(t)\|\Big| \sum_{ij} u_{ij}(t) = 0\Big]\Big)\Pr_{\bar{\mathbf{X}}}\Big\{ \sum_{ij} u_{ij}(t) = 0\Big\} \\
&\Rightarrow \mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{u}(t)\|\|\mathbf{q}_{\perp}(t+1)\|\Big] \\
&\leq n^2\Big(\epsilon - \Pr{}_{\bar{\mathbf{X}}}\{r(t) > 0\}\Big)\Big(\mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{q}_{\perp}(t)\|\Big] + a_{\max}\Big) \\
&\quad + n\Big(\mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{q}_{\perp}(t)\|\Big] - \mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{q}_{\perp}(t)\|\Big| \sum_{ij} u_{ij}(t) = 0\Big]\Big),
\end{aligned}$$

where (11) is applied.

We then have

$$\begin{aligned}
\mathcal{T}_4 \leq 4n^3&\Big(\epsilon - \Pr{}_{\bar{\mathbf{X}}}\{r(t) > 0\}\Big)\Big(\mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{q}_{\perp}(t)\|\Big] + a_{\max}\Big) \\
&+ 4n^2\Big(\mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{q}_{\perp}(t)\|\Big] - \mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{q}_{\perp}(t)\|\Big| \sum_{ij} u_{ij}(t) = 0\Big]\Big)
\end{aligned} \tag{13}$$

and

$$\begin{aligned}
\mathcal{T}_4 \geq -6n^3&\Big(\epsilon - \Pr{}_{\bar{\mathbf{X}}}\{r(t) > 0\}\Big)\Big(\mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{q}_{\perp}(t)\|\Big] + a_{\max}\Big) \\
&- 6n^2\Big(\mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{q}_{\perp}(t)\|\Big] - \mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{q}_{\perp}(t)\|\Big| \sum_{ij} u_{ij}(t) = 0\Big]\Big).
\end{aligned} \tag{14}$$

Applying (9), (10), (12), and (13) into (8), we obtain the following upper bound:

$$
\begin{aligned}
0 \le &-2\Big(\epsilon - \Pr_{\bar{\mathbf{X}}}\{r(t) > 0\}\Big)\mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} q_{ij}(t)\Big] + \left(2 - \frac{1}{n}\right)\|\boldsymbol{\sigma}\|^2 + n\epsilon^2 \\
&+ n(1-2\epsilon)\Pr_{\bar{\mathbf{X}}}\{r(t) > 0\} + n\Big(\epsilon - \Pr_{\bar{\mathbf{X}}}\{r(t) > 0\}\Big) \\
&+ 2\Big(\mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} q_{ij}(t)\Big] - \mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} q_{ij}(t)\Big| r(t) = 0\Big]\Big)\Pr_{\bar{\mathbf{X}}}\{r(t) = 0\} \\
&+ 4n^3\Big(\epsilon - \Pr_{\bar{\mathbf{X}}}\{r(t) > 0\}\Big)\Big(\mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{q}_\perp(t)\|\Big] + a_{\max}\Big) \\
&+ 2n^2\Big(\mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{q}_\perp(t)\|\Big] - \mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{q}_\perp(t)\|\Big|\sum_{ij} u_{ij}(t) = 0\Big]\Big) \\
\Rightarrow &\Big(\epsilon - \Pr_{\bar{\mathbf{X}}}\{r(t) > 0\}\Big)\Big(\mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} q_{ij}(t)\Big] - 2n^3\mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{q}_\perp(t)\|\Big]\Big) \\
\le &(1 - \frac{1}{2n})\|\boldsymbol{\sigma}\|^2 + \Big|\mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} q_{ij}(t)\Big] - \mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} q_{ij}(t)\Big| r(t) = 0\Big]\Big| + \frac{n\epsilon(1+\epsilon)}{2} \\
&+ 2n^3\Big(\epsilon - \Pr_{\bar{\mathbf{X}}}\{r(t) > 0\}\Big)a_{\max} \\
&+ 2n^2\Big|\mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{q}_\perp(t)\|\Big] - \mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{q}_\perp(t)\|\Big|\sum_{ij} u_{ij}(t) = 0\Big]\Big|,
\end{aligned} \tag{15}
$$

and similarly applying (9), (10), (12), and (14) into (8), we obtain a lower bound

$$
\begin{aligned}
&\Big(\epsilon - \Pr_{\bar{\mathbf{X}}}\{r(t) > 0\}\Big)\Big(\mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} q_{ij}(t)\Big] + 3n^3\mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{q}_\perp(t)\|\Big]\Big) \\
&\ge \left(1 - \frac{1}{2n}\right)\|\boldsymbol{\sigma}\|^2 - \Big|\mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} q_{ij}(t)\Big] - \mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} q_{ij}(t)\Big| r(t) = 0\Big]\Big| - n\epsilon(1+\epsilon) \\
&\quad - 3n^3\Big(\epsilon - \Pr_{\bar{\mathbf{X}}}\{r(t) > 0\}\Big)a_{\max} \\
&\quad - 3n^2\Big|\mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{q}_\perp(t)\|\Big] - \mathbb{E}_{\bar{\mathbf{X}}}\Big[\|\mathbf{q}_\perp(t)\|\Big|\sum_{ij} u_{ij}(t) = 0\Big]\Big|.
\end{aligned} \tag{16}
$$

Note that with the ergodicity of the Markov chain $\mathbf{X}(t)$, $\mathbb{E}_{\bar{\mathbf{X}}}[\sum_{ij} q_{ij}(t)]$ and $\mathbb{E}_{\bar{\mathbf{X}}}[\sum_{ij} q_{ij}(t)|r(t) = 0]$ equal the time average of sum of queue lengths, and that sum under $r(t) = 0$, respectively. From (11) and the nonnegativity of the unused service, we can derive $\Pr_{\bar{\mathbf{X}}}\{r(t) > 0\} \le \epsilon$, and thus the probability of $r(t) = 0$ approaches 1 as $\epsilon \to 0$. This means that the time average under $r(t) = 0$ only excludes a diminishing number of time instances. Then, since the change in the sum of queue lengths $|\sum_{ij} q_{ij}(t+1) - \sum_{ij} q_{ij}(t)| < n^2 a_{\max}$ is bounded, we have that the difference $\mathbb{E}_{\bar{\mathbf{X}}}[\sum_{ij} q_{ij}(t)] - \mathbb{E}_{\bar{\mathbf{X}}}[\sum_{ij} q_{ij}(t)|r(t) = 0] \to 0$ as $\epsilon \to 0$.

Similarly, since the probability of $\sum_{ij} u_{ij}(t) = 0$ approaches 1 as $\epsilon \to 0$, the time average under $\sum_{ij} u_{ij}(t) = 0$ only excludes a diminishing number of time instances.

Then, since the change in $\|\mathbf{q}_\perp(t)\|$ is bounded, *i.e.*, $\Big|\|\mathbf{q}_\perp(t+1)\| - \|\mathbf{q}_\perp(t)\|\Big| < na_{\max}$, we have that $\mathbb{E}_{\bar{\mathbf{X}}}[\|\mathbf{q}_\perp(t)\|] - \mathbb{E}_{\bar{\mathbf{X}}}[\|\mathbf{q}_\perp(t)\|| \sum_{ij} u_{ij}(t) = 0] \to 0$ as $\epsilon \to 0$.

Note that the drift analysis up to this point only uses the fact that each schedule is a maximal schedule and that the switch systems are positive recurrent under the scheduling policy. We now apply the WSSC result of Adaptive MaxWeight. From Proposition 1, we can derive $\mathbb{E}_{\bar{\mathbf{X}}}[\|\mathbf{q}_\perp(t)\|] \le \theta \mathbb{E}_{\bar{\mathbf{X}}}[\|\mathbf{q}(t)\|] + M_\theta \le \theta \mathbb{E}_{\bar{\mathbf{X}}}[\sum_{ij} q_{ij}(t)] + M_\theta$ for any θ satisfying $0 < \theta < 1/2$. Together with (15), we have

$$\Big(\epsilon - \Pr\{\bar{r} > 0\}\Big)\Big(1 - 2n^3\theta\Big)\Big(\mathbb{E}\Big[\sum_{ij} \bar{q}_{ij}\Big]\Big) \le \Big(1 - \frac{1}{2n}\Big)\|\boldsymbol{\sigma}\|^2 + B_1'(\theta, \epsilon, n),$$

where $B_1'(\theta, \epsilon, n) = \Big|\mathbb{E}_{\bar{\mathbf{X}}}[\sum_{ij} q_{ij}(t)] - \mathbb{E}_{\bar{\mathbf{X}}}[\sum_{ij} q_{ij}(t)|r(t) = 0]\Big| + \frac{n\epsilon(1+\epsilon)}{2} + 2n^2\Big|\mathbb{E}_{\bar{\mathbf{X}}}[\|\mathbf{q}_\perp(t)\|] - \mathbb{E}_{\bar{\mathbf{X}}}[\|\mathbf{q}_\perp(t)\|| \sum_{ij} u_{ij}(t) = 0]\Big| + 2n^3\epsilon(M_\theta + a_{\max})$. Since M_θ is a fixed constant according to Proposition 1, and the first and the third terms of $B_1'(\theta, \epsilon, n)$ approach 0 as $\epsilon \downarrow 0$, we then have that $\lim_{\epsilon \to 0} B_1'(\theta, \epsilon, n) = 0$. Dividing both sides by $(1 - 2n^3\theta)$ and defining $B_1(\theta, \epsilon, n) = B_1'(\theta, \epsilon, n)/(1 - 2n^3\theta)$ then gives (4).

Similarly, for the lower bound, we have

$$\Big(\epsilon - \Pr\{\bar{r} > 0\}\Big)\Big(1 + 3n^3\theta\Big)\Big(\mathbb{E}\Big[\sum_{ij} \bar{q}_{ij}\Big]\Big) \ge \Big(1 - \frac{1}{2n}\Big)\|\boldsymbol{\sigma}\|^2 + B_2'(\theta, \epsilon, n),$$

where $B_2'(\theta, \epsilon, n) = -\Big|\mathbb{E}_{\bar{\mathbf{X}}}[\sum_{ij} q_{ij}(t)] - \mathbb{E}_{\bar{\mathbf{X}}}[\sum_{ij} q_{ij}(t)|r(t) = 0]\Big| - n\epsilon(1 + \epsilon) - 2n^2\Big|\mathbb{E}_{\bar{\mathbf{X}}}[\|\mathbf{q}_\perp(t)\|] - \mathbb{E}_{\bar{\mathbf{X}}}[\|\mathbf{q}_\perp(t)\|| \sum_{ij} u_{ij}(t) = 0]\Big| - 3n^3\epsilon(M_\theta + a_{\max})$, and $\lim_{\epsilon \to 0} B_2'(\theta, \epsilon, n) = 0$. Dividing both sides by $(1 + 3n^3\theta)$ and defining $B_2(\theta, \epsilon, n) = B_2'(\theta, \epsilon, n)/(1 + 3n^3\theta)$ then gives (5). □

4.3 Expected schedule duration

In the previous subsection, we derived a bound on steady-state queue lengths. Note that this bound depends on the probability of reconfiguration $\Pr\{\bar{r}^{(\epsilon)} > 0\}$. In this subsection, we derive the mean schedule duration in order to evaluate this probability.

Recall that in the previous subsection, we set the drift of $\sum_{ij} \bar{q}_{ij}^{(\epsilon)}$ to zero to obtain (11), which is an expression of the total unused service. Here, we consider the drift of another Lyapunov function to obtain a different expression for the total unused service and combine the two expressions to derive the expected schedule duration.

We start by defining the schedule reconfiguration times as follows: Given the Markov chain $\mathbf{X}^{(\epsilon)}(t) = (\mathbf{q}^{(\epsilon)}(t), \mathbf{s}^{(\epsilon)}(t), r^{(\epsilon)}(t))$ representing the state of a switch with reconfiguration delay, the schedule reconfiguration times $\{t_k\}_{k=0}^{\infty}$ are defined as $t_0 = 0$ and $t_k = \min\{t : t > t_{k-1},\ \text{and}\ \mathbf{s}^{(\epsilon)}(t) \ne \mathbf{s}^{(\epsilon)}(t-1)\}$ for $k > 0$.

Now consider $\mathbf{X}^{(\epsilon)}(t)$ sampled at the schedule reconfiguration times $\{t_k\}$ and denoted as $\mathbf{X}_k^{(\epsilon)} = \mathbf{X}^{(\epsilon)}(t_k)$. Note that $\{t_k\}$ are stopping times with respect to $\mathbf{X}^{(\epsilon)}(t)$, hence, by the strong Markov property, $\mathbf{X}_k^{(\epsilon)}$ is also a Markov chain. Furthermore, the positive recurrence of $\mathbf{X}^{(\epsilon)}(t)$ implies the positive recurrence of $\mathbf{X}_k^{(\epsilon)}$. We then denote the steady-state distribution of $\mathbf{X}_k^{(\epsilon)}$ as $\hat{\mathbf{X}}^{(\epsilon)} = (\hat{\mathbf{q}}^{(\epsilon)}, \hat{\mathbf{s}}^{(\epsilon)}, \hat{r}^{(\epsilon)})$.

With the schedule reconfiguration times $\{t_k\}$, we now define the schedule duration $T^S(\mathbf{X}_k^{(\epsilon)}) = t_{k+1} - t_k$ as the number of time slots to the next schedule reconfiguration starting from state $\mathbf{X}_k^{(\epsilon)}$. We let $T^S(\hat{\mathbf{X}}^{(\epsilon)})$ denote the schedule duration when $\mathbf{X_k}^{(\epsilon)}$ follows the steady-state distribution $\hat{\mathbf{X}}^{(\epsilon)}$, and thus $\mathbb{E}[T^S(\hat{\mathbf{X}}^{(\epsilon)})]$ is the expected steady-state schedule duration. Since each schedule reconfiguration is followed by a reconfiguration delay of Δ_r time slots, the probability of reconfiguration is then determined by the expected steady-state schedule duration as $\Pr\{\bar{r}^{(\epsilon)} > 0\} = \Delta_r / \mathbb{E}[T^S(\hat{\mathbf{X}}^{(\epsilon)})]$.

The following theorem establishes a relation between the schedule duration and the queue length.

Theorem 2 *Consider a switch system with a fixed reconfiguration delay $\Delta_r > 0$, and the arrival process $\mathbf{a}(t)$ as described in Sect. 2 with the mean arrival rate vector given by $\boldsymbol{\lambda} = (1-\epsilon)\boldsymbol{\nu}$ for some $\boldsymbol{\nu} \in \mathcal{F}$. Suppose the switch system is operated under the Adaptive MaxWeight policy with hysteresis function $g(\cdot)$, where $g(\cdot)$ is a sublinear and strictly increasing function. Define the MaxWeight function $W^*(\mathbf{X}^{(\epsilon)}) = \max_{\mathbf{S}\in\mathcal{S}} \langle \mathbf{q}^{(\epsilon)}, \mathbf{S}\rangle$ for each state $\mathbf{X}^{(\epsilon)} = (\mathbf{q}^{(\epsilon)}, \mathbf{s}^{(\epsilon)}, r^{(\epsilon)}) \in \mathcal{X}$, and denote $\hat{\mathbf{W}}^{*^{(\epsilon)}} = W^*(\hat{\mathbf{X}}^{(\epsilon)})$. Then the following relation holds:*

$$\mathbb{E}\Big[T^S(\hat{\mathbf{X}}^{(\epsilon)})\Big] = \frac{\mathbb{E}\Big[g(\hat{\mathbf{W}}^{*^{(\epsilon)}}) + \delta_W\Big]}{(n-\alpha^{(\epsilon)})(1-\epsilon)}, \tag{17}$$

where δ_W satisfies $0 \le \delta_W < n(a_{\max}+1)$, and $\alpha^{(\epsilon)} = \Big\langle \boldsymbol{\nu}, \mathbb{E}_{\hat{\mathbf{X}}^{(\epsilon)}}[\mathbf{s}^{(\epsilon)}(t_k)]\Big\rangle$.

Proof For simplicity of the notation, we drop the superscript (ϵ) in the following proof. First, we define the Lyapunov function W on the state $\mathbf{X} = (\mathbf{q}, \mathbf{s}, r)$:

$$W(\mathbf{X}) = \Big\langle \mathbf{q}, \mathbf{s} \Big\rangle = \sum_{ij} q_{ij} s_{ij},$$

which is simply the schedule weight function. Note that $W(\mathbf{X}) \le \sum_{ij} q_{ij}$, hence the steady-state mean of $W(\mathbf{X})$ is finite. We may then set the drift of $W(\mathbf{X})$ between two schedule reconfiguration times to be zero. With an abuse of notation, we let $\mathbb{E}_{\hat{\mathbf{X}}}[\cdot]$ denote the expectation given that $\mathbf{X}_k = \mathbf{X}(t_k)$ is distributed as $\hat{\mathbf{X}}$.

$$\begin{aligned}
&\mathbb{E}_{\hat{\mathbf{X}}}\Big[W(\mathbf{X}(t_{k+1})) - W(\mathbf{X}(t_k))\Big] \\
&= \mathbb{E}_{\hat{\mathbf{X}}}\Big[\sum_{ij} q_{ij}(t_{k+1}) s_{ij}(t_{k+1}) - \sum_{ij} q_{ij}(t_k) s_{ij}(t_k)\Big]
\end{aligned}$$

$$
\begin{aligned}
&= \mathbb{E}_{\hat{\mathbf{X}}}\Big[\sum_{ij} q_{ij}(t_{k+1})\Big(s_{ij}(t_{k+1}) - s_{ij}(t_k)\Big) + \sum_{ij}\Big(q_{ij}(t_{k+1}) - q_{ij}(t_k)\Big)s_{ij}(t_k)\Big] \\
&= 0. \qquad (18)
\end{aligned}
$$

Note that the first term in (18) is the weight difference between the schedules $s(t_{k+1})$ and $s(t_k)$. From the definition of Adaptive MaxWeight, the weight difference exceeds the threshold $g(W^*(t_{k+1}))$ at time t_{k+1} and is less than the threshold at time $t_{k+1}-1$. Since at most a_{max} packets can arrive at a queue and at most 1 packet can leave a queue in one time slot, the maximum change of the weight difference in a time slot cannot exceed $n(a_{\max}+1)$. We may then write the first term as $\sum_{ij} q_{ij}(t_{k+1})\Big(s_{ij}(t_{k+1}) - s_{ij}(t_k)\Big) = g(W^*(t_{k+1})) + \delta_W$, where $0 \le \delta_W < n(a_{\max}+1)$.

For the second term in (18), we have

$$
\begin{aligned}
&\mathbb{E}_{\hat{\mathbf{X}}}\Big[\sum_{ij}\Big(q_{ij}(t_{k+1}) - q_{ij}(t_k)\Big)s_{ij}(t_k)\Big] \\
&\quad = \mathbb{E}_{\hat{\mathbf{X}}}\Big[\sum_{t=t_k}^{t_{k+1}-1}\sum_{ij}\Big(q_{ij}(t+1) - q_{ij}(t)\Big)s_{ij}(t_k)\Big] \\
&\quad = \mathbb{E}_{\hat{\mathbf{X}}}\Big[\sum_{t=t_k}^{t_{k+1}-1}\sum_{ij}\Big(a_{ij}(t) - s_{ij}(t)\mathbb{1}_{\{r(t)=0\}} + u_{ij}(t)\Big)s_{ij}(t_k)\Big] \\
&\quad = \mathbb{E}_{\hat{\mathbf{X}}}\Big[\sum_{t=t_k}^{t_{k+1}-1}\Big(\sum_{ij}\lambda_{ij}s_{ij}(t_k) - n\mathbb{1}_{\{r(t)=0\}} + \sum_{ij}u_{ij}(t)\Big)\Big],
\end{aligned}
$$

where the last equality is given by the fact that the schedule remains $s(t_k)$ for any time slot between t_k and t_{k+1}, therefore $\sum_{ij} s_{ij}(t)s_{ij}(t_k) = n$ and $\sum_{ij} u_{ij}(t)s_{ij}(t_k) = \sum_{ij} u_{ij}(t)s_{ij}(t) = \sum_{ij} u_{ij}(t)$ for any time slot $t \in [t_k, t_{k+1})$.

Since the arrival processes are independent from the scheduling decisions, we have $\mathbb{E}_{\hat{\mathbf{X}}}\big[\sum_{ij}\lambda_{ij}s_{ij}(t_k)\big] = (1-\epsilon)\big\langle \boldsymbol{\nu}, \mathbb{E}_{\hat{\mathbf{X}}}[\mathbf{s}(t_k)]\big\rangle$. Now define

$$
\alpha = \big\langle \boldsymbol{\nu}, \mathbb{E}_{\hat{\mathbf{X}}}[\mathbf{s}(t_k)]\big\rangle, \qquad (19)
$$

and we may then write (18) as

$$
\begin{aligned}
&\mathbb{E}_{\hat{\mathbf{X}}}\Big[g(W^*(t_{k+1})) + \delta_W + (t_{k+1} - t_k)\alpha(1-\epsilon) - n(t_{k+1} - t_k - \Delta_r) \\
&\qquad + \sum_{t=t_k}^{t_{k+1}-1}\sum_{ij} u_{ij}(t)\Big] = 0.
\end{aligned}
$$

We then rearrange the terms to obtain

$$\mathbb{E}_{\hat{\mathbf{X}}}\Bigg[\sum_{t=t_k}^{t_{k+1}-1}\sum_{ij} u_{ij}(t)\Bigg] = \Big(n-\alpha(1-\epsilon)\Big)\mathbb{E}_{\hat{\mathbf{X}}}\Big[t_{k+1}-t_k\Big] \\ - n\Delta_r - \mathbb{E}_{\hat{\mathbf{X}}}\Big[g(W^*(t_{k+1}))+\delta_W\Big]. \tag{20}$$

On the other hand, we may set the drift of $\sum_{ij} q_{ij}$ to zero and obtain

$$\mathbb{E}_{\hat{\mathbf{X}}}\Bigg[\sum_{t=t_k}^{t_{k+1}-1}\sum_{ij} u_{ij}(t)\Bigg] = n\Big(\epsilon\mathbb{E}_{\hat{\mathbf{X}}}\Big[t_{k+1}-t_k\Big]-\Delta_r\Big). \tag{21}$$

Combining (20) and (21), we then have

$$\mathbb{E}_{\hat{\mathbf{X}}}\Big[t_{k+1}-t_k\Big] = \frac{\mathbb{E}_{\hat{\mathbf{X}}}\Big[g(W^*(t_{k+1}))+\delta_W\Big]}{(n-\alpha)(1-\epsilon)} \tag{22}$$

and since $\mathbf{X}(t_{k+1})$ also follows the steady-state distribution when $\mathbf{X}(t_k)$ does, we have

$$\mathbb{E}\Big[T^S(\hat{\mathbf{X}})\Big] = \frac{\mathbb{E}\Big[g(\hat{\mathbf{W}}^*)+\delta_W\Big]}{(n-\alpha)(1-\epsilon)}. \tag{23}$$

□

Theorem 2 immediately implies the following lower bound on expected queue length when the hysteresis function $g(\cdot)$ is a concave function:

Corollary 1 *Given the switch system as described in Theorem 2, and in addition to being sublinear and strictly increasing, suppose the hysteresis function $g(\cdot)$ is also concave. Then then the expected maximum weight $\mathbb{E}\Big[\hat{\mathbf{W}}^{*(\epsilon)}\Big]$ satisfies*

$$\mathbb{E}\Big[\hat{\mathbf{W}}^{*(\epsilon)}\Big] \geq g^{-1}\Bigg(\frac{(n-\alpha^{(\epsilon)})(1-\epsilon)\Delta_r}{\epsilon} - \mathbb{E}\Big[\delta_W\Big]\Bigg), \tag{24}$$

where $\alpha^{(\epsilon)} = \Big\langle \boldsymbol{\nu}, \mathbb{E}_{\hat{\mathbf{X}}^{(\epsilon)}}[\mathbf{s}^{(\epsilon)}(t_k)]\Big\rangle$ and $0 \leq \delta_W < n(a_{\max}+1)$.

Proof For simplicity of the notation, we drop the superscript (ϵ) in the following proof. By the nonnegativity of the unused service $u_{ij}(t)$, we derive the following lower bound from (21):

$$\epsilon\mathbb{E}\Big[T^S(\hat{\mathbf{X}})\Big] - \Delta_r = \frac{1}{n}\mathbb{E}_{\hat{\mathbf{X}}}\Bigg[\sum_{t=t_k}^{t_{k+1}-1}\sum_{ij} u_{ij}(t)\Bigg] \geq 0$$

$$\Rightarrow \mathbb{E}\Big[T^S(\hat{\mathbf{X}})\Big] \geq \frac{\Delta_r}{\epsilon},$$

and using Theorem 2, we then have

$$\mathbb{E}\Big[g(\hat{\mathbf{W}}^*) + \delta_W\Big] = (n-\alpha)(1-\epsilon)\mathbb{E}\Big[T^S(\hat{\mathbf{X}})\Big] \geq \frac{(n-\alpha)(1-\epsilon)\Delta_r}{\epsilon}, \tag{25}$$

where $\alpha = \langle \boldsymbol{\nu}, \mathbb{E}_{\hat{\mathbf{X}}}[\mathbf{s}(t_k)]\rangle$. Since $g(\cdot)$ is strictly increasing and concave, its inverse function $g^{-1}(\cdot)$ is convex. Then by Jensen's inequality we have

$$\begin{aligned}\mathbb{E}\Big[\hat{\mathbf{W}}^*\Big] = \mathbb{E}\Big[g^{-1}\big(g(\hat{\mathbf{W}}^*)\big)\Big] &\geq g^{-1}\Big(\mathbb{E}\Big[g(\hat{\mathbf{W}}^*)\Big]\Big) \\ &\geq g^{-1}\bigg(\frac{(n-\alpha)(1-\epsilon)\Delta_r}{\epsilon} - \mathbb{E}\Big[\delta_W\Big]\bigg).\end{aligned} \tag{26}$$

□

In the rest of this section, we utilize Theorems 1 and 2 to derive bounds on the expected sum of queue lengths for switches with reconfiguration operating under the Adaptive MaxWeight policy, as shown in the following theorem.

Theorem 3 *Consider a set of switch systems with a fixed reconfiguration delay $\Delta_r > 0$, parametrized by $0 < \epsilon < 1$, where the arrival process of each system follows the same assumptions as in Theorem 1. In addition, the limiting arrival rate matrix $\boldsymbol{\nu} \in \mathcal{F}$ is assumed to satisfy $\|\boldsymbol{\nu}\|^2 < n$. Each switch system is operated under the Adaptive MaxWeight policy with sublinear hysteresis function $g(\cdot)$, where $g(\cdot)$ is strictly increasing and also concave. Let $\hat{\mathbf{W}}^{*(\epsilon)} = W^*(\hat{\mathbf{X}}^{(\epsilon)})$ be the maximum weight of the steady-state $\hat{\mathbf{X}}^{(\epsilon)}$. Then in the heavy traffic limit we have*

$$\lim_{\epsilon \downarrow 0} \epsilon \mathbb{E}\Big[g(\hat{\mathbf{W}}^{*(\epsilon)})\Big] = (n - \|\boldsymbol{\nu}\|^2)\Delta_r. \tag{27}$$

For the queue length sum $\sum_{ij} \hat{q}_{ij}^{(\epsilon)}$, we have the following bounds:

$$\limsup_{\epsilon \downarrow 0} \epsilon \mathbb{E}\Big[g\Big(\sum_{ij} \hat{q}_{ij}^{(\epsilon)}\Big)\Big] \leq n(n - \|\boldsymbol{\nu}\|^2)\Delta_r \tag{28}$$

and, for any θ such that $0 < \theta < 1/2$,

$$\mathbb{E}\Big[\sum_{ij} \hat{q}_{ij}^{(\epsilon)}\Big] \geq \frac{n}{1+n(n-1)\theta} g^{-1}\Big(\frac{(n-\alpha^{(\epsilon)})(1-\epsilon)\Delta_r}{\epsilon}\Big) - \frac{n(n-1)}{1+n(n-1)\theta} M_\theta, \tag{29}$$

where M_θ is as defined in Proposition 1, which is a function of $\theta, \tilde{\sigma}, a_{\max}, \nu_{\min}$ and n, but is independent of ϵ.

We first note that (29) determines the scaling of the queue length sum with respect to $\epsilon \downarrow 0$ as $\Omega\left(ng^{-1}\left(\frac{(n-\|\boldsymbol{\nu}\|^2)\Delta_r}{\epsilon}\right)\right)$. Recall the queue length upper bound (7) derived from Theorem 1, which includes a term independent of the reconfiguration delay with scaling $O\left(\frac{1}{\epsilon}\right)$, and the second term depending on the reconfiguration delay. The scaling result from (29) then suggests that the scaling of the second term of (7) is $\Omega\left(ng^{-1}\left(\frac{(n-\|\boldsymbol{\nu}\|^2)\Delta_r}{\epsilon}\right)\right)$, since the first term is negligible compared to the scaling.

On the other hand, while unfortunately (28) does not provide an asymptotically tight upper bound, it suggests an approximate scaling of $O\left(g^{-1}\left(\frac{n(n-\|\boldsymbol{\nu}\|^2)\Delta_r}{\epsilon}\right)\right)$, which becomes more accurate when g is closer to a linear function. Despite the lack of actual scaling behavior, we consider this scaling behavior as a conjecture and compare this scaling to the simulation results in Sect. 6.

Note that the hysteresis function g getting closer to a linear function is also the regime of interest when pursuing an optimal queue length bound, since $g^{-1}(n(n-\|\boldsymbol{\nu}\|^2)\Delta_r/\epsilon)$ provides a better scaling with respect to the traffic load when g is closer to a linear function. In other words, if we consider $g(x) = (1-\gamma)x^{1-\delta}$ as in [20] and take $\delta \to 0$, we not only get a tighter asymptotic bounds but also a better delay scaling. The only caveat here is that selecting g as exactly a linear function does not fit the analysis in this paper. In fact, it is even unclear whether throughput optimality could be guaranteed if g is linear.

The approximate queue length upper bound not only depends on ϵ and the reconfiguration delay Δ_r, but also depends on the limiting arrival rate matrix $\boldsymbol{\nu} \in \mathcal{F}$. It is not hard to show that $1 \leq \|\boldsymbol{\nu}\|^2 < n$. When the arrival rate is uniform, i.e., $\nu_{ij} = 1/n, \forall i, j$, we have that $\|\boldsymbol{\nu}\|^2 = 1$, and the expected queue length has the worst scaling $O\big(g^{-1}(n^2/\epsilon)\big)$. This result makes sense as the uniform arrival rate means that the switch system has to reconfigure between several different schedules in order to cover arrivals in every queue. On the other extreme, $\|\boldsymbol{\nu}\|^2$ attains maximum when $\boldsymbol{\nu}$ is close to a permutation matrix. Hence, if the arrival rate is highly nonuniform such that $\|\boldsymbol{\nu}\|^2$ is close to n, then the switch is more likely to stay on a few dominating schedules and the reconfiguration does not need to occur as often, and the expected queue length has better scaling as $O\big(g^{-1}(n/\epsilon)\big)$.

We now proceed with the proof of Theorem 3.

Proof We first start with the derivation for (27). Note the lower bound (25) derived in Corollary 1, which we rewrite as follows:

$$\epsilon\mathbb{E}\left[g(\hat{\mathbf{W}}^{*(\epsilon)})\right] \geq (n-\alpha^{(\epsilon)})(1-\epsilon)\Delta_r - \epsilon\mathbb{E}\left[\delta_W\right],$$

where $\alpha^{(\epsilon)} = \big\langle \boldsymbol{\nu}, \mathbb{E}_{\hat{\mathbf{X}}^{(\epsilon)}}[\mathbf{s}^{(\epsilon)}(t_k)]\big\rangle$. Note that by setting the expected drift of $s_{ij}^{(\epsilon)}(t_k)$ at steady state to be zero for each $i, j \in \{1, 2, \ldots, n\}$, we can show that for each i, j, $\mathbb{E}_{\hat{\mathbf{X}}^{(\epsilon)}}[s_{ij}^{(\epsilon)}(t_k)] \geq (1-\epsilon)\nu_{ij}$, which then implies $\lim_{\epsilon\downarrow 0} \mathbb{E}_{\hat{\mathbf{X}}^{(\epsilon)}}[\mathbf{s}(t_k)] = \boldsymbol{\nu}$. We thus have

$\lim_{\epsilon \downarrow 0} \alpha^{(\epsilon)} = \langle \boldsymbol{\nu}, \boldsymbol{\nu} \rangle = \|\boldsymbol{\nu}\|^2$ and the following asymptotic lower bound:

$$\liminf_{\epsilon \downarrow 0} \epsilon \mathbb{E}\Big[g(\hat{\mathbf{W}}^{*^{(\epsilon)}}) \Big] \geq (n - \|\boldsymbol{\nu}\|^2)(1-\epsilon)\Delta_r. \tag{30}$$

We now combine the results from Theorems 1 and 2 to derive an upper bound for $\mathbb{E}\Big[g(\hat{\mathbf{W}}^{*^{(\epsilon)}}) \Big]$. Recalling the relation between the probability of reconfiguration and the expected steady-state schedule duration, and using Theorem 2, we have

$$\Pr\{\bar{r}^{(\epsilon)} > 0\} = \frac{\Delta_r}{\mathbb{E}[T^S(\hat{\mathbf{X}}^{(\epsilon)})]} = \frac{(n - \alpha^{(\epsilon)})(1-\epsilon)\Delta_r}{\mathbb{E}[g(\hat{\mathbf{W}}^{*^{(\epsilon)}}) + \delta_W]}. \tag{31}$$

We can then apply (31) into (4) from Theorem 1. For the simplicity of notation, we denote $\hat{\mathbf{g}} = \mathbb{E}\Big[g(\hat{W}^*) + \delta_W \Big]$, $\bar{\mathbf{q}}_s = \mathbb{E}_{\bar{\mathbf{X}}}\Big[\sum_{ij} \bar{q}_{ij} \Big]$, $\beta = (1 - \frac{1}{2n})\|\tilde{\boldsymbol{\sigma}}\|^2/(1 - 2n^3\theta) + B_1(\theta, \epsilon, n)$ and obtain

$$\begin{aligned}
&\left(\epsilon - \frac{(n - \alpha^{(\epsilon)})(1-\epsilon)\Delta_r}{\hat{\mathbf{g}}} \right) \bar{\mathbf{q}}_s \leq \beta \\
&\quad \Rightarrow \left(\epsilon \hat{\mathbf{g}} - (n - \alpha^{(\epsilon)})(1-\epsilon)\Delta_r \right) \bar{\mathbf{q}}_s \leq \beta \hat{\mathbf{g}} \\
&\quad \Rightarrow \left(\epsilon - \frac{\beta}{\bar{\mathbf{q}}_s} \right) \hat{\mathbf{g}} \leq (n - \alpha^{(\epsilon)})(1-\epsilon)\Delta_r.
\end{aligned}$$

We thus have

$$\hat{\mathbf{g}} \leq (n - \alpha^{(\epsilon)})(1-\epsilon)\Delta_r \frac{\bar{\mathbf{q}}_s}{\epsilon \bar{\mathbf{q}}_s - \beta} = \frac{(n - \alpha^{(\epsilon)})(1-\epsilon)\Delta_r}{\epsilon} \left(1 + \frac{\beta}{\epsilon \bar{\mathbf{q}}_s - \beta} \right).$$

Note that $\sum_{ij} q_{ij} \geq \mathbf{W}^*$ for any state $\mathbf{X}$. We can then use Corollary 1 to obtain $\sum_{ij} \hat{\mathbf{q}}_{ij} \sim \Omega\big(g^{-1}(1/\epsilon)\big)$, and thus $\bar{\mathbf{q}}_s \sim \Omega\big(g^{-1}(1/\epsilon)\big)$ when $\epsilon \downarrow 0$. Therefore, we have $\frac{\beta}{\epsilon \bar{\mathbf{q}}_s - \beta} \to 0$ as $\epsilon \downarrow 0$, which then implies

$$\limsup_{\epsilon \downarrow 0} \epsilon \hat{\mathbf{g}} = \limsup_{\epsilon \downarrow 0} (n - \alpha^{(\epsilon)})\Delta_r = (n - \|\boldsymbol{\nu}\|^2)\Delta_r. \tag{32}$$

Since $\delta_W < n(a_{\max} + 1)$, we then have

$$\limsup_{\epsilon \downarrow 0} \epsilon \mathbb{E}\Big[g(\hat{\mathbf{W}}^{*^{(\epsilon)}}) \Big] = (n - \|\boldsymbol{\nu}\|^2)\Delta_r. \tag{33}$$

Combining (30) and (33), we obtain

$$\lim_{\epsilon \downarrow 0} \epsilon \mathbb{E}\Big[g(\hat{\mathbf{W}}^{*^{(\epsilon)}}) \Big] = (n - \|\boldsymbol{\nu}\|^2)\Delta_r. \tag{34}$$

From (34), we can then derive (28) and (29) using relations between the total queue length and the maximum weight, which we establish in the following paragraphs.

Note that the sum of queue length can be seen as the inner product between the queue length matrix $\mathbf{q}$ and an all-one matrix $\mathbf{1}$, namely $\sum_{ij} q_{ij} = \langle \mathbf{q}, \mathbf{1} \rangle$. Also, the all-one matrix $\mathbf{1}$ can be written as the sum of n disjoint (maximal) schedules. In particular, we can find a set of $n-1$ schedules $\{\mathbf{s}^{(1)}, \mathbf{s}^{(2)}, \ldots, \mathbf{s}^{(n-1)}\}$ such that $\mathbf{s}^* + \sum_{k=1}^{n-1} s^{(k)} = \mathbf{1}$.

By the definition of the MaxWeight schedule, we have that $\langle \mathbf{q}, \mathbf{s} \rangle \le \langle \mathbf{q}, \mathbf{s}^* \rangle = W^*$ for any schedule s, hence we can obtain

$$\sum_{ij} q_{ij} = \left\langle \mathbf{q}, \mathbf{s}^* + \sum_{k=1}^{n-1} s^{(k)} \right\rangle \le n W^*.$$

Combining this with (34), we then obtain (28) as

$$\begin{aligned} \limsup_{\epsilon \downarrow 0} \epsilon \mathbb{E}\Big[g\Big(\sum_{ij} \hat{q}_{ij}^{(\epsilon)}\Big)\Big] &\le \lim_{\epsilon \downarrow 0} \epsilon \mathbb{E}\Big[g(n \hat{\mathbf{W}}^{*^{(\epsilon)}})\Big] \\ &\le \lim_{\epsilon \downarrow 0} n \epsilon \mathbb{E}\Big[g(\hat{\mathbf{W}}^{*^{(\epsilon)}})\Big] = n(n - \|\boldsymbol{\nu}\|^2)\Delta_r, \end{aligned}$$

where we use the fact that $g(nx) \le n g(x)$ for sufficiently large x from the sublinearity of $g(\cdot)$.

Now, for (29), we apply another relation between the total queue length and the maximum weight, which utilize the property of cone $\mathcal{K}$. We can write the maximum weight as $W^* = \langle \mathbf{q}, \mathbf{s}^* \rangle = \langle \mathbf{q}_\perp + \mathbf{q}_\parallel, \mathbf{s}^* \rangle$ for the MaxWeight schedule $\mathbf{s}^*$ and $W = \langle \mathbf{q}, \mathbf{s} \rangle = \langle \mathbf{q}_\perp + \mathbf{q}_\parallel, \mathbf{s} \rangle$ for any other schedule $\mathbf{s}$. Note that $\mathbf{q}_\parallel$ is in cone $\mathcal{K}$, which implies $\langle \mathbf{q}_\parallel, \mathbf{s}^* \rangle = \langle \mathbf{q}_\parallel, \mathbf{s} \rangle$ for any schedule $\mathbf{s}$. We then have

$$W = W^* - \langle \mathbf{q}_\perp, \mathbf{s}^* - \mathbf{s} \rangle. \tag{35}$$

Since all the elements in a schedule $\mathbf{s}$ are either 0 or 1, we have $\langle \mathbf{q}_\perp, \mathbf{s}^* - \mathbf{s} \rangle \le \sum_{ij} |q_{\perp ij}| \le n\|\mathbf{q}\|$. On the other hand, recall that we can find a set of $n-1$ schedules $\{\mathbf{s}^{(1)}, \mathbf{s}^{(2)}, \ldots, \mathbf{s}^{(n-1)}\}$ such that $\mathbf{s}^* + \sum_{k=1}^{n-1} s^{(k)} = \mathbf{1}$ and write the sum of the queue lengths as $\sum_{ij} q_{ij} = \langle \mathbf{q}, \mathbf{s}^* + \sum_{k=1}^{n-1} s^{(k)} \rangle$. For each of these $n-1$ schedules, we have $W = W^* - \langle \mathbf{q}_\perp, \mathbf{s}^* - \mathbf{s} \rangle \ge W^* - n\|\mathbf{q}\|$. We can then lower bound $\sum_{ij} q_{ij}$ as

$$\sum_{ij} q_{ij} = \left\langle \mathbf{q}, \mathbf{s}^* + \sum_{k=1}^{n-1} s^{(k)} \right\rangle \ge n W^* - n(n-1)\|\mathbf{q}_\perp\|. \tag{36}$$

We now combine (34) and (36) to derive (29). Recall from Proposition 1, we have that $\mathbb{E}[\|\bar{\mathbf{q}}_\perp^{(\epsilon)}\|] \le \theta \mathbb{E}[\|\bar{\mathbf{q}}^{(\epsilon)}\|] + M_\theta$ for any $\theta \in (0, \frac{1}{2})$. We then use (36) and take expectation on the steady-state distribution $\hat{\mathbf{X}}^{(\epsilon)}$ to obtain

$$\mathbb{E}\Big[\sum_{ij}\hat{q}_{ij}^{(\epsilon)}\Big] \geq n\mathbb{E}\Big[\hat{\mathbf{W}}^{*(\epsilon)}\Big] - n(n-1)\mathbb{E}\Big[\|\hat{\mathbf{q}}_{\perp}^{(\epsilon)}\|\Big]$$

$$\geq n\mathbb{E}\Big[\hat{\mathbf{W}}^{*(\epsilon)}\Big] - n(n-1)\theta\mathbb{E}\Big[\sum_{ij}\hat{q}_{ij}^{(\epsilon)}\Big] - n(n-1)M_\theta$$

$$\Rightarrow \mathbb{E}\Big[\sum_{ij}\hat{q}_{ij}^{(\epsilon)}\Big] \geq \frac{n}{1+n(n-1)\theta}\mathbb{E}\Big[\hat{\mathbf{W}}^{*(\epsilon)}\Big] - \frac{n(n-1)}{1+n(n-1)\theta}M_\theta$$

$$\geq \frac{n}{1+n(n-1)\theta}g^{-1}\Big(\frac{(n-\alpha^{(\epsilon)})(1-\epsilon)\Delta_r}{\epsilon}\Big) - \frac{n(n-1)}{1+n(n-1)\theta}M_\theta.$$

□

5 Benchmark queue length behavior under reconfiguration delay

In this section, we derive some benchmark queue length behavior of switches with reconfiguration delay for the Adaptive MaxWeight policy to compare with. We start with a queue length lower bound for switch systems with reconfiguration delay, which determines a limit on the performance for any scheduling policy. In a later subsection, we then derive a queue length upper bound for a benchmark policy known as the Fixed Frame MaxWeight (FFMW) [11] policy. Although it is shown that the FFMW policy may achieve the optimal queue length scaling in the heavy traffic regime, this optimality would require perfect knowledge of the traffic load, which restricts its feasibility in practice.

5.1 Queue length lower bound with reconfiguration delay

The first proposition extends the analysis from [15, Proposition 1], which gives a universal lower bound on the expected queue length for switch systems without reconfiguration delay. The proof of the proposition couples the queue length process $\mathbf{q}(t)$ to that of a queueing system with less restricted schedule constraint and is given in Appendix B.

Proposition 2 *Consider a switch system with the arrival process $\mathbf{a}(t)$, which has mean $\boldsymbol{\lambda} = (1-\epsilon)\boldsymbol{\nu}$ for some $\boldsymbol{\nu} \in \mathcal{F}$ and variance $\boldsymbol{\sigma}^2$. Let $\mathbf{q}(t)$ denote the queue lengths process of the switch system. Suppose the switch system is stable under its scheduling policy, where the queue lengths process $\mathbf{q}(t)$ converges in distribution to a steady-state random vector $\bar{\mathbf{q}}$. The expected sum of queue lengths is lower bounded by*

$$\mathbb{E}\Big[\sum_{ij}\bar{q}_{ij}\Big] \geq \frac{\|\boldsymbol{\sigma}\|^2}{2(\epsilon-p)} - \frac{n(1-\epsilon)(\epsilon-2p)}{2(\epsilon-p)}, \tag{37}$$

where $p = \mathbb{E}[\mathbb{1}_{\{r(t)>0\}}]$ is the probability that the switch system is in reconfiguration under the given scheduling policy.

Note that the lower bound in Proposition 2 coincides with the lower bound in [15] when $p = 0$, and monotonically increases as p increases. This result is not surprising since the probability of reconfiguration p represents the portion of overhead caused by reconfiguration delay and should degrade the performance when p increases. However, the minimum of the lower bound occurring at $p = 0$ contradicts the intuition that not switching the schedule also hurts the performance. In fact, for $p = 0$, the switch is always stuck at one schedule and any queues that are not served by the schedule would increase without bound. In other words, the lower bound in Proposition 2 does not capture the effect of infrequent schedule reconfiguration.

To capture the effect of infrequent schedule reconfiguration, the following proposition lower bounds the expected queue lengths by examining the unserved queues when the switch is fixed at one schedule between two reconfiguration times. The proof of Proposition 3 is given in Appendix C.

Proposition 3 *Given a switch system with the arrival process* $\mathbf{a}(t)$*, which has mean* $\boldsymbol{\lambda} = (1-\epsilon)\boldsymbol{\nu}$ *for some* $\boldsymbol{\nu} \in \mathcal{F}$ *and variance* $\boldsymbol{\sigma}^2$*. For any scheduling policy under which the switch system is stable, and the system state* $\mathbf{X}(t) = \big(\mathbf{q}(t), \mathbf{s}(t), r(t)\big)$ *converges in distribution to a steady-state random vector* $\bar{\mathbf{X}} = \big(\bar{\mathbf{q}}, \bar{\mathbf{s}}, \bar{r}\big)$*, the average sum of queue lengths is lower bounded by*

$$\mathbb{E}\Bigg[\sum_{ij} \bar{q}_{ij}\Bigg] \geq \frac{\Delta_r}{2p}(1-\epsilon)(n-\bar{\alpha}), \tag{38}$$

where $\bar{\alpha} = \max_{\mathbf{S}\in\mathcal{S}} \langle \boldsymbol{\nu}, \mathbf{S} \rangle$*, and* $p = \mathbb{E}[\mathbb{1}_{\{r(t)>0\}}]$ *is the probability that the switch system is in reconfiguration under the given scheduling policy.*

With the two propositions above, we may then derive an optimal lower bound for a given switch system. In particular, for each reconfiguration probability p, the lower bound is given by the maximum of Eqs. (37) and (38). Due to the monotonicity with p of Eqs. (37) and (38), the reconfiguration probability p that minimizes the joint lower bound can be easily solved by equating Eqs. (37) and (38).

In this paper, we are particularly interested in the queue length scaling in the heavy traffic regime, where ϵ approaches 0. The following corollary provides a queue length lower bound in the heavy traffic regime.

Corollary 2 *Consider a sequence of switch systems with a fixed reconfiguration delay* $\Delta_r > 0$*, parametrized by* $0 < \epsilon < 1$*. For each switch system, let* $p^{(\epsilon)}$ *be the reconfiguration probability that minimizes the joint lower bounds of Eqs.* (37) *and* (38)*. Then in the heavy traffic regime, the reconfiguration probability satisfies*

$$\lim_{\epsilon\downarrow 0} \frac{p^{(\epsilon)}}{\epsilon} = \frac{\Delta_r(n-\bar{\alpha})}{\|\boldsymbol{\sigma}\|^2 + \Delta_r(n-\bar{\alpha})}, \tag{39}$$

where $\bar{\alpha} = \max_{\mathbf{S}\in\mathcal{S}} \langle \boldsymbol{\nu}, \mathbf{S} \rangle$*.*

Therefore, we have the following queue length lower bound in the heavy traffic limit:

$$\liminf_{\epsilon \downarrow 0} \epsilon \mathbb{E}\left[\sum_{ij} \bar{q}_{ij}\right] \geq \frac{\|\boldsymbol{\sigma}\|^2 + \Delta_r(n - \bar{\alpha})}{2}. \tag{40}$$

Proof Since Eq. (37) is monotonically increasing in p and Eq. (38) is monotonically decreasing in p, the minimizer $p^{(\epsilon)}$ can be solved by equating (37) and (38):

$$p^{(\epsilon)2} + \left(\frac{\|\boldsymbol{\sigma}\|^2}{2n(1-\epsilon)} - \frac{\epsilon}{2} + \frac{\Delta_r(n-\bar{\alpha})}{2n}\right)p^{(\epsilon)} - \frac{\epsilon\Delta_r(n-\bar{\alpha})}{2n} = 0.$$

Since $p^{(\epsilon)} \geq 0$, the only feasible solution for $p^{(\epsilon)}$ is given by

$$p^{(\epsilon)} = \frac{\sqrt{C^2 + x} - C}{2} = \frac{C}{2}\left(\sqrt{1 + \frac{x}{C^2}} - 1\right),$$

where $C = \frac{\|\boldsymbol{\sigma}\|^2}{2n(1-\epsilon)} - \frac{\epsilon}{2} + \frac{\Delta_r(n-\bar{\alpha})}{2n}$ and $x = 2\epsilon\Delta_r\frac{n-\bar{\alpha}}{n}$. Note that when $\frac{x}{C^2} << 1$, we have $p^{(\epsilon)} \approx \frac{x}{4C}$. Also, since $\frac{x}{C^2} \to 0$ as $\epsilon \downarrow 0$, we thus obtain

$$\lim_{\epsilon \downarrow 0} \frac{p^{(\epsilon)}}{\epsilon} = \lim_{\epsilon \downarrow 0} \frac{x}{4C} = \frac{\Delta_r(n-\bar{\alpha})}{\|\boldsymbol{\sigma}\|^2 + \Delta_r(n-\bar{\alpha})}.$$

□

Corollary 2 generalizes the lower bound from [15, Proposition 1] and characterizes the effect of the reconfiguration delay Δ_r on the delay performance. Note that $\bar{\alpha}$ in Corollary 2 is different from α defined in Eq. (19), and by definition $\alpha \leq \bar{\alpha} \leq n$ for any $\boldsymbol{\nu} \in \mathcal{F}$. Compare the queue length bound of Adaptive MaxWeight in Eq. (28) with Corollary 2, and we may see that when the hysteresis function $g(\cdot)$ approaches a linear function, the queue length behavior of Adaptive MaxWeight approximates the optimal scaling with respect to ϵ as well as the reconfiguration delay Δ_r in the heavy traffic limit $\epsilon \to 0$.

5.2 Queue length behavior of fixed frame MaxWeight

In this subsection, we analyze the queue length behavior of the Fixed Frame Max-Weight (FFMW) as a benchmark policy and compare it with that of the Adaptive MaxWeight policy.

The FFMW policy is a simple extension of the MaxWeight policy, which sets a fixed parameter T, and periodically reconfigures to the MaxWeight schedule every T time slots. It is shown in [11] that given the traffic load $\rho = 1 - \epsilon$, the switch system is stabilized by the FFMW policy with any period $T > \frac{\Delta_r}{\epsilon}$. Note that the FFMW policy requires knowledge of the traffic load, which limits the applicability of the policy in practice.

The following proposition extends the heavy traffic queue length analysis of the MaxWeight policy in [15] and gives an upper bound on the expected sum of queue lengths for switches with reconfiguration delay.

Proposition 4 *Consider a switch system with a fixed reconfiguration delay* $\Delta_r > 0$, *and the arrival process* $\mathbf{a}(t)$ *as described in Sect. 2. Suppose the mean arrival rate vector is given by* $\boldsymbol{\lambda} = (1-\epsilon)\boldsymbol{\nu}$, *where* $\boldsymbol{\nu} \in \mathcal{F}$ *is such that* $\nu_{\min} \triangleq \min_{ij} \nu_{ij} > 0$, *and for some* $\epsilon > 0$. *The variance of* $\mathbf{a}(t)$ *is* $\boldsymbol{\sigma}^2$. *For any* ϵ *that satisfies* $\epsilon < \frac{\nu_{\min}}{4n}$, *suppose that the switch system is operated under the Fixed Frame MaxWeight policy with schedule duration* $T > \frac{\Delta_r}{\epsilon}$. *Then the expected queue length satisfies*

$$\mathbb{E}\Big[\sum_{ij} \bar{q}_{ij}\Big] \leq \left(1 - \frac{1}{2n}\right) \frac{T}{\epsilon T - \Delta_r} \|\boldsymbol{\sigma}\|^2 + T\Big(\frac{n(1+\epsilon)}{2} + n^2(a_{\max} + 2M)\Big), \tag{41}$$

where $M = \frac{4n(a_{\max}+1)+4(\|\boldsymbol{\lambda}\|^2+\|\boldsymbol{\sigma}\|^2+n)+16\sqrt{2}n^2 a_{\max}^2}{\nu_{\min}} + 2n(2\sqrt{2}a_{\max} + 1)$.

We may then further minimize the upper bound over T *and derive the minimizing schedule duration as* $T^* = \frac{\Delta_r}{\epsilon}\left(1 + \sqrt{\frac{(1-\frac{1}{2n})\|\boldsymbol{\sigma}\|^2}{\Delta_r M'}}\right)$ *and the corresponding heavy traffic queue length upper bound is given by*

$$\limsup_{\epsilon \downarrow 0} \epsilon \mathbb{E}\Big[\sum_{ij} \bar{q}_{ij}\Big] \leq \left(\sqrt{(1 - \frac{1}{2n})\|\boldsymbol{\sigma}\|^2} + \sqrt{\Delta_r M'}\right)^2, \tag{42}$$

where $M' = \frac{n}{2} + n^2(a_{\max} + 2M)$.

From Proposition 4, we can see that given the traffic load information, the FFMW policy may achieve the optimal scaling with respect to ϵ and Δ_r in the heavy traffic limit $\epsilon \to 0$. Note that since Eq. (41) is not necessarily a tight bound, the derived minimizing schedule duration T^* may not be the true minimizer of the expected queue length. However, it does guarantee the optimal scaling and may be a good estimate for the true minimizer. From the expression of T^*, we can see that the minimizer is close to the boundary of stability $T = \frac{\Delta_r}{\epsilon}$; this also implies that the optimal delay scaling of the FFMW policy is rather intolerant to estimation error of the traffic load. Moreover, this issue becomes more prominent for large Δ_r, since it decreases the distance between T^* and the boundary of stability.

In the next section, we compare the average queue length performance between the Adaptive MaxWeight policy and the FFMW policy. We show that with the hysteresis function $g(\cdot)$ that is close to a linear function, the Adaptive MaxWeight policy has a comparable performance to the FFMW policy and does not require any knowledge of traffic load.

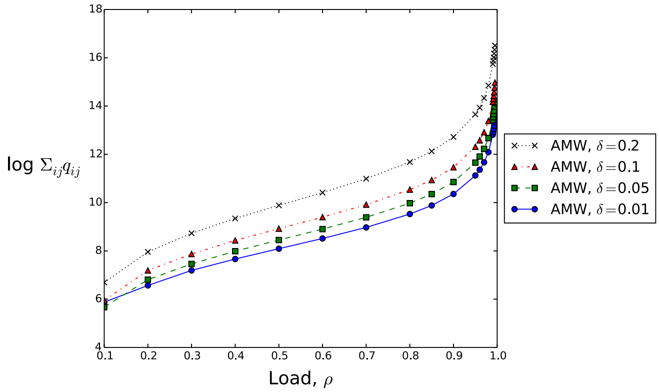

Fig. 1 Mean total queue length versus traffic load ρ under uniform traffic. Number of ports is $n = 16$, and reconfiguration delay $\Delta_r = 20$

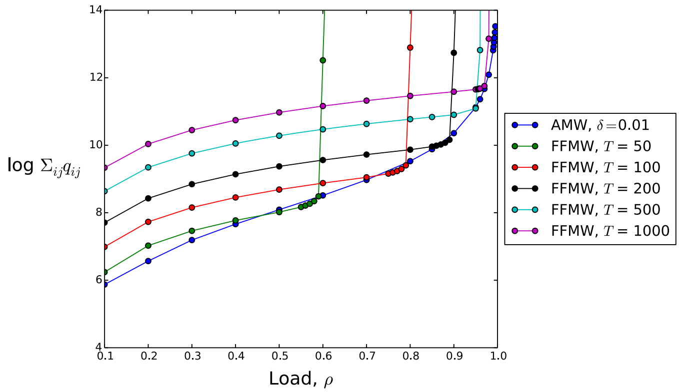

Fig. 2 Mean total queue length versus traffic load ρ under uniform traffic. Number of ports is $n = 16$, and reconfiguration delay $\Delta_r = 20$

6 Simulations

In this section, we show simulation results for switches with reconfiguration delay operated under the Adaptive MaxWeight policy, with hysteresis function $g(x) = (1-\gamma)x^{1-\delta}$. We first compare the simulation result with the Fixed Frame MaxWeight policy, and then determine the scaling of the average queue length with respect to different system parameters and compare with the queue length scaling derived in Sect. 4.

We now briefly describe the simulation setup. The arrival processes are assumed to be Poisson processes, all with the same arrival rate, which is also known as uniform traffic. More specifically, the matrix $\boldsymbol{\nu} \in \mathcal{F}$ satisfies $\nu_{ij} = \frac{1}{n}, \forall i, j \in \{1, \ldots, n\}$. Under uniform traffic, we have $\|\boldsymbol{\nu}\|^2 = 1$, and thus $\alpha = n - 1$ in Eq. (33). For the parameter of the hysteresis function g, since we are only interested in the scaling, we fix $\gamma = 0.1$, and consider $\delta \in \{0.01, 0.05, 0.1, 0.2\}$ for average queue length comparison.

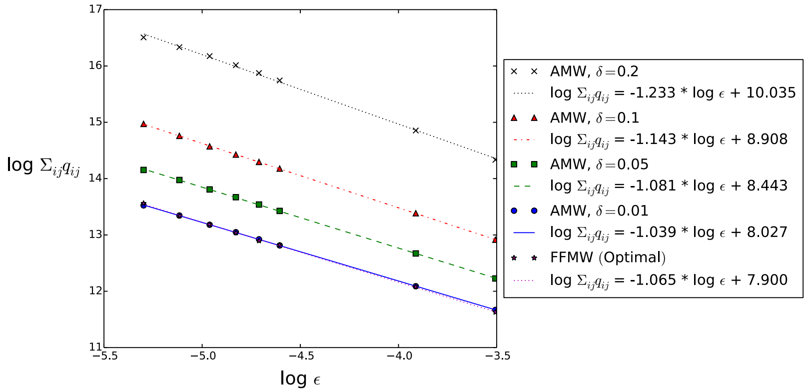

Fig. 3 Mean total queue length versus traffic load near the capacity region. Number of ports is $n = 16$, and reconfiguration delay $\Delta_r = 20$

Figure 1 shows the average queue length under various traffic loads $\rho \in [0.1, 1]$. From Fig. 1, we can see that the average queue length is smaller with smaller δ, in other words, the delay performance improves when the hysteresis function $g(x)$ approaches a linear function. The result implies that while the analysis in this work focuses on the heavy traffic regime, the conclusion that δ close to zero gives better delay performance also applies for lower traffic loads.

In Fig. 2, we keep Adaptive MaxWeight with $\delta = 0.01$, which has the best performance among the set of considered δ values, and compare the average queue length to the FFMW policy with various schedule durations $T \in \{50, 100, 200, 500, 1000\}$. We first note that for each schedule duration T, the average queue length grows quickly when ρ approaches $1 - \frac{\Delta_r}{T}$.

We now focus on simulations in the heavy traffic regime. In Fig. 3, we plot the average total queue length of Adaptive MaxWeight for various $\epsilon \in [0.005, 0.03]$ (corresponding to $\rho \in [0.97, 0.995]$), and take log scale for both axes. For the FFMW policy, we consider the optimal queue length performance over schedule durations. In other words, for each ϵ, we consider the FFMW policy with different schedule durations and take the one that minimizes the average queue length for comparison. We see that Adaptive MaxWeight with $\delta = 0.01$ closely follows the optimal performance of the FFMW policy. We then use linear regression to determine the scaling (i.e., the exponent) of the average queue length with respect to ϵ in the heavy traffic regime. With the scaling result from (33), the scaling with respect to ϵ is close to $g^{-1}(1/\epsilon)$, hence the theoretical exponent should be $-1/(1 - \delta)$, which would be $\{-1.010, -1.053, -1.111, -1.250\}$ for $\delta = \{0.01, 0.05, 0.1, 0.2\}$, respectively.

Figures 4 and 5 show the queue length scaling behavior under varying reconfiguration delays Δ_r and varying number of ports n, while the traffic load is fixed at $\rho = 0.96$ (or $\epsilon = 0.04$). For the reconfiguration delay, the scaling is $g^{-1}(\Delta_r)$, hence the theoretical exponents are $\{1.010, 1.053, 1.111, 1.250\}$ for $\delta \in \{0.01, 0.05, 0.1, 0.2\}$, respectively. We may see from Fig. 4 that the exponents obtained from the simulation result are close to our derived scaling. On the other hand, the scaling with respect to n is $g^{-1}(n^2)$, hence the theoretical exponents should be $\{2.020, 2.105, 2.222, 2.500\}$.

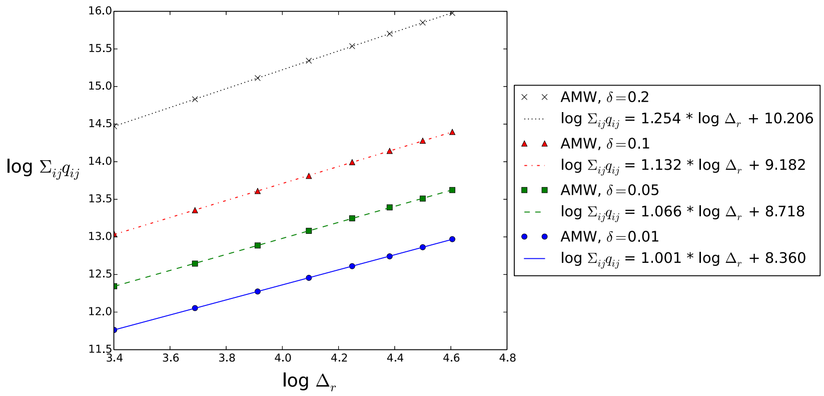

Fig. 4 Mean total queue length versus reconfiguration delay Δ_r. Number of ports is $n = 16$, and traffic load is $\rho = 0.96$

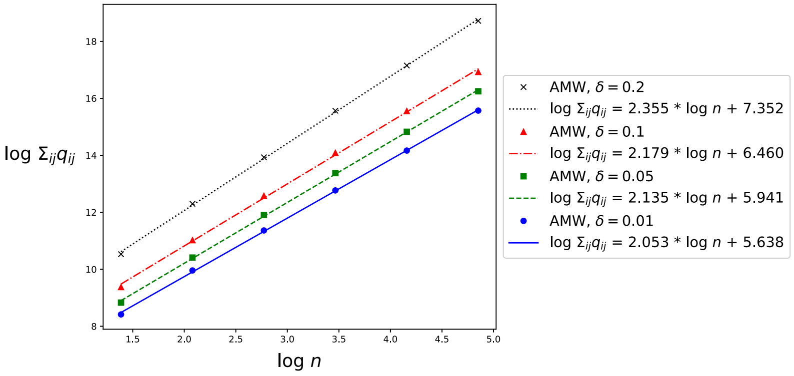

Fig. 5 Mean total queue length versus number of ports n. Reconfiguration delay is $\Delta_r = 20$, and traffic load is $\rho = 0.96$

We can see that the exponents derived from the simulation result are slightly larger than our derived scaling.

7 Conclusions

We consider the heavy traffic queue length behavior in input-queued switches with reconfiguration delay, operating under the Adaptive MaxWeight policy. It is shown that the Adaptive MaxWeight exhibits weak state space collapse behavior, which could be considered as an inheritance from the MaxWeight policy in the regime of zero reconfiguration delay. Utilizing the Lyapunov drift technique introduced in [15], we obtain a queue length upper bound in heavy traffic, which depends on the expected schedule duration. We then discover a relation between the expected schedule duration and the expected queue length, which then implies asymptotically tight bounds for the expected schedule duration in heavy traffic limit, thus determining its scaling. The scaling of the expected schedule duration then implies the dependence of the queue length scaling on the selection of the hysteresis function g, and this this scaling improves as g becomes closer to linear. Simulation results are also presented to illustrate the queue length scaling with respect to several system parameters (for example traffic load, number of ports, reconfiguration delay) and for comparison to the derived queue length scaling in heavy traffic.

The results obtained in this paper apply to traffic patterns in which all input and output ports are saturated. It would be interesting to consider the queue length behavior of Adaptive MaxWeight under incompletely saturated traffic, for example, the traffic conditions considered in [14]. The weak state space collapse result might be similar due to the inheritance from the MaxWeight policy, but the characterization for the expected schedule duration remains unclear at this point.

Acknowledgements This material is based upon work supported by the National Science Foundation under Grants CNS-1329819 and CCF-1850439.

Appendices

A Proof of Proposition 1

Proof For ease of notation, we drop the superscript (ϵ) in the following derivation. For each state $\mathbf{X} = (\mathbf{q}, \mathbf{s}, r)$, we define the Lyapunov function $Z(\mathbf{X}) = \max\{\|\mathbf{q}_\perp\| - \theta\|\mathbf{q}_\|\|, 0\}$. We then apply Lemma 1 with the Lyapunov function Z to obtain the result. Note that the selection of the Lyapunov function is such that Z is a nonnegative function. Since $\|\mathbf{q}_\perp\| - \theta\|\mathbf{q}_\|\| \le Z(\mathbf{X})$ for any state $\mathbf{X} = (\mathbf{q}, \mathbf{s}, r)$, the statement of Proposition 1 follows from a bound on $\mathbb{E}[Z(\bar{\mathbf{X}})]$.

We first verify Condition C.2 for $Z(\mathbf{X})$. Since $|\max\{a, 0\} - \max\{b, 0\}| \le |a - b|$ for any $a, b \in \mathbb{R}$, we have

$$\begin{aligned}
|\Delta^{\mathrm{T}} Z(\mathbf{X})| &\le \Big|\Big(\|\mathbf{q}_\perp(t+T)\| - \theta\|\mathbf{q}_\|(t+T)\|\Big) - \Big(\|\mathbf{q}_\perp(t)\| - \theta\|\mathbf{q}_\|(t)\|\Big)\Big| \\
&\le \Big|\|\mathbf{q}_\perp(t+T)\| - \|\mathbf{q}_\perp(t)\|\Big| + \theta\Big|\|\mathbf{q}_\|(t+T)\| - \|\mathbf{q}_\|(t)\|\Big| \\
&\le \|\mathbf{q}_\perp(t+T) - \mathbf{q}_\perp(t)\| + \theta\|\mathbf{q}_\|(t+T) - \mathbf{q}_\|(t)\| \\
&\le (1+\theta)\,\|\mathbf{q}(t+T) - \mathbf{q}(t)\| \\
&\le (1+\theta) n a_{\max} T \triangleq D.
\end{aligned} \tag{43}$$

Here, we use the fact that $\mathbf{q}_\perp$ is a projection onto $\mathcal{K}^\circ = \{\mathbf{x} \in \mathbb{R}^{n^2} : \langle \mathbf{x}, \mathbf{y}\rangle \le 0, \forall \mathbf{y} \in \mathcal{K}\}$, the polar cone of $\mathcal{K}$. Since the projection onto a cone is nonexpansive, we have $\|\mathbf{x}_\perp - \mathbf{y}_\perp\| \le \|\mathbf{x} - \mathbf{y}\|$ and $\|\mathbf{x}_\| - \mathbf{y}_\|\| \le \|\mathbf{x} - \mathbf{y}\|$, $\forall \mathbf{x}, \mathbf{y}$.

To verify condition C.1, we need to bound the expected T-step drift for $Z(\mathbf{X})$. For ease of notation, we denote $\mathbb{E}[\,\cdot\,|\mathbf{X}(t) = \mathbf{X}]$ as $\mathbb{E}_{\mathbf{X}}[\,\cdot\,]$.

Since (43) provides an upper bound on the magnitude of the T-step drift, we know that if $\mathbf{X}(t)$ satisfies $Z(\mathbf{X}) > D$, then $Z(\mathbf{X}(t+T))$ is positive, and we may drop the $\max\{\cdot, 0\}$ expression in the definition of $Z(\mathbf{X})$ in this case. In other words, for all $\mathbf{X}$

such that $Z(\mathbf{X}) > D$, we may write the expected T-step drift as

$$\mathbb{E}_{\mathbf{X}}\Big[\Delta^{\mathrm{T}} Z(\mathbf{X})\Big] = \mathbb{E}_{\mathbf{X}}\Big[\Big(\|\mathbf{q}_\perp(t+T)\| - \|\mathbf{q}_\perp(t)\|\Big) - \theta\Big(\|\mathbf{q}_\parallel(t+T)\| - \|\mathbf{q}_\parallel(t)\|\Big)\Big]. \tag{44}$$

Therefore, we need only to consider the T-step expected drift of $\|\mathbf{q}_\perp\|$ and $\|\mathbf{q}_\parallel\|$.

We first consider the drift of $\|\mathbf{q}_\perp\|$. The derivation follows along the lines of [15], where the relation in [15, Lemma 4] is used: Let $V(\mathbf{X}) = ||\mathbf{q}||^2$, $V_\parallel(\mathbf{X}) = ||\mathbf{q}_\parallel||^2$, and ΔV, $\Delta V_\parallel$ denote the one-step drift of V, $V_\parallel$, respectively. Then

$$\|\mathbf{q}_\perp(t+1)\| - \|\mathbf{q}_\perp(t)\| \le \frac{1}{2\|\mathbf{q}_\perp(t)\|}\big(\Delta V(\mathbf{X}(t)) - \Delta V_\parallel(\mathbf{X}(t))\big). \tag{45}$$

The inequality could be derived as follows:

$$\|\mathbf{q}_\perp(t+1)\| - \|\mathbf{q}_\perp(t)\| = \sqrt{\|\mathbf{q}_\perp(t+1)\|^2} - \sqrt{\|\mathbf{q}_\perp(t)\|^2} \le \frac{\|\mathbf{q}_\perp(t+1)\|^2 - \|\mathbf{q}_\perp(t)\|^2}{2\|\mathbf{q}_\perp(t)\|},$$

where the inequality follows from the concavity of the square root function: Since $f(x) = \sqrt{x}$, $x > 0$ is concave, we have $f(y) - f(x) \le (y-x)f'(x) = \frac{y-x}{2\sqrt{x}}$. Setting $x = \|\mathbf{q}_\perp(t)\|^2$ and $y = \|\mathbf{q}_\perp(t+1)\|^2$ gives the inequality. Then, with the orthogonality between $\mathbf{q}_\perp$ and $\mathbf{q}_\parallel$, we have $\|\mathbf{q}_\perp\|^2 = \|\mathbf{q}\|^2 - \|\mathbf{q}_\parallel\|^2$, and (45) follows by applying the relation for $\mathbf{q}_\perp(t+1)$ and $\mathbf{q}_\perp(t)$ and then rearranging the terms.

With (45), we have the following inequality for the T-step drift of $\|\mathbf{q}_\perp(t)\|$:

$$\begin{aligned}
\mathbb{E}_{\mathbf{X}}\Big[\|\mathbf{q}_\perp(t+T)\| - \|\mathbf{q}_\perp(t)\|\Big] &= \mathbb{E}_{\mathbf{X}}\Bigg[\sum_{\tau=t}^{t+T-1}\big(\|\mathbf{q}_\perp(\tau+1)\| - \|\mathbf{q}_\perp(\tau)\|\big)\Bigg] \\
&\le \mathbb{E}_{\mathbf{X}}\Bigg[\sum_{\tau=t}^{t+T-1}\frac{\Delta V(\mathbf{X}(\tau)) - \Delta V_\parallel(\mathbf{X}(\tau))}{2\|\mathbf{q}_\perp(\tau)\|}\Bigg] \\
&= \mathbb{E}_{\mathbf{X}}\Bigg[\sum_{\tau=t}^{t+T-1}\mathbb{E}\Big[\frac{\Delta V(\mathbf{X}(\tau)) - \Delta V_\parallel(\mathbf{X}(\tau))}{2\|\mathbf{q}_\perp(\tau)\|}\Big|\mathbf{X}(\tau)\Big]\Bigg].
\end{aligned} \tag{46}$$

We now derive bounds for ΔV and $\Delta V_\parallel$:

$$\begin{aligned}
&\mathbb{E}\Big[\Delta V(\mathbf{X}(\tau))\Big|\mathbf{X}(\tau)\Big] \\
&\quad = \mathbb{E}\Big[\|\mathbf{q}(\tau+1)\|^2 - \|\mathbf{q}(\tau)\|^2\Big|\mathbf{X}(\tau)\Big] \\
&\quad = \mathbb{E}\Big[\|\mathbf{q}(\tau) + \mathbf{a}(\tau) - \mathbf{s}(\tau)\mathbb{1}_{\{r(\tau)=0\}}\|^2 + \|\mathbf{u}(\tau)\|^2 \\
&\qquad + 2\Big\langle\mathbf{q}(\tau+1) - \mathbf{u}(\tau), \mathbf{u}(\tau)\Big\rangle - \|\mathbf{q}(\tau)\|^2\Big|\mathbf{X}(\tau)\Big]
\end{aligned}$$

$$
\begin{aligned}
&\leq \mathbb{E}\Big[\|\mathbf{q}(\tau)+\mathbf{a}(\tau)-\mathbf{s}(\tau)\mathbb{1}_{\{r(\tau)=0\}}\|^2-\|\mathbf{q}(\tau)\|^2\Big|\mathbf{X}(\tau)\Big]\\
&=\sum_{i,j}\mathbb{E}\Big[a_{ij}^2(\tau)+s_{ij}(\tau)\mathbb{1}_{\{r(\tau)=0\}}-2a_{ij}(\tau)s_{ij}(\tau)\mathbb{1}_{\{r(\tau)=0\}}\Big|\mathbf{X}(\tau)\Big]\\
&\quad+\mathbb{E}\Big[2\Big\langle\mathbf{q}(\tau),\boldsymbol{\lambda}-\mathbf{s}(\tau)\mathbb{1}_{\{r(\tau)=0\}}\Big\rangle\Big|\mathbf{X}(\tau)\Big]\\
&\overset{(a)}{\leq}\sum_{ij}(\lambda_{ij}^2+\sigma_{ij}^2)+n+2\Big\langle\mathbf{q}(\tau),(1-\epsilon)\boldsymbol{\nu}-\mathbf{s}(\tau)\Big\rangle+2\Big\langle\mathbf{q}(\tau),\mathbf{s}(\tau)\Big\rangle\mathbb{1}_{\{r(\tau)>0\}}\\
&=\|\boldsymbol{\lambda}\|^2+\|\boldsymbol{\sigma}\|^2+n-2\epsilon\Big\langle\mathbf{q}(\tau),\boldsymbol{\nu}\Big\rangle+2\Big\langle\mathbf{q}(\tau),\boldsymbol{\nu}-\mathbf{s}(\tau)\Big\rangle+2\Big\langle\mathbf{q}(\tau),\mathbf{s}(\tau)\Big\rangle\mathbb{1}_{\{r(\tau)>0\}},
\end{aligned}
$$

where (a) follows from $\mathbb{E}[a_{ij}^2]=\lambda_{ij}^2+\sigma_{ij}^2$, $a_{ij}(t)s_{ij}(t)\geq 0$ for all i,j, and $\sum_{ij}s_{ij}(t)=1$ for all t.

Suppose g is the sublinear hysteresis function for Adaptive MaxWeight, then by the sublinearity, there exists a constant K_θ such that $g(x)<\frac{\theta}{\alpha}x$ for any $x>K_\theta$, where $\alpha=\frac{8\|\boldsymbol{\nu}\|}{\nu_{\min}}$. Hence, by the definition of Adaptive MaxWeight, we have, for any $\mathbf{X}(\tau)$ such that $\langle\mathbf{q}(\tau),\mathbf{s}^*(\tau)\rangle>K_\theta$,

$$
\begin{aligned}
\Big\langle\mathbf{q}(\tau),\boldsymbol{\nu}-\mathbf{s}(\tau)\Big\rangle&=\Big\langle\mathbf{q}(\tau),\boldsymbol{\nu}-\mathbf{s}^*(\tau)\Big\rangle+\Big\langle\mathbf{q}(\tau),\mathbf{s}^*(\tau)-\mathbf{s}(\tau)\Big\rangle\\
&\leq\Big\langle\mathbf{q}(\tau),\boldsymbol{\nu}-\mathbf{s}^*(\tau)\Big\rangle+g\left(\Big\langle\mathbf{q}(\tau),\mathbf{s}^*(\tau)\Big\rangle\right)\\
&\leq\Big\langle\mathbf{q}(\tau),\boldsymbol{\nu}-\mathbf{s}^*(\tau)\Big\rangle+\frac{\theta}{\alpha}\Big\langle\mathbf{q}(\tau),\mathbf{s}^*(\tau)\Big\rangle\\
&=\Big(1-\frac{\theta}{\alpha}\Big)\Big\langle\mathbf{q}(\tau),\boldsymbol{\nu}-\mathbf{s}^*(\tau)\Big\rangle+\frac{\theta}{\alpha}\Big\langle\mathbf{q}(\tau),\boldsymbol{\nu}\Big\rangle.
\end{aligned}
$$

From [15, Claim 2], we have $\Big\langle\mathbf{q}(\tau),\boldsymbol{\nu}-\mathbf{s}^*(\tau)\Big\rangle\leq-\nu_{\min}\|\mathbf{q}_\perp(\tau)\|$. Therefore,

$$
\begin{aligned}
\mathbb{E}\Big[\Delta V(\mathbf{X}(\tau))\Big|\mathbf{X}(\tau)\Big]\leq\|\boldsymbol{\lambda}\|^2+\|\boldsymbol{\sigma}\|^2+n-2\epsilon\Big\langle\mathbf{q}(\tau),\boldsymbol{\nu}\Big\rangle-2\Big(1-\frac{\theta}{\alpha}\Big)\nu_{\min}\|\mathbf{q}_\perp(\tau)\|\\
+2\frac{\theta}{\alpha}\Big\langle\mathbf{q}(\tau),\boldsymbol{\nu}\Big\rangle+2\Big\langle\mathbf{q}(\tau),\mathbf{s}(\tau)\Big\rangle\mathbb{1}_{\{r(\tau)>0\}}.
\end{aligned}
\tag{47}
$$

For $\Delta V_\parallel$, we have

$$
\begin{aligned}
\mathbb{E}\Big[\Delta V_\parallel(\mathbf{X}(\tau))\Big|\mathbf{X}(\tau)\Big]&=\mathbb{E}\Big[\|\mathbf{q}_\parallel(\tau+1)\|^2-\|\mathbf{q}_\parallel(\tau)\|^2\Big|\mathbf{X}(\tau)\Big]\\
&=\mathbb{E}\Big[\Big\langle\mathbf{q}_\parallel(\tau+1)+\mathbf{q}_\parallel(\tau),\mathbf{q}_\parallel(\tau+1)-\mathbf{q}_\parallel(\tau)\Big\rangle\Big|\mathbf{X}(\tau)\Big]\\
&=\mathbb{E}\Big[\|\mathbf{q}_\parallel(\tau+1)-\mathbf{q}_\parallel(\tau)\|^2+2\Big\langle\mathbf{q}_\parallel(\tau),\mathbf{q}_\parallel(\tau+1)-\mathbf{q}_\parallel(\tau)\Big\rangle\Big|\mathbf{X}(\tau)\Big]\\
&\geq 2\mathbb{E}\Big[\Big\langle\mathbf{q}_\parallel(\tau),\mathbf{q}_\parallel(\tau+1)-\mathbf{q}_\parallel(\tau)\Big\rangle\Big|\mathbf{X}(\tau)\Big]\\
&=2\mathbb{E}\Big[\Big\langle\mathbf{q}_\parallel(\tau),\mathbf{q}(\tau+1)-\mathbf{q}(\tau)\Big\rangle-\Big\langle\mathbf{q}_\parallel(\tau),\mathbf{q}_\perp(\tau+1)-\mathbf{q}_\perp(\tau)\Big\rangle\Big|\mathbf{X}(\tau)\Big]
\end{aligned}
$$

$$
\begin{aligned}
&\overset{(b)}{\geq} 2\mathbb{E}\Big[\Big\langle \mathbf{q}_{\|}(\tau), \mathbf{a}(\tau) - \mathbf{s}(\tau)\mathbb{1}_{\{r(\tau)=0\}} + \mathbf{u}(\tau)\Big\rangle \Big| \mathbf{X}(\tau)\Big] \\
&\geq 2\Big\langle \mathbf{q}_{\|}(\tau), \boldsymbol{\lambda}\Big\rangle - 2\Big\langle \mathbf{q}_{\|}(\tau), \mathbf{s}(\tau)\mathbb{1}_{\{r(\tau)=0\}}\Big\rangle \\
&= -2\epsilon\Big\langle \mathbf{q}_{\|}(\tau), \boldsymbol{\nu}\Big\rangle + 2\Big\langle \mathbf{q}_{\|}(\tau), \boldsymbol{\nu} - \mathbf{s}(\tau)\Big\rangle + 2\Big\langle \mathbf{q}_{\|}(\tau), \mathbf{s}(\tau)\Big\rangle \mathbb{1}_{\{r(\tau)>0\}} \\
&= -2\epsilon\Big\langle \mathbf{q}_{\|}(\tau), \boldsymbol{\nu}\Big\rangle + 2\Big\langle \mathbf{q}_{\|}(\tau), \mathbf{s}(\tau)\Big\rangle \mathbb{1}_{\{r(\tau)>0\}}.
\end{aligned} \tag{48}
$$

For (b), we use the following properties of the projection onto cone $\mathcal{K}$: For $\mathbf{q} \in \mathbb{R}^{n^2}$, $\langle \mathbf{q}_{\|}, \mathbf{q}_{\perp}\rangle = 0$, and $\mathbf{q}_{\perp} \in \mathcal{K}^{\circ}$. Therefore $\langle \mathbf{q}_{\|}(t), \mathbf{q}_{\perp}(t)\rangle = 0$, and $\langle \mathbf{q}_{\|}(t), \mathbf{q}_{\perp}(t+1)\rangle \leq 0$.

Applying (47) and (48) in (46), we obtain

$$
\begin{aligned}
&\mathbb{E}_{\mathbf{X}}\Big[\|\mathbf{q}_{\perp}(t+T)\| - \|\mathbf{q}_{\perp}(t)\|\Big] \\
&\quad \leq \mathbb{E}_{\mathbf{X}}\Bigg[\sum_{\tau=t}^{t+T-1}\Big(\frac{\|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\sigma}\|^2 + n}{2\|\mathbf{q}_{\perp}(\tau)\|} - \epsilon\Big\langle \frac{\mathbf{q}_{\perp}(\tau)}{\|\mathbf{q}_{\perp}(\tau)\|}, \boldsymbol{\nu}\Big\rangle - \big(1 - \frac{\theta}{\alpha}\big)\nu_{\min} + \frac{\theta\big\langle \mathbf{q}(\tau), \boldsymbol{\nu}\big\rangle}{\alpha\|\mathbf{q}_{\perp}(\tau)\|} \\
&\qquad + \Big\langle \frac{\mathbf{q}_{\perp}(\tau)}{\|\mathbf{q}_{\perp}(\tau)\|}, \mathbf{s}(\tau)\Big\rangle \mathbb{1}_{\{r(\tau)>0\}}\Big)\Bigg] \\
&\quad \leq \mathbb{E}_{\mathbf{X}}\Bigg[T\bigg(\frac{\|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\sigma}\|^2 + n}{\min\limits_{\tau\in[t,t+T]} 2\|\mathbf{q}_{\perp}(\tau)\|} + \epsilon\|\boldsymbol{\nu}\| - \big(1-\theta\big)\nu_{\min} + \frac{1+\theta}{\alpha}\|\boldsymbol{\nu}\|\bigg) + \sqrt{n}\sum_{\tau=t}^{t+T-1}\mathbb{1}_{\{r(\tau)>0\}}\Bigg],
\end{aligned} \tag{49}
$$

where we have used the fact that $\|\mathbf{q}_{\perp}\| \geq \theta\|\mathbf{q}_{\|}\|$ implies $\theta\|\mathbf{q}\| \leq \theta(\|\mathbf{q}_{\|}\| + \|\mathbf{q}_{\perp}\|) \leq (1+\theta)\|\mathbf{q}_{\perp}\|$, and $\|\mathbf{s}(\tau)\| \leq \sqrt{n}$ for any schedule $\mathbf{s}(\tau) \in \mathcal{S}$.

On the other hand, the drift of $\|\mathbf{q}_{\|}\|$ can be obtained following (48):

$$
\begin{aligned}
\mathbb{E}_{\mathbf{X}}\Big[\|\mathbf{q}_{\|}(t+T)\| - \|\mathbf{q}_{\|}(t)\|\Big] &= \mathbb{E}_{\mathbf{X}}\Bigg[\sum_{\tau=t}^{t+T-1}\mathbb{E}\Big[\|\mathbf{q}_{\|}(\tau+1)\| - \|\mathbf{q}_{\|}(\tau)\|\Big|\mathbf{X}(\tau)\Big]\Bigg] \\
&\geq \mathbb{E}_{\mathbf{X}}\Bigg[\sum_{\tau=t}^{t+T-1}\mathbb{E}\Big[\frac{\|\mathbf{q}_{\|}(\tau+1)\|^2 - \|\mathbf{q}_{\|}(\tau)\|^2}{\|\mathbf{q}_{\|}(\tau+1)\| + \|\mathbf{q}_{\|}(\tau)\|}\Big|\mathbf{X}(\tau)\Big]\Bigg] \\
&\geq \mathbb{E}_{\mathbf{X}}\Bigg[\sum_{\tau=t}^{t+T-1}\mathbb{E}\Big[\frac{-2\epsilon\big\langle \mathbf{q}_{\|}(\tau), \boldsymbol{\nu}\big\rangle}{\|\mathbf{q}_{\|}(\tau+1)\| + \|\mathbf{q}_{\|}(\tau)\|}\Big|\mathbf{X}(\tau)\Big]\Bigg] \\
&\geq \mathbb{E}_{\mathbf{X}}\Bigg[\sum_{\tau=t}^{t+T-1}\frac{-2\epsilon\big\langle \mathbf{q}_{\|}(\tau), \boldsymbol{\nu}\big\rangle}{\|\mathbf{q}_{\|}(\tau)\|}\Bigg] \geq -2T\epsilon\|\boldsymbol{\nu}\|,
\end{aligned} \tag{50}
$$

where the last inequality follows from the Cauchy–Schwartz inequality.

Now, applying (49) and (50) in (44), we obtain

$$\mathbb{E}_{\mathbf{X}}\Big[\Delta^{\mathrm{T}} Z(\mathbf{X})\Big] \le \mathbb{E}_{\mathbf{X}}\Bigg[T\Bigg(\frac{\|\boldsymbol{\lambda}\|^2+\|\boldsymbol{\sigma}\|^2+n}{\min\limits_{\tau\in[t,t+T]} 2\|\mathbf{q}_{\perp}(\tau)\|}+(1+2\theta)\epsilon\|\boldsymbol{\nu}\|-\big(1-\theta\big)\nu_{\min}+\frac{1+\theta}{\alpha}\|\boldsymbol{\nu}\|\Bigg) + \sqrt{n}\sum_{\tau=t}^{t+T-1}\mathbb{1}_{\{r(\tau)>0\}}\Bigg]. \tag{51}$$

From [19, Lemma 1], we know that for any fixed $T > 0$, if $W^*(t) > g^{-1}\big(nT(a_{\max}+1)\big)+nT$, then at most one reconfiguration could occur within $[t, t+T]$, which gives $\sum_{\tau=t}^{t+T-1}\mathbb{1}_{r(\tau)>0} \le \Delta_r$.

Select $T = \frac{8\sqrt{n}\Delta_r}{\nu_{\min}}$, then set $D = \frac{3}{2}na_{\max}T = \frac{12n^{3/2}a_{\max}\Delta_r}{\nu_{\min}}$. Then, $\forall \mathbf{X}$ such that $Z(\mathbf{X}) > \kappa = \Big\{D, na_{\max}T + \frac{4(\|\boldsymbol{\lambda}\|^2+\|\boldsymbol{\sigma}\|^2+n)}{\nu_{\min}}, nK_\theta, ng^{-1}\big(nT(a_{\max}+1)\big)+n^2T\Big\}$, and $\forall\epsilon$ such that $0<\epsilon\le\frac{\nu_{\min}}{16\|\boldsymbol{\nu}\|}$, we have $\mathbb{E}_{\mathbf{X}}\Big[\Delta^{\mathrm{T}}Z(\mathbf{X})\Big] \le -\frac{(1-\theta)\nu_{\min}}{4} \le -\frac{\nu_{\min}}{8}$.

Hence, by Lemma 1, we have $\forall\epsilon$ such that $0<\epsilon\le\frac{\nu_{\min}}{16\|\boldsymbol{\nu}\|}$,

$$\mathbb{E}\left[\|\bar{\mathbf{q}}_{\perp}\| - \theta\|\bar{\mathbf{q}}_{\|}\|\right] \le \mathbb{E}\left[Z(\bar{\mathbf{X}})\right] \le \kappa + \frac{16D^2}{\nu_{\min}}.$$

Letting $M_\theta = \kappa + \frac{16D^2}{\nu_{\min}}$, we then have the result. □

B Proof of Proposition 2

Proof We will derive the lower bound by bounding the expected queue length sum at each input port, in particular, $\mathbb{E}\big[\sum_j \bar{q}_{ij}(t)\big]$ for each input port i.

The queue length dynamics of the coupled queue $\phi_i(t)$ are given by

$$\begin{aligned}\phi_i(t+1) &= \left[\phi_i(t)+b_i(t)-\mathbb{1}_{\{r(t)=0\}}\right]^+\\ &= \phi_i(t)+b_i(t)-\mathbb{1}_{\{r(t)=0\}}+v_i(t),\end{aligned}$$

where $v_i(t)$ is the unused service and satisfies $\phi_i(t+1)v_i(t)=0$.

We may show by induction that $\mathbb{E}\big[\sum_j \bar{q}_{ij}(t)\big] \ge \mathbb{E}\big[\bar{\phi}_i(t)\big]$. It then remains to lower bound $\mathbb{E}\big[\bar{\phi}_i(t)\big]$. We consider the expected drift of $\big(\bar{\phi}_i(t)\big)^2$ as follows:

$$\begin{aligned}&\mathbb{E}\big[\big(\bar{\phi}_i(t+1)\big)^2-\big(\bar{\phi}_i(t)\big)^2\big]\\ &= \mathbb{E}\big[\big(\bar{\phi}_i(t)+b_i(t)-\mathbb{1}_{\{r(t)=0\}}\big)^2-\big(v_i(t)\big)^2-\big(\bar{\phi}_i(t)\big)^2\big]\\ &= \mathbb{E}\big[2\bar{\phi}_i(t)\big(b_i(t)-\mathbb{1}_{\{r(t)=0\}}\big)+\big(b_i(t)-\mathbb{1}_{\{r(t)=0\}}\big)^2-\big(v_i(t)\big)^2\big]\\ &= \mathbb{E}\big[2\bar{\phi}_i(t)\big((1-\epsilon)-(1-\mathbb{1}_{\{r(t)>0\}})\big)+\big(b_i(t)-(1-\epsilon)+(\mathbb{1}_{\{r(t)>0\}}-\epsilon)\big)^2-v_i(t)\big]\\ &= -2\epsilon\mathbb{E}\big[\bar{\phi}_i(t)\big]+2\mathbb{E}\big[\bar{\phi}_i(t)\big]\mathbb{E}\big[\mathbb{1}_{\{r(t)>0\}}\big]+\mathrm{Var}(b_i(t))+\mathbb{E}\big[(\mathbb{1}_{\{r(t)>0\}}-\epsilon)^2\big]-\mathbb{E}\big[v_i(t)\big],\end{aligned}$$

where the last equality follows from the independence between the queue length process $\phi_i(t)$ and the schedule reconfiguration decision. We have

$$2(\epsilon - p)\mathbb{E}\big[\bar{\phi}_i(t)\big] = \sum_j \sigma_{ij}^2 + p - 2p\epsilon + \epsilon^2 - \mathbb{E}\big[v_i(t)\big].$$

Considering the drift of $\bar{\phi}_i(t)$, we can derive $\mathbb{E}\big[v_i(t)\big]$ as follows:

$$\begin{aligned}&\mathbb{E}\big[\bar{\phi}_t(t+1) - \bar{\phi}_i(t)\big] = (1-\epsilon) - (1 - \mathbb{E}\big[\mathbb{1}_{\{r(t)>0\}}\big]) + \mathbb{E}\big[v_i(t)\big] = 0\\ &\Rightarrow \mathbb{E}\big[v_i(t)\big] = \epsilon - \mathbb{E}\big[\mathbb{1}_{\{r(t)>0\}}\big] = \epsilon - p.\end{aligned}$$

We thus have

$$\mathbb{E}\big[\bar{\phi}_i(t)\big] = \frac{\sum_j \sigma_{ij}^2}{2(\epsilon - p)} - \frac{(1-\epsilon)(\epsilon - 2p)}{2(\epsilon - p)}.$$

Using $\mathbb{E}\big[\sum_j \bar{q}_{ij}(t)\big] \geq \mathbb{E}\big[\bar{\phi}_i(t)\big]$, and summing over each input port, we obtain

$$\mathbb{E}\Big[\sum_{ij} \bar{q}_{ij}(t)\Big] \geq \mathbb{E}\Big[\sum_i \bar{\phi}_i(t)\Big] \geq \frac{\sum_{ij} \sigma_{ij}^2}{2(\epsilon - p)} - \frac{(1-\epsilon)(\epsilon - 2p)}{2(\epsilon - p)}.$$

□

C Proof of Proposition 3

Proof Let t_k, $k = 1, \ldots$, denote the kth schedule reconfiguration time. By the assumption that the system state $\mathbf{X}(t) = \big(\mathbf{q}(t), \mathbf{s}(t), r(t)\big)$ converges in distribution to a steady-state random vector $\bar{\mathbf{X}} = \big(\bar{\mathbf{q}}, \bar{\mathbf{s}}, \bar{r}\big)$, we may consider a renewal-reward process associated with $\mathbf{X}(t)$ where t_k are the arrival epochs, while the reward function is the sum of queue lengths given by $R(t) = \sum_{ij} q_{ij}(t)$.

From renewal-reward theory, we have

$$\mathbb{E}\Big[\sum_{ij} q_{ij}(t)\Big] = \lim_{T\to\infty} \frac{1}{T}\mathbb{E}\Big[\sum_{t=0}^{T} q_{ij}(t)\Big] = \frac{\mathbb{E}\Big[\sum_{t=t_k}^{t_{k+1}-1} \sum_{ij} q_{ij}(t)\Big]}{\mathbb{E}[t_{k+1} - t_k]}.$$

For any (i, j) such that $S_{ij}(t_k) = 0$, we have $q_{ij}(t) = q_{ij}(t_k) + \sum_{\tau=t_k}^{t-1} a_{ij}(\tau) \geq \sum_{\tau=t_k}^{t-1} a_{ij}(\tau)$, while for any (i, j) such that $S_{ij}(t_k) = 1$, we have $q_{ij}(t) \geq 0$.

$$
\begin{aligned}
\mathbb{E}\left[\sum_{t=t_k}^{t_{k+1}-1} \sum_{ij} q_{ij}(t)\right] &= \mathbb{E}\left[\sum_{t=t_k}^{t_{k+1}-1} \mathbb{E}\left[\sum_{ij} q_{ij}(t) \middle| \mathbf{q}(t_k)\right]\right] \\
&\geq \mathbb{E}\left[\sum_{t=t_k}^{t_{k+1}-1} \sum_{(i,j):S_{ij}(t_k)=0} \lambda_{ij}(t - t_k)\right] \\
&\geq \mathbb{E}\left[\sum_{t=t_k}^{t_{k+1}-1} \left(t - t_k\right)\left(n(1-\epsilon) - \max_{\mathbf{S}} \langle \mathbf{S}, \Lambda \rangle\right)\right] \\
&= \mathbb{E}\left[\frac{(t_{k+1} - t_k)^2}{2}\right](n - \bar{\alpha})(1 - \epsilon).
\end{aligned}
$$

We then have a lower bound on the expected queue length sum as follows:

$$
\begin{aligned}
\mathbb{E}\left[\sum_{ij} q_{ij}(t)\right] &= \frac{\mathbb{E}[(t_{k+1} - t_k)^2]}{2\mathbb{E}[t_{k+1} - t_k]}(n - \bar{\alpha})(1 - \epsilon) \\
&\geq \frac{\mathbb{E}[t_{k+1} - t_k]}{2}(n - \bar{\alpha})(1 - \epsilon) \\
&= \frac{\Delta_r}{2p}(n - \bar{\alpha})(1 - \epsilon).
\end{aligned}
$$

□

D Proof of Proposition 4

Proof Given the period T, we consider the Markov chain $\mathbf{X}(t)$ being sampled at times $t_k = kT, k = 0, 1, \ldots$. Since the Fixed Frame MaxWeight policy stabilizes the system if $T > \frac{\Delta_r}{\epsilon}$, we know that $\mathbf{X}(t)$ converges to a steady-state distribution, and so does $\mathbf{X}(t_k)$. Let $\bar{\mathbf{X}} = (\bar{\mathbf{q}}, \bar{\mathbf{s}}, \bar{r})$ and $\hat{\mathbf{X}} = (\hat{\mathbf{q}}, \hat{\mathbf{s}}, \hat{r})$ denote the steady-state distribution of $\mathbf{X}(t)$ and $\mathbf{X}(t_k)$, respectively. By the assumption on the maximum arrival, we immediately have that $\mathbb{E}[\sum_{ij} \bar{q}_{ij}] \leq \mathbb{E}[\sum_{ij} \hat{q}_{ij}] + n^2 a_{\max} T$. It then remains to bound $\mathbb{E}[\sum_{ij} \hat{q}_{ij}]$ following the similar procedures in [15]:

1. Derive an upper bound on $\mathbb{E}[\|\mathbf{q}_{\perp}(t_k)\|^2]$.
2. Derive the queue length upper bound which depends on $\mathbb{E}[\|\mathbf{q}_{\perp}(t_k)\|^2]$.

Consider the Lyapunov function $Z(\mathbf{X}) = \|\mathbf{q}_\perp\|$. By the assumption on the maximum arrival, we have

$$|\Delta Z(\mathbf{X})| = \Big| \|\mathbf{q}_\perp(t_{k+1})\| - \|\mathbf{q}(t_k)\| \Big| \leq \|\mathbf{q}_\perp(t_{k+1}) - \mathbf{q}(t_k)\| = \sqrt{\sum_{ij} |q_{ij}(t_{k+1} - q_{ij}(t_k)|^2} \\ \leq n a_{\max} T. \tag{52}$$

For the expected drift at steady state, we have

$$\mathbb{E}\Big[\|\mathbf{q}_\perp(t_{k+1})\| - \|\mathbf{q}(t_k)\| \Big] \leq \mathbb{E}\Big[\sum_{\tau=t_k}^{t_{k+1}-1} \mathbb{E}\Big[\frac{\Delta V(\mathbf{X}(\tau)) - \Delta V_\|(\mathbf{X}(\tau))}{2\|\mathbf{q}_\perp\|} \Big| \mathbf{X}(\tau) \Big] \Big],$$

where $\Delta V(\mathbf{X})$ and $\Delta V_\|(\mathbf{X})$ are the drift of Lyapunov functions $V(\mathbf{X}) = \|\mathbf{q}\|^2$ and $V_\|(\mathbf{X}) = \|\mathbf{q}_\|\|^2$, respectively. Now, for each $\tau \in [t_k, t_{k+1} - 1]$, we have

$$\begin{aligned} \mathbb{E}\Big[\Delta V(\mathbf{X}(\tau)) \Big| \mathbf{X}(\tau) \Big] &\leq \|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\sigma}\|^2 + n - 2\epsilon\langle \mathbf{q}(\tau), \boldsymbol{\nu}\rangle + 2\langle \mathbf{q}(\tau), \boldsymbol{\nu} - \mathbf{S}^*(\tau)\rangle \\ &\quad + 2\langle \mathbf{q}(\tau), \mathbf{S}^*(\tau) - \mathbf{S}(\tau)\rangle + 2\langle \mathbf{q}(\tau), \mathbf{s}(\tau)\rangle \mathbb{1}_{\{r(\tau)>0\}} \\ &\leq \|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\sigma}\|^2 + n - 2\epsilon\langle \mathbf{q}(\tau), \boldsymbol{\nu}\rangle - 2\nu_{\min}\|\mathbf{q}_\perp(\tau)\| \\ &\quad + 2n(a_{\max} + 1)\tau + 2\langle \mathbf{q}(\tau), \mathbf{s}(\tau)\rangle \mathbb{1}_{\{r(\tau)>0\}} \end{aligned}$$

and

$$\mathbb{E}\Big[\Delta V_\|(\mathbf{X}(\tau)) \Big| \mathbf{X}(\tau) \Big] \geq -2\epsilon\langle \mathbf{q}_\|(\tau), \boldsymbol{\nu}\rangle + 2\langle \mathbf{q}_\|(\tau), \mathbf{s}(\tau)\rangle.$$

We then have the expected drift of $Z(\mathbf{X})$ at steady state given by

$$\begin{aligned} &\mathbb{E}\Big[\|\mathbf{q}_\perp(t_{k+1})\| - \|\mathbf{q}(t_k)\| \Big] \\ &\quad \leq \mathbb{E}\Big[\sum_{\tau=t_k}^{t_{k+1}-1} \Big(\frac{\|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\sigma}\|^2 + n}{2\|\mathbf{q}(\tau)\|} - \epsilon \Big\langle \frac{\mathbf{q}_\perp(\tau)}{\|\mathbf{q}_\perp(\tau)\|}, \boldsymbol{\nu} \Big\rangle - \nu_{\min} + \frac{n(a_{\max} + 1)\tau}{\|\mathbf{q}(\tau)\|} \\ &\qquad + 2\Big\langle \frac{\mathbf{q}_\perp(\tau)}{\|\mathbf{q}_\perp(\tau)\|}, \mathbf{s}(\tau) \Big\rangle \mathbb{1}_{\{r(\tau)>0\}} \Big) \Big] \\ &\quad \leq T\Big(\frac{\|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\sigma}\|^2 + n + n(a_{\max} + 1)T}{2(\|\mathbf{q}(t_k)\| - nT)} + \epsilon\|^\circ\| - \nu_{\min} \Big) + \sqrt{n}\Delta_r \\ &\quad \leq T\Big(\frac{\|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\sigma}\|^2 + n + n(a_{\max} + 1)T}{2(\|\mathbf{q}(t_k)\| - nT)} - \frac{\nu_{\min}}{4} \Big), \end{aligned}$$

where the last inequality follows from $\epsilon \leq \frac{\nu_{\min}}{4\|\boldsymbol{\nu}\|}$ and $T > \frac{\Delta_r}{\epsilon} \geq \frac{4\|\boldsymbol{\nu}\|\Delta_r}{\nu_{\min}} \geq \frac{4\sqrt{n}\Delta_r}{\nu_{\min}}$.

Let $\kappa = \frac{2(\|\boldsymbol{\lambda}\|^2+\|\boldsymbol{\sigma}\|^2+n)}{\nu_{\min}} + nT(\frac{2(a_{\max}+1)}{\nu_{\min}} + 1)$. We have that $\|\mathbf{q}(t_k)\| > \kappa$ implies

$$\mathbb{E}\Big[\|\mathbf{q}_\perp(t_{k+1})\| - \|\mathbf{q}(t_k)\|\Big] \leq -\frac{\nu_{\min}}{4}T,$$

then using [15, Lemma 3] with $D = na_{\max}T$ and $\eta = \frac{\nu_{\min}T}{4}$, we have

$$\begin{aligned}\mathbb{E}\Big[\|\mathbf{q}_\perp(t_k)\|^2\Big] &\leq 4\kappa^2 + 32D^2\Big(1+\frac{D}{\eta}\Big)^2 \leq \Big(2\kappa + 4\sqrt{2}D(1+\frac{D}{\eta})\Big)^2 \\ &\leq T^2\Big(\frac{4na_{\max} + 4(\|\boldsymbol{\lambda}\|^2 + \|\boldsymbol{\sigma}\|^2 + 2n) + 16\sqrt{2}n^2a_{\max}^2}{\nu_{\min}} + 2n + 4\sqrt{2}na_{\max}\Big)^2.\end{aligned}$$

Letting $M = \frac{4na_{\max}+4(\|\boldsymbol{\lambda}\|^2+\|\boldsymbol{\sigma}\|^2+2n)+16\sqrt{2}n^2a_{\max}^2}{\nu_{\min}} + 2n + 4\sqrt{2}na_{\max}$, we have $\mathbb{E}\Big[\|\mathbf{q}_\perp(t_k)\|^2\Big] \leq T^2M^2$.

Consider the Lyapunov function $W(\mathbf{X}) = \sum_i\Big(\sum_j q_{ij}\Big)^2 + \sum_i\Big(\sum_i q_{ij}\Big)^2 - \frac{1}{n}\Big(\sum_{ij} q_{ij}\Big)^2$, and set the corresponding Lyapunov drift at steady state to zero, $\mathbb{E}\Big[W(\mathbf{X}(t_{k+1})) - W(\mathbf{X}(t_k))\Big] = 0$. We then have $T_1 = T_2 + T_3 + T_4$, where

$$\begin{aligned}T_1 = 2\mathbb{E}\Big[&\sum_i\Big(\sum_j q_{ij}(t_k)\Big)\Big(\sum_j\sum_{\tau=t_k}^{t_{k+1}-1}(s_{ij}(\tau)\mathbb{1}_{r(\tau)=0} - a_{ij}(\tau))\Big) \\ &+ \sum_j\Big(\sum_i q_{ij}(t_k)\Big)\Big(\sum_i\sum_{\tau=t_k}^{t_{k+1}-1}(s_{ij}(\tau)\mathbb{1}_{r(\tau)=0} - a_{ij}(\tau))\Big) \\ &- \frac{1}{n}\Big(\sum_{ij} q_{ij}(t_k)\Big)\Big(\sum_{ij}\sum_{\tau=t_k}^{t_{k+1}-1}(s_{ij}(\tau)\mathbb{1}_{r(\tau)=0} - a_{ij}(\tau))\Big)\Big] \\ =\ & 2(\epsilon T - \Delta_r)\mathbb{E}\Big[\sum_{ij} q_{ij}(t_k)\Big],\end{aligned}$$

$$\begin{aligned}T_2 = \mathbb{E}\Big[&\sum_i\Big(\sum_j\sum_{\tau=t_k}^{t_{k+1}-1}(a_{ij}(\tau) - s_{ij}(\tau)\mathbb{1}_{\{r(\tau)=0\}})\Big)^2 \\ &+ \sum_j\Big(\sum_i\sum_{\tau=t_k}^{t_{k+1}-1}(a_{ij}(\tau) - s_{ij}(\tau)\mathbb{1}_{\{r(\tau)=0\}})\Big)^2 \\ &- \frac{1}{n}\Big(\sum_{ij}\sum_{\tau=t_k}^{t_{k+1}-1}(a_{ij}(\tau) - s_{ij}(\tau)\mathbb{1}_{\{r(\tau)=0\}})\Big)^2\Big] \\ =\ & \Big(2-\frac{1}{n}\Big)T\sum_{ij}\sigma_{ij}^2 + n(\epsilon T - \Delta_r)^2,\end{aligned}$$

$$T_3 = \mathbb{E}\Big[-\sum_i\Big(\sum_j\sum_{\tau=t_k}^{t_{k+1}-1}u_{ij}(\tau)\Big)^2 - \sum_j\Big(\sum_i\sum_{\tau=t_k}^{t_{k+1}-1}u_{ij}(\tau)\Big)^2 + \frac{1}{n}\Big(\sum_{ij}\sum_{\tau=t_k}^{t_{k+1}-1}u_{ij}(\tau)\Big)^2\Big]$$

$$\leq nT(\epsilon T - \Delta_r),$$

$$T_4 = 2\mathbb{E}\Big[\sum_i \Big(\sum_j q_{ij}(t_{k+1})\Big)\Big(\sum_j \sum_{\tau=t_k}^{t_{k+1}-1} u_{ij}(\tau)\Big) + \sum_j \Big(\sum_i q_{ij}(t_{k+1})\Big)\Big(\sum_i \sum_{\tau=t_k}^{t_{k+1}-1} u_{ij}(\tau)\Big)$$

$$- \frac{1}{n}\Big(\sum_{ij} q_{ij}(t_{k+1})\Big)\Big(\sum_{ij} \sum_{\tau=t_k}^{t_{k+1}-1} u_{ij}(\tau)\Big)\Big]$$

$$\leq 2na_{\max} T\mathbb{E}\Big[\sum_{ij} \sum_{\tau=t_k}^{t_{k+1}-1} u_{ij}(\tau)\Big] + 4n \sum_{\tau=t_k}^{t_{k+1}-1} \mathbb{E}\Big[\sum_{ij} u_{ij}(\tau)\Big]\sqrt{\mathbb{E}[\|\mathbf{q}_\perp(t_k+\tau)\|^2]}$$

$$\leq 2n^2(\epsilon T - \Delta_r)T(a_{\max} + 2M).$$

Hence, we have the upper bound

$$\mathbb{E}\Big[\sum_{ij} \bar{q}_{ij}\Big] \leq \Big(1 - \frac{1}{2n}\Big)\frac{T}{\epsilon T - \Delta_r}\|\boldsymbol{\sigma}\|^2 + \frac{n((1+\epsilon)T - \Delta_r)}{2} + n^2 T(a_{\max} + 2M)$$

$$\leq \left(1 - \frac{1}{2n}\right)\frac{T}{\epsilon T - \Delta_r}\|\tilde{\boldsymbol{\sigma}}\|^2 + T\Big(\frac{n(1+\epsilon)}{2} + n^2(a_{\max} + 2M)\Big).$$

□

References

1. Bertsimas, D., Gamarnik, D., Tsitsiklis, J.N.: Performance of multiclass Markovian queueing networks via piecewise linear Lyapunov functions. Ann. Appl. Probab. **11**, 1384–1428 (2001)
2. Bramson, M.: State space collapse with application to heavy traffic limits for multiclass queueing networks. Queueing Syst. **30**(1), 89–140 (1998). https://doi.org/10.1023/A:1019160803783
3. Çelik, G., Borst, S., Whiting, P., Modiano, E.: Variable frame based max-weight algorithms for networks with switchover delay. In: 2011 IEEE International Symposium on Information Theory Proceedings, pp. 2537–2541 (2011)
4. Celik, G., Borst, S.C., Whiting, P.A., Modiano, E.: Dynamic scheduling with reconfiguration delays. Queueing Syst. Theory Appl. **83**(1–2), 87–129 (2016). https://doi.org/10.1007/s11134-016-9471-4
5. Chan, C.W., Armony, M., Bambos, N.: Maximum weight matching with hysteresis in overloaded queues with setups. Queueing Syst. Theory Appl. **82**(3–4), 315–351 (2016). https://doi.org/10.1007/s11134-015-9468-4
6. De Dobbelaere, P., Falta, K., Gloeckner, S., Patra, S.: Digital mems for optical switching. IEEE Commun. Mag. **40**(3), 88–95 (2002). https://doi.org/10.1109/35.989763
7. Eryilmaz, A., Srikant, R.: Asymptotically tight steady-state queue length bounds implied by drift conditions. Queueing Syst. Theory Appl. **72**(3–4), 311–359 (2012). https://doi.org/10.1007/s11134-012-9305-y
8. Kang, W.N., Kelly, F.P., Lee, N.H., Williams, R.J.: State space collapse and diffusion approximation for a network operating under a fair bandwidth sharing policy. Ann. Appl. Probab. **19**(5), 1719–1780 (2009)
9. Keyworth, B.: Roadm subsystems and technologies. In: Optical Fiber Communication Conference, 2005. Technical Digest. OFC/NFOEC, vol. 3, pp. 3–4 (2005). https://doi.org/10.1109/OFC.2005.192706
10. Li, X., Hamdi, M.: On scheduling optical packet switches with reconfiguration delay. IEEE J. Sel. Areas Commun. **21**(7), 1156–1164 (2003). https://doi.org/10.1109/JSAC.2003.815843
11. Li, Y., Panwar, S., Chao, H.: Frame-based matching algorithms for optical switches. In: 2003, HPSR. Workshop on High Performance Switching and Routing, pp. 97–102 (2003). https://doi.org/10.1109/HPSR.2003.1226687

12. Liu, H., Lu, F., Forencich, A., Kapoor, R., Tewari, M., Voelker, G.M., Papen, G., Snoeren, A.C., Porter, G.: Circuit switching under the radar with reactor. In: Proceedings of the 11th USENIX Conference on Networked Systems Design and Implementation, NSDI'14, pp. 1–15. USENIX Association, Berkeley, CA, USA (2014). http://dl.acm.org/citation.cfm?id=2616448.2616450
13. Liu, H., Mukerjee, M.K., Li, C., Feltman, N., Papen, G., Savage, S., Seshan, S., Voelker, G.M., Andersen, D.G., Kaminsky, M., Porter, G., Snoeren, A.C.: Scheduling techniques for hybrid circuit/packet networks. In: ACM CoNEXT, ACM CoNEXT'15. ACM, New York (2015)
14. Maguluri, S.T., Burle, S.K., Srikant, R.: Optimal heavy-traffic queue length scaling in an incompletely saturated switch. SIGMETRICS Perform. Eval. Rev. **44**(1), 13–24 (2016). https://doi.org/10.1145/2964791.2901466
15. Maguluri, S.T., Srikant, R.: Heavy traffic queue length behavior in a switch under the maxweight algorithm. Stoch. Syst. **6**(1), 211–250 (2016). https://doi.org/10.1214/15-SSY193
16. McKeown, N., Anantharam, V., Walrand, J.: Achieving 100% throughput in an input-queued switch. In: INFOCOM'96. Fifteenth Annual Joint Conference of the IEEE Computer Societies. Networking the Next Generation. Proceedings IEEE, vol. 1, pp. 296–302 (1996). https://doi.org/10.1109/INFCOM.1996.497906
17. Tassiulas, L., Ephremides, A.: Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. IEEE Trans. Autom. Control **37**(12), 1936–1948 (1992). https://doi.org/10.1109/9.182479
18. Venkatakrishnan, S.B., Alizadeh, M., Viswanath, P.: Costly circuits, submodular schedules and approximate carathéodory theorems. In: Proceedings of the 13th ACM SIGMETRICS/PERFORMANCE Joint International Conference on Measurement and Modeling of Computer Systems, SIGMETRICS'16. ACM, New York (2016)
19. Wang, C., Javidi, T.: Adaptive policies for scheduling with reconfiguration delay: an end-to-end solution for all-optical data centers. IEEE/ACM Trans. Netw. **25**(3), 1555–1568 (2017). https://doi.org/10.1109/TNET.2016.2644617
20. Wang, C.H., Javidi, T., Porter, G.: End-to-end scheduling for all-optical data centers. In: INFOCOM, 2015 Proceedings IEEE (2015)
21. Wang, W., Maguluri, S.T., Srikant, R., Ying, L.: Heavy-traffic delay insensitivity in connection-level models of data transfer with proportionally fair bandwidth sharing. SIGMETRICS Perform. Eval. Rev. **45**(3), 232–245 (2018). https://doi.org/10.1145/3199524.3199565