Queueing Syst (2016) 83:87–129
DOI 10.1007/s11134-016-9471-4

Dynamic scheduling with reconfiguration delays

G. Celik[1,2] · **S. C. Borst**[3] · **P. A. Whiting**[4] · **E. Modiano**[5]

Received: 20 November 2014 / Revised: 12 December 2015 / Published online: 23 January 2016

Abstract We consider scheduling in networks with interference constraints and reconfiguration delays, which may be incurred when one service schedule is dropped and a distinct service schedule is adopted. Reconfiguration delays occur in a variety of communication settings, such as satellite, optical, or delay-tolerant networks. In the absence of reconfiguration delays it is well known that the celebrated Max-Weight scheduling algorithm guarantees throughput optimality without requiring any knowledge of arrival rates. As we will show, however, the Max-Weight algorithm may fail to achieve throughput optimality in case of nonzero reconfiguration delays. Motivated by the latter issue, we propose a class of adaptive scheduling algorithms which persist with the current schedule until a certain stopping criterion is reached, before switching to the next schedule. While earlier proposed Variable Frame-Based Max-Weight (VFMW) policies belong to this class, we also present Switching-Curve-Based (SCB) policies that are more adaptive to bursts in arrivals. We develop novel Lyapunov drift

✉ P. A. Whiting
philip.whiting@mq.edu.au

G. Celik
guner@alum.mit.edu

S. C. Borst
sem@research.bell-labs.com

E. Modiano
modiano@mit.edu

1 LIDS, MIT, Cambridge, MA 02139, USA

2 400 Oracle Parkway, Rm. 1385, Redwood Shores, CA 94065, USA

3 Bell Labs, 600–700 Mountain Avenue, Murray Hill, NJ 07974, USA

4 Department of Engineering, Macquarie University, North Ryde, NSW 2109, Australia

5 MIT, Room 33-412, 77 Massachusetts Avenue, Cambridge, MA 02139, USA

techniques to prove that this class of algorithms under certain conditions achieves throughput optimality by dynamically adapting the durations of the interswitching intervals. Numerical results demonstrate that these algorithms significantly outperform the ordinary Max-Weight algorithm, and that SCB policies yield a better delay performance than VFMW policies.

Keywords Max-Weight algorithm · Scheduling · Reconfiguration · Switching delay · Stability · Throughput optimality

Mathematics Subject Classification Primary: 68M20 (Performance evaluation; queueing; scheduling) · Secondary: 60K25 (Queueing theory) · 90B18 (Communication networks) · 90B22 (Queues and service)

1 Introduction

Dynamic scheduling of servers in stochastic networks with interference constraints has been a very active field of research for over two decades (for example, [12,19,22,23,27,30,31,37]). Perhaps the most familiar scenario in which such interference constraints are manifested is that of traffic light controlled two-way intersections [6,35]. Here, the interference constraints reflect the necessary switching delays which lead to some loss of throughput in the case of fixed cycle controls [6], and also for the adaptive control considered in [35].

In communications, the significant effects of server switchover delays or the time durations to reconfigure schedules have been largely ignored, even though such reconfiguration delays commonly arise in a variety of settings [1,4,18,38]. In satellite networks where multiple mechanically steered antennas are servicing ground stations, the time to rotate from one station to another can be around 10 ms [4,33], which is on the same order as the time to transmit multiple packets. Similarly, in optical communication systems, laser tuning delay for different lightpaths at transceivers can take from tens of μs to ms, which, at data rates of giga-bits-per-second, corresponds to the transmission time of hundreds of packets [7,18]. Large switchover delays also occur in delay-tolerant networks where mobile servers (for example, robots) are used as data gatherers from sensors in a field [24].

The effects of switchover delays have been extensively analyzed in the literature on polling systems [5,28,36]. However, in polling systems it is typically assumed that there is just a single server which can only serve one queue at a time. In contrast, the above-mentioned stochastic networks with interference constraints usually involve several servers and various subsets of queues that can be served simultaneously. This latter scenario was considered in the seminal work of Tassiulas and Ephremides [30,31] that characterized the stability region and proved throughput optimality of the Max-Weight scheduling algorithm. These results were later extended to joint power allocation and routing in wireless networks in [22,23] as well as optimal scheduling for packet switches in [26,27]. More recently, various distributed scheduling algorithms have been developed that achieve throughput optimality [12,19]. We refer the reader to [15,20] for a detailed review of scheduling in wireless networks. However, these papers do not consider the server switchover delays and, in fact, the algorithms proposed in these papers fail to achieve throughput optimality under nonzero switching delays, as we demonstrate in Lemma 1. The intuition behind this observation is that the Max-Weight scheduling algorithm tends to switch between configurations often (i.e., as soon as another schedule has a greater weight); thus, the fraction of time spent on reconfigurations does not diminish as the total queue length grows large, leading to instability.

Of particular relevance to our work are the papers [2,7,16,29,32], which consider perturbations such as switchover delays in a network setting. However, the system model and the assumptions in [7,29,32] are substantially different from those in the present paper. In particular, [7] assumes knowledge of the arrival rate vectors which significantly simplifies the scheduling problem, [29] considers a deterministic setting where servers do not interfere, and [32] proposes a policy similar to one of the policies proposed in the present paper, for a system consisting of a single-server serving multiple queues with asynchronous transmission opportunities in the absence of switching delays. The work in [2] addresses the problem of switchover delays by aggregating jobs into batches. Batching is driven through a sequence of load estimates; see [2] Proposition 4.2, page 239. The authors of [16] propose a dynamic cone-based scheduling algorithm for a network model similar to the one in this paper. While the proposed algorithm may achieve good throughput performance in practice, the proof arguments in [16] do not seem entirely sound however, as we discuss in greater detail in Sect. 5. Finally, for a simple two-queue system, [8] showed that the simultaneous presence of randomly varying connectivity and switchover delays reduces the system stability region significantly, whereas in the absence of randomly varying connectivity, switchover delays do not reduce the stability region [7,9].

The fundamental limitation of the Max-Weight scheduling algorithm and its variations under nonzero switchover delays suggests that instead of reconfiguring the system as soon as another schedule with larger weight is found, one should persist with the current schedule for a certain period of time, in order to avoid spending too much time on reconfiguration, which is also the main idea behind the dynamic cone-based policy proposed in [16]. By letting the expected amount of time between reconfigurations grow with the total queue length, the fraction of time spent on reconfiguration is reduced when the queues get large, leading to throughput optimality. Our main theorem, see Sect. 3, establishes conditions on the amount of time between reconfigurations to guarantee throughput optimality. These conditions include lower bounds on the expected time to stop (not necessarily a reconfiguration time) and an upper bound on its second moment. Both of these bounds are expressed in terms of a sublinear function.

We propose three classes of algorithms and prove their throughput optimality by showing that they indeed satisfy the sufficient conditions mentioned above. The first class of algorithms, termed the *Variable Frame-Based Max-Weight* (VFMW) algorithms, persist with a Max-Weight schedule during a frame whose duration is calculated at the beginning of the frame, as a sublinear function of the queue sizes. The VFMW algorithms were first introduced and analyzed in [9]. In this paper we show that the throughput optimality of the VFMW algorithms follows directly from the sufficient stability conditions mentioned above. We also consider two classes which are more responsive to queue dynamics—the Bias-Based (BB) and Switching-Curve-Based (SCB) algorithms. BB algorithms make reconfiguration decisions based on the event that there is a sufficiently large change in queue lengths as compared to the queue sizes at the time of the previous reconfiguration event. SCB algorithms are a natural extension of the Max-Weight scheme, and make reconfiguration decisions when the weight of an alternative schedule exceeds that of the current schedule by an amount that is a function of the current queue lengths.

The approach we take is to work with a randomly stopped Lyapunov drift function. This is not straightforward to treat because the stopped sum and the stopping time are usually dependent and first and second moments of the stopped sum must be considered. Bounds for stopped queues are also needed in establishing the sufficient criteria for the BB and SCB algorithms. Additionally, to establish throughput optimality for SCB we show that if a policy has a negative drift throughout a service interval, and if another policy with an earlier switching decision is throughput-optimal, then the former policy is also throughput-optimal. Subsequently, lower bounds for mean delay are developed which may be of additional interest.

The remainder of the paper is organized as follows. In Sect. 2 we introduce the system model, provide technical preliminaries, and show the instability of the ordinary Max-Weight algorithm. Section 3 contains our main result, Theorem 1, which provides conditions on the amount of time between reconfiguration to guarantee stability. Section 4 shows that the three classes of algorithms described above all meet the criteria of Theorem 1. In Sect. 5, we discuss performance considerations of the three proposed algorithms and the analysis of [16]. We present numerical experiments in Sect. 6, and briefly discuss parameter choices based on delay results for the proposed algorithms. In Sect. 7 we make concluding remarks and offer suggestions for further research.

2 Model description and preliminaries

This section provides a model description followed by a brief example, after which the instability of the ordinary Max-Weight algorithm for scheduled networks with switching delay is demonstrated. Following this, a class of generalized Max-Weight policies is defined. These are exemplified and analyzed in the following sections.

2.1 Model description

We consider a discrete-time (slotted) system of N parallel infinite-buffer queues labeled by the set $\mathcal{N} = \{1, 2, \ldots, N\}$. The time slot indices take values in the set $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$ of nonnegative integers.

The vector $\mathbf{A}(t) \in \mathbb{N}_0^N$ denotes the numbers of arriving packets at the various queues at the end of time slot t. (Here and throughout the paper, we represent vectors by bold letters.) The sequences of arrivals to the various queues are mutually independent, and $A_\ell(1), A_\ell(2), \ldots$ are i.i.d. copies of a random variable A_ℓ, with $\mathbb{P}\{A_\ell = 0\} > 0$,

Table 1 State variable updates in slot t

(1)	If $\phi(t) > 0$, set $Q'_\ell(t) = Q_\ell(t), \forall \ell \in \mathcal{N}$, then proceed to step 3)
(2)	A decision is taken whether or not to invoke a reconfiguration, based on the process history
(a)	If a reconfiguration takes place to service vector $J \in \mathcal{I}$, then set $\mathbf{I}(t) = J, \phi(t) = T_r$ and $Q'_\ell(t) = Q_\ell(t), \forall \ell \in \mathcal{N}$ and proceed to step 3)
(b)	If no reconfiguration takes place, then apply the service vector $\mathbf{I}(t)$ to the queue length vector $\mathbf{Q}(t)$, i.e., set $Q'_\ell(t) = [Q_\ell(t) - I_\ell(t)]^+ \; \forall \ell \in \mathcal{N}$, and proceed to step 3)
(3)	Set $\mathbf{I}(t+1) = \mathbf{I}(t)$, $\mathbf{Q}(t+1) = \mathbf{Q}'(t) + \mathbf{A}(t)$, and $\phi(t+1) = [\phi(t) - 1]^+$.

$\mathbb{P}\{A_\ell > a\} > 0$ for all $a \in \mathbb{N}$ (unbounded support), $\mathbb{E}[A_\ell] = \lambda_\ell$, and $\mathbb{E}\left[A_\ell^2\right] \leq A_{\max}^2 < \infty$ (finite second moment) for all $\ell \in \mathcal{N}$.

Let $\mathbf{Q}(t) = (Q_1(t), \ldots, Q_N(t)) \in \mathbb{N}_0^N$ denote the lengths of the various queues at the beginning of time slot t. The queues are initially all assumed to be empty, i.e., $\mathbf{Q}(0) = \mathbf{0}$. Let $\mathbf{I}(t) = (I_1(t), \ldots, I_N(t)) \in \mathcal{I}$ be the applicable service vector in time slot t, where $\mathcal{I} \subseteq \mathbb{N}_0^N$ is the set of all feasible service vectors. For convenience, we assume that $\mathcal{I}$ includes the all-zero vector, and also suppose that $\mathcal{I}$ is finite, implying that there exists a $\mu_{\max} \in \mathbb{N}_0$ such that $I_\ell \leq \mu_{\max}$ for all $\mathbf{I} \in \mathcal{I}$, $\ell \in \mathcal{N}$.

During each time slot, the system may either be in service mode or in reconfiguration mode. A reconfiguration involves a fixed duration of T_r time slots, and allows a switch from one service vector to another. During reconfiguration, a counter $\phi(t)$ is used to count down the number of slots that remain until service can continue. On the other hand, if the system is not being reconfigured or switched during time slot t, but is in service mode, then service proceeds according to the vector $\mathbf{I}(t)$, meaning that $\min\{I_\ell(t), Q_\ell(t)\}$ packets depart from queue ℓ at the end of the time slot. The lengths of the various queues then evolve according to the usual Lindley recursion

$$Q_\ell(t+1) = [Q_\ell(t) - I_\ell(t)]^+ + A_\ell(t), \qquad \forall \ell \in \mathcal{N}, \tag{1}$$

where $[x]^+ \doteq \max\{x, 0\}$.

We now summarize the sequence of events which occur in time slot t. Table 1 indicates the order in which state variables are updated, in chronological order within slot t.

We now go through the above steps in detail. If it is determined in step (1) that a reconfiguration period is in progress, no further reconfiguration can take place nor can any service be rendered. At step (2) if no reconfiguration is in progress, then a decision is taken whether to invoke a reconfiguration or not based on the process history. (For some policies, the decision may be based on just the current queue length vector and current service vector. See Definition 1 for the specification of a general class of (nonrandom) scheduling policies and Sect. 2.5 for the class of Generalized Max-Weight policies proposed in the present paper.) If the decision is to reconfigure in case (a), then a counter is set for the duration T_r of the reconfiguration delay. Otherwise, in case (b), packets are removed from the various queues in accordance with the current service vector. Finally, in step (3), the current service vector is propagated as the default service vector for the next time slot, arriving packets are added to the queue lengths, and the reconfiguration counter is decremented by 1, if not already 0. Note that the service vector $\mathbf{I}(t)$ is also defined during a reconfiguration period, even though it does not take effect until the reconfiguration ends and service may resume.

For convenience, we assume that the process always starts with a reconfiguration in time slot 0 as in step (2a). The sequence of slots $(S_k)_{k\in\mathbb{N}_0}$ in which a reconfiguration is invoked are referred to as the reconfiguration times (taken to be ∞ if the kth such does not occur); see also Sect. 2.2. Finally note that it is possible for one reconfiguration to be followed by another back-to-back. In the case $T_r = 1$ there can be a reconfiguration in one slot and in the next slot.

Unless otherwise stated, $||\mathbf{Q}|| = \sum_{\ell=1}^{N} Q_\ell$, $\mathbf{Q} \in \mathbb{N}_0^N$, denotes the $\mathcal{L}_1$ norm. The inner product of two N-dimensional vectors is denoted as $\mathbf{u} \cdot \mathbf{v} = \sum_{\ell=1}^{N} u_\ell v_\ell$. We denote a sequence of scalars or vectors up to time t using a superscript, for example $\mathbf{A}^{(t)} = (\mathbf{A}(0), \ldots, \mathbf{A}(t))$.

2.2 Mathematical preliminaries

Only the arrivals are random and the sample points lie in the set of sequences $\Omega = \prod_{\mathbb{N}_0} \mathbb{N}_0^N$. A filtration (stochastic basis) is obtained via the cylinder σ-algebras

$$\mathcal{F}_t = \sigma\left(\left\{\prod_{s=0}^{\infty} U_s : \emptyset \neq U_s \subseteq \mathbb{N}_0^N,\ 0 \leq s \leq t, U_s = \mathbb{N}_0^N, s > t\right\}\right),\ t \in \mathbb{N}_0.$$

Additionally, define

$$\mathcal{F} \doteq \sigma\left(\bigcup_{t\in\mathbb{N}_0} \mathcal{F}_t\right).$$

The pre-T σ-algebra for a stopping time S with respect to $\mathcal{F}_t$ is by definition

$$\mathcal{F}_S \doteq \{F \in \mathcal{F} : F \cap \{S \leq t\} \in \mathcal{F}_t, \forall t \in \mathbb{N}_0\},$$

which we denote by $\mathcal{F}_S$. Now let S, T be $\mathcal{F}_t$ stopping times, then it is readily shown that

$$\{\omega : T(\omega) = S(\omega)\} \in \mathcal{F}_T \text{ (and } \mathcal{F}_S). \tag{2}$$

Moreover, since $S+t, t = 0, 1, 2, \ldots$ are themselves a sequence of increasing stopping times, we obtain the S-filtration,

$$\mathcal{F}_S \subseteq \mathcal{F}_{S+1} \subseteq \ldots.$$

A stopping time defined with respect to this filtration is called a $\mathcal{F}_S$ stopping time. It is straightforward to show that if S is a $\mathcal{F}_t$ stopping time and χ is a $\mathcal{F}_S$ stopping time, then $S + \chi$ is a $\mathcal{F}_t$ stopping time.

Finally if $\mathcal{G}$ is a σ-algebra, $X, Y \in m\mathcal{G}$ (see [34, p. 29] for definition) and $A \in \mathcal{G}$, then we say that $X \leq Y$ on occurrence of A almost surely if

$$X\mathbb{1}_A \leq Y\mathbb{1}_A, \quad a.s.$$

An important case is when $X = \mathbb{E}[Z|\mathcal{G}]$ is a version of a conditional expectation of a random variable Z.

Given a (aperiodic, time-homogeneous) Markov chain on a countable state space $\mathcal{A}$, a state $a \in \mathcal{A}$ is ergodic iff $\mathbb{E}_a[\tau_a] < \infty$ where τ_a is the number of steps to return (see [13, p. 321]). The chain is said to be a recurrent chain if it enters an *ergodic state* almost surely. It can readily be shown that such a chain is a Harris chain; see [11, pp. 325–326] for the definition. In the case where there is a *given* ergodic state which the Markov chain enters almost surely, we say the Markov chain is a *1-recurrent Harris chain*. Such chains always have a unique limiting stationary measure which can be deduced from [11, Theorem 4.7, p. 307].

Definition 1 (*Scheduling Policy*) A scheduling policy π is any sequence of deterministic mappings, $\pi \doteq (\pi_t)_{t\in\mathbb{N}}$,

$$\pi_t : \left(\mathbb{N}_0^N\right)^t \times \mathcal{I}^t \times \{0, 1, \ldots, T_r\} \to \mathcal{I} \times \{0, 1, \ldots, T_r\}, \; t \in \mathbb{N}, \tag{3}$$

determining the current service vector $\mathbf{I}(t)$ and reconfiguration state $\phi(t)$. As in Table 1, $\phi(t) = [\phi(t-1) - 1]^+$ or $\phi(t) = T_r$ according to whether a reconfiguration is invoked in slot t. Given the history up to time t, the mapping π_{t+1} determines whether or not a reconfiguration takes place in slot $t+1$,

$$\pi_{t+1}(\mathbf{A}^{(t)}, \mathbf{I}^{(t)}, \phi(t)) = (\mathbf{I}(t+1), \phi(t+1)), \; t \in \mathbb{N}_0,$$

together with any constant map $\pi_0(\cdot) = (\mathbf{I}(0), T_r)$ for slot 0.

The following consistency conditions must be satisfied, for all $t \in \mathbb{N}_0$:

$$(\mathbf{I}(t+1), \phi(t+1)) = (\mathbf{I}(t), \phi(t) - 1), \text{ if } \phi(t) \geq 1; \tag{4}$$

$$\phi(t+1) = T_r, \text{ if } \phi(t) = 0, \mathbf{I}(t) \neq \mathbf{I}(t+1); \tag{5}$$

$$\phi(t+1) \in \{0, T_r\}, \text{ if } \phi(t) = 0, \tag{6}$$

where $\mathbf{I}(t)$ is the applicable service vector in slot t.

Note that the above definition does not rule out mappings which reconfigure at time $t \geq 1$ with $\mathbf{I}(t) = \mathbf{I}(t-1)$. π_0 simply chooses the initial service vector which may be taken arbitrarily, and $\mathbf{Q}(0) = 0$ as already mentioned.

The above definition follows the rubric given in Table 1. Equation (4) is the case where the mapping has elected not to (or cannot) reconfigure in slot $t+1$, so the service vector must be unchanged. Equation (5) is the case where the mapping has elected to reconfigure to a distinct service vector, which can only take place if $\phi(t) = 0$. (6) covers the possibility that the service vector remains the same, with the mapping electing for reconfiguration iff $\phi(t+1) = T_r$.

For any such scheduling policy π, the sequence of reconfiguration times $\{S_k\}_{k\in\mathbb{N}_0}$ satisfy $S_0 = 0$ by definition and obviously (since the scheduling policy is deterministic) $\{\omega : S_k(\omega) \leq t\} \in \mathcal{F}_{t-1}$, $k \geq 1$, $1 \leq t \leq \infty$, so that $\{S_k\}_{k\in\mathbb{N}_0}$ is a sequence of $\mathcal{F}_t$ stopping times.

Definition 2 (*Strong Stability* [15,20,23]) The system is strongly stable under a given scheduling policy π if for any initial state

$$\limsup_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}\left[||\mathbf{Q}^{\pi}(t)||\right] < \infty, \tag{7}$$

where $\mathbf{Q}^{\pi}(t)$ is the random vector of queue lengths obtained under policy π.

We will show the above starting from the empty state but the result is readily seen to hold for arbitrary initial states.

Since the arrival and service variables are integer, this stability criterion implies the existence of a long-run stationary measure with bounded first moments [20].

Definition 3 (*Stability Region* [20,23]) The stability region Λ is the closure of the set of all arrival rate vectors $\lambda = [\lambda_1, \lambda_2, \ldots, \lambda_N]'$ such that there exists a scheduling policy that renders the system strongly stable for λ.

Define $\Lambda_0 = \text{conv}(\mathcal{I})$ as the convex hull of the set of all service vectors, with interior Λ_0^o. Note that the long-term average service rates for the various queues must be a convex combination of the feasible service vectors, i.e., belong to Λ_0. Thus stability cannot possibly be achieved for any arrival rate vector outside $\Lambda_{\mathbf{0}}$, i.e., $\lambda \in \Lambda_0$ is a necessary condition for the system to be strongly stable under any policy, which implies $\Lambda \subseteq \Lambda_0$.

Definition 4 (*Throughput optimality*) A policy is said to be throughput-optimal if it renders the system strongly stable for all arrival rate vectors in the interior of the stability region Λ.

2.3 Example: an optical switch

Figure 1 shows an input-queued optical switch, with 3 input ports and 3 output ports. At each input port $i \in \{1, 2, 3\}$ packets arrive for the 3 output ports, with the mean number of packets arriving for output $o \in \{1, 2, 3\}$ per time slot being $\lambda_{i,o}$. A number of packet arrivals for each input-output pair are i.i.d. over time and mutually independent, over all slots. The switch requires the input ports to each be connected to a distinct output port, in any given time slot, except when it is being reconfigured. For example Fig. 1 shows input ports 1, 2, 3 are currently connected to output ports 3, 1, 2, respectively. In all there are $N = 9$ queues, and the set of nontrivial service vectors corresponds to the set of all possible $6 = 3!$ matchings between the various input and output ports, where one packet is transferred per time slot (if one is present). Hence, the above arrival rates $\lambda_{i,o}$ are in the stability region Λ (see Sect. 2.2) iff $\sum_{o=1}^{3} \lambda_{i,o} \leq 1, \forall i$, $\sum_{i=1}^{3} \lambda_{i,o} \leq 1$, $\forall o$. A new set of connections requires packet transmission to be suspended while the switch is being reconfigured [26] for a period of $T_r > 0$ slots.

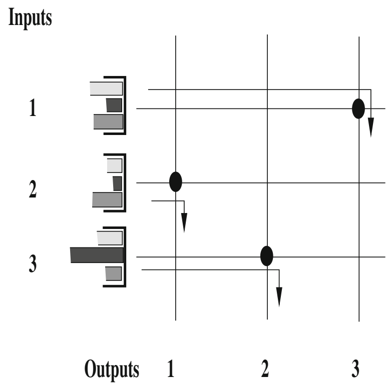

Fig. 1 Optical switch with input to output matching

2.4 Instability of the ordinary Max-Weight policy

We now show that the ordinary Max-Weight policy fails to be throughput-optimal in the case of nonzero reconfiguration delays. Consider a single-server system with $N = 2$ queues and i.i.d. Bernoulli arrivals, i.e., $\mathbb{P}\{A_\ell = 1\} = 1 - \mathbb{P}\{A_\ell = 0\} = p < 1/2$. The set of service vectors is $\mathcal{I} = \{(0,0), (0,1), (1,0)\}$, and the switching delay is $T_r = 1$ time slot.

The stability region of this simple system is given by $p \in [0, 1/2)$. Indeed, the policy of serving each queue to exhaustion, and only then switching, can be shown to render this system stable for any such value of p.

Since the set of service vectors is $\mathcal{I} = \{(0,0), (0,1), (1,0)\}$, the Max-Weight service vector simply corresponds to serving the larger of the two queues. Hence, the ordinary Max-Weight policy reconfigures whenever the queue being served is shorter than the other at step 2) in Table 1. For example if $Q_1(t) < Q_2(t)$ and the current service vector is $(1,0)$, then $(0,1)$ is selected which is used in the following slot (giving one slot where there is no service). Reconfiguration cannot take place in two subsequent slots except initially because there is at most one arrival in each queue per slot. Finally, without loss of generality, the Max-Weight policy sets $I(0) = (1,0)$ initially. It is readily verified that the queue lengths are confined to stay in the set $\mathcal{D} = \{(Q_1, Q_2) : |Q_1 - Q_2| \leq 3\}$. The loss in throughput (defined to be $1 - 2\overline{p}_{MaxWt}$ where $\overline{p}_{MaxWt}$ is the supremum of probabilities p for which the above network is strongly stable for all $q \in (0, p)$) is at least $1/(1 + A)$ where A is a uniform bound on the expected time between switches. The following lemma gives a precise statement.

Lemma 1 *The ordinary Max-Weight algorithm is not throughput-optimal for the above-described two-queue system. Specifically, for all arrival rates* $p > 0.42049$ *both queues grow to infinity almost surely.*

The proof of the above lemma is provided in Appendix 1. As further discussed in the appendix, the expected time between reconfigurations is uniformly bounded over the sum queue length. Starting states well away from the main diagonal are irrelevant since the diagonal will almost surely be crossed at some point, and the queue lengths must inevitably stay in the set $\mathcal{D}$ from then on.

In order to avoid the above pitfall, a scheduling policy must maintain the current service vector for an extended period of time. Moreover, the further the queue state is from the origin, the longer this time should be. In this way the fraction of time spent on switching is gradually reduced as the queues become longer, which is also the key driver for the dynamic cone-based policy proposed in [16]. The question that must be addressed is how to dynamically arrange the durations of the interswitching intervals. In the next sections we will further address this question, and propose a class of adaptive scheduling policies which achieve throughput optimality by defining service intervals based on queue lengths or the evolution of the queues over time.

2.5 Generalized Max-Weight scheduling policies

In the remainder of the paper we consider a particular class of scheduling policies which we refer to as Generalized Max-Weight scheduling policies. These policies differ only in the specific definition of the stopping criterion, but at the kth reconfiguration time, with finite stopping time $S_k < \infty$, it is always the Max-Weight service vector

$$\mathbf{I}(S_k) = \arg\max_{\mathbf{I}\in\mathcal{I}} \mathbf{Q}(S_k)\cdot\mathbf{I}$$

which is selected. Possible ties between several service vectors are assumed to be resolved in a deterministic but otherwise arbitrary manner. It follows that reconfiguration at queue state $\mathbf{Q}$ always results in the same decision under these policies. Following T_r reconfiguration slots, service vector $\mathbf{I}(S_k)$ is invoked unless it is determined at time $S_k + T_r$ that a new service vector should again be adopted and that there should therefore be a further period of reconfiguration.

The reconfiguration times (see text below Definition 1) are defined for these policies according to the following recursion with $S_0 = 0$: if $S_k = \infty$, then $S_{k+1} = \infty$; otherwise,

$$S_{k+1} = \inf\left\{t \geq S_k + T_r :\ f_\pi(\cdot) \geq 0\right\}, k \in \mathbb{N}_0,$$

where $f_\pi(\cdot)$ is a (nonrandom) mapping involving S_k, $\mathbf{Q}(S_k)$, $\mathbf{I}(S_k)$, $\mathbf{Q}(t)$, and t (only through $s = t - S_k$ and depending not at all on k). We refer to $f_\pi(\cdot)$ as the *reconfiguration rule*—determining when it occurs.

As $f_\pi(\cdot)$ is deterministic, it follows that the reconfiguration times $\{S_k\}_{k=0}^{\infty}$ form an increasing sequence of $\mathcal{F}_t$ stopping times. With a slight abuse of notation, define

$$S_k(t,\omega) = \inf\left\{S_k : S_k < t\right\},\ t \geq 1,$$

so that $S_k(1,\omega) = S_0 = 0$ for example. The process defined with respect to the sequence of states with t

$$(t, \mathbf{Q}(t,\omega), \mathbf{I}(t,\omega), \mathbf{Q}(S_k(t,\omega),\omega), S_k(t,\omega))$$

as the time variable, is, by the definition of $f_\pi(\cdot)$, a Markov chain. Among these processes we will only consider those where $f_\pi(\cdot)$ depends on time only as time elapsed since the last reconfiguration. Furthermore we will only consider those where there is an auxiliary state variable $\varrho(t) \in \mathbb{N}_0$ such that

$$(\mathbf{Q}(t), \mathbf{I}(t), \mathbf{Q}(S_k(t, \omega)), \varrho(t)) \tag{8}$$

is a Markov chain with time index t and where $\varrho(t) \in \mathbb{N}_0$ depends only on time elapsed since the last reconfiguration, or to the next reconfiguration (when known in advance). Wherever the above is the case, we will refer to this process as a f_π-chain (defined on $(\Omega, \mathcal{F})$ and induced by $f_\pi(\cdot)$).

It is easy to see that the above stochastic process determined by $f_\pi(\cdot)$ and restricted to the stopping times S_k forms a time-homogeneous Markov chain, with state variable $(\mathbf{Q}(S_k), \mathbf{I}(S_k))$. To deal with the possibility that there are only a finite number of steps, we will suppose that the current event is absorbing with probability equal to the chance that the next reconfiguration occurs at $S_{k+1} = \infty$.

It is clear that more general scheduling policies than those described above can be contemplated. However, constructions involving only the most recent reconfiguration and independent of both the time slot t and of the number of switches k are natural to consider. In particular, there are scheduling policies π of this type which achieve throughput optimality.

3 Throughput-optimal policies and main results

The main purpose of this section is the presentation and proof of Theorem 1. Before doing so, we describe three policies all of which, as we will show, satisfy the conditions of Theorem 1, and then some necessary lemmas will be presented. We begin with the following definition.

Definition 5 A strictly positive increasing function $F(\cdot)$ with a uniformly bounded continuous derivative is called *sublinear* if it satisfies the following two conditions:

$$\lim_{y\to\infty} F(y) = \infty,$$
$$\lim_{y\to\infty} \frac{F(y)}{y} = 0.$$

It is easy to see that any strictly positive increasing concave function with continuous derivatives in $(0, \infty)$ such that $F'(y) \to 0$ as $y \to \infty$ and $F(y) \uparrow \infty$ as $y \uparrow \infty$ is sublinear, for example, $F(y) = 1 + \sqrt{y}$. In the following section the function $F(\cdot)$ is any sublinear function as just defined.

3.1 Example policies

3.1.1 Variable Frame-Based Max-Weight (VFMW) policy

In the Variable Frame-Based Max-Weight (VFMW) policy introduced in [9], the next time to reconfigure following reconfiguration at time S_k is determined in advance as

$$f_\pi = t - F\left(||\mathbf{Q}(S_k)||\right) - T_r - S_k = s - F\left(||\mathbf{Q}(S_k)||\right) - T_r,\ t \geq S_k,$$

where $s \doteq t - S_k$ is the number of slots since the most recent reconfiguration. Reconfiguration will occur at $S_{k+1} = \lceil F\left(||\mathbf{Q}(S_k)||\right)\rceil + T_r + S_k$, and the service interval is referred to as a frame, from whence the name of this policy is derived.

Observe that the process determined at slot t by the state variable

$$(\mathbf{Q}(t), \mathbf{I}(t), \upsilon(t), \phi(t)) \tag{9}$$

is a Markov chain. Here $\varrho(t) := \upsilon(t) = S_{k+1} - S_k - s \in \mathbb{N}_0$ is the time-to-go variable. Clearly $\upsilon(t)$ is decremented as the frame proceeds, and a reconfiguration is invoked when it becomes 0. The variable $\phi(t) \in \{0, \ldots, T_r\}$ counts the slots during a reconfiguration; see Definition 1. It is easy to see that each state (υ, ϕ) can be injectively mapped into $\mathbb{N}_0$, and so the VFMW policy gives rise to a f_π-chain.

3.1.2 Bias-Based (BB) policy

The Bias-Based (BB) policy makes a reconfiguration decision whenever there is a sufficiently large change in the queue length, hence the name. Given a constant $\theta_{BB} > 0$ and setting $s = t - S_k$, $t \geq S_k$,

$$f_\pi = ||\mathbf{Q}(s + S_k) - \mathbf{Q}(S_k)|| - \theta_{BB} F\left(||\mathbf{Q}(S_k)||\right).$$

Unlike the VFMW policy, the next stopping time is not known in advance. However, it is the case that $S_k < \infty$ for all $k \in \mathbb{N}_0$. To see this we may suppose that a reconfiguration does not occur at $S_k + T_r$, which implies

$$||\mathbf{Q}(S_k + T_r) - \mathbf{Q}(S_k)|| < \theta_{BB} F\left(||\mathbf{Q}(S_k)||\right).$$

But then in any slot $t \geq S_k + T_r$, with strictly positive probability, there will be a sufficiently large number of arrivals $||\mathbf{A}(t)||$ so that

$$||\mathbf{Q}(t+1)|| \geq ||\mathbf{A}(t)|| > \theta_{BB} F\left(||\mathbf{Q}(S_k)||\right) + ||\mathbf{Q}(S_k)||, \tag{10}$$

as the support of the distribution of the number of arrivals is unbounded. But (10) implies that $f_\pi > 0$ in slot $t + 1$. Hence $S_{k+1} - S_k < \infty$ almost surely for all $k \geq 0$.

Under the BB policy, the process with state variable

$$(\mathbf{Q}(t), \mathbf{I}(t), \mathbf{Q}(S_k(t, \omega)), \varrho(t))$$

is a Markov chain, where $\varrho(t) := \phi(t)$ is the reconfiguration variable as above. Hence the BB policy again gives rise to a f_π-chain; see (8).

3.1.3 Switching-Curve-Based (SCB) policy

The SCB policy reconfigures whenever the weight of a given service vector $\mathbf{I}^*$, necessarily the current Max-Weight vector, exceeds that of the current service vector, $\mathbf{I}(t)$, by an amount determined by the current queue state

$$f_\pi = \max_{\mathbf{I}\in\mathcal{I}} \mathbf{I} \cdot \mathbf{Q}(t) - \mathbf{I}(s + S_k) \cdot \mathbf{Q}(S_k + s) - F(||\mathbf{Q}(S_k + s)||),$$

where $s = t - S_k$, $t \geq S_k$, as before. If there are only two queues, the second term can be thought of as a curve which must be crossed in order for the process to reconfigure, hence the name. Unlike the BB policy, only finitely many reconfigurations may occur.

The process with state variable

$$(\mathbf{Q}(t), \mathbf{I}(t), \varrho(t))$$

is a Markov chain, with $\varrho(t) := \phi(t)$ once again indicating a reconfiguration step as with the VFMW and BB policies. Hence the SCB policy also gives rise to a f_π-chain.

3.2 Main theorem

We now establish sufficient generic conditions for the class of Generalized Max-Weight policies to achieve throughput optimality.

As already mentioned, the reconfiguration rule $f_\pi(\cdot)$ depends on time t only as time elapsed since the previous reconfiguration. It thus gives rise to a time-homogeneous Markov chain with state as given in (8).

The proof of Theorem 1, to be stated in a moment, relies on drift calculations for the Lyapunov function

$$L(\mathbf{Q}) \doteq \sum_{\ell=1}^{N} Q_\ell^2. \tag{11}$$

Drift analysis for a quadratic Lyapunov function is a classical approach in establishing stability of Max-Weight policies. However, in contrast to the usual set-up, we cannot hope to obtain negative drift on a slot-by-slot basis because of the forced inactivity during reconfiguration. In order to show negative drift, we therefore need to consider the evolution of the Lyapunov function over longer random periods defined in terms of suitably constructed stopping times.

Theorem 1 *Fix $\lambda \in \Lambda_0^o$ and let π be a Generalized Max-Weight scheduling policy defined by the reconfiguration rule $f_\pi(\cdot)$, so that λ, f_π determine a f_π-chain.*

Denote the reconfiguration sequence by $\{S_k\}_{k\in\mathbb{N}_0}$. Suppose that for each k, a $\mathcal{F}_{S_k}$ stopping time χ_k is given, together with a compact set $\mathcal{C} \subset \mathbb{N}_0^N$, a sublinear function $F(\cdot)$, a nonnegative function $\delta(\cdot)$ with $\lim_{x\to\infty} \delta(x) = 0$, and positive constants $\epsilon, c_1, c_2 > 0$, so that the following four conditions are satisfied:

(i) $S_{k+1} \geq S_k + \chi_k \; \forall k \in \mathbb{N}_0, \forall \omega \in \mathcal{F}$, *and on occurrence of* $S_k < \infty$,
(ii) $\mathbb{E}\left[\chi_k | \mathcal{F}_{S_k}\right] \geq c_1(1 - \delta(||\mathbf{Q}(S_k)||))F(||\mathbf{Q}(S_k)||)$ *a.s.*,
(iii) $\mathbb{E}\left[\chi_k^2 | \mathcal{F}_{S_k}\right] \leq T_r^2 + c_2(F(||\mathbf{Q}(S_k)||))^2$ *a.s.*,
(iv) $\mathbb{E}\left[L(\mathbf{Q}(S_k + t + 1)) | \mathcal{F}_{S_k+t}\right] - L(\mathbf{Q}(S_k + t)) \leq -\epsilon||\mathbf{Q}(S_k + t)||$ *a.s. if it also occurs that* $\mathbf{Q}(S_k + t) \notin \mathcal{C}$, *for any* t, $S_k + \chi_k \leq S_k + t < S_{k+1}$.

Then the system is strongly stable in the sense of Definition 2, and since this holds for arbitrary $\lambda \in \Lambda_0^o$, *it follows that the policy* π *is throughput-optimal.*

Condition (ii) simply ensures that the mean time between reconfigurations goes to infinity as the total queue length grows, which is a necessary condition for stability. Condition (iii) is more closely connected to the use of a quadratic Lyapunov function. It has the effect of ensuring that the selected service vector remains "very nearly" Max-Weight and so there is no loss in quadratic Lyapunov drift. Condition (iv) allows for policies which do not reconfigure at the stopping time χ_k. The policy will still be strongly stable under Theorem 1 provided it is shown that the negative drift condition stated in (iv) holds for each $t \in [S_k + \chi_k, S_{k+1})$ whenever $\mathbf{Q}$ is outside a given compact set. (Later on, we will exhibit such a stopping time χ_k together with a compact set, to show that condition (iv) holds for the SCB policy.)

Before presenting the proof of Theorem 1, we first make a few brief comments. First, recall that $\Lambda \subseteq \Lambda_0$. The above result shows that $\Lambda = \Lambda_0$.

Second, observe that if condition (i) holds with equality, so that $S_{k+1} = S_k + \chi_k$, $k \in \mathbb{N}_0$, then condition (iv) is redundant. Moreover, the conditional form of Jensen's inequality [34, p. 88] applied to condition (iii) shows that

$$\mathbb{E}\left[\chi_k | \mathcal{F}_{S_k}\right] \leq T_r + \sqrt{c_2} F\left(||\mathbf{Q}(S_k)||\right) < \infty, \tag{12}$$

so that if $S_k < \infty$ almost surely so is S_{k+1}. Hence if (i) holds with equality, then the sequence of reconfiguration epochs satisfies $S_k < \infty$ for all $k \in \mathbb{N}_0$ almost surely. (If condition (i) does not hold with equality, then it may be the case that there are only a finite number of reconfigurations, so that $S_k = \infty, \forall k > k_0$.)

In what follows, we will consider the drift during a service interval, and hence introduce the shorthand notation

$$\Delta_{S_k} = \mathbb{E}\left[L(\mathbf{Q}(S_k + \chi_k)) | \mathcal{F}_{S_k}\right] - L(\mathbf{Q}(S_k)), k \in \mathbb{N}_0, \tag{13}$$

if $S_k < \infty$, and $\Delta_{S_k} = 0$ otherwise.

We now state two auxiliary lemmas which will enable us to obtain upper bounds for certain expressions which arise in bounding the Lyapunov drift. First we obtain an upper bound on the drift of the Lyapunov function between reconfiguration at time S_k and stopping at $S_k + \chi_k$. The proof of the lemma is deferred to Appendix 2.

Lemma 2 *Let* π *be a Generalized Max-Weight scheduling policy determined by a reconfiguration rule* $f_\pi(\cdot)$. *Also let* $\{S_k\}_{k \in \mathbb{N}_0}$ *be the associated reconfiguration sequence and suppose that the* f_π*-chain satisfies the conditions of Theorem 1.*

Then for any arrival rate vector $\lambda \in \Lambda_0^o$*, there exist fixed constants* $c_3 < \infty$, $\eta > 0$*, such that on occurrence of* $S_k < \infty$*,*

$$\Delta_{S_k} \leq c_3 - \eta F\left(||\mathbf{Q}(S_k)||\right)||\mathbf{Q}(S_k)||, \quad k \in \mathbb{N}_0.$$

The second lemma provides an upper bound on the expectation of the square of a randomly stopped sum. We omit a detailed proof, since the lemma is only a minor extension of a result in [25].

Lemma 3 *Let* $\{\mathcal{F}_t\}_{t\in\mathbb{N}}$ *be a filtration and* $\mathbf{Y}_t \in \mathbb{N}_0^N$, $t \in \mathbb{N}$*, a sequence of independent random vectors adapted to* $\{\mathcal{F}_t\}$ *and such that the components of* $\mathbf{Y}_t$ *are mutually independent. Let* $\mathbb{E}[\mathbf{Y}_t] = \lambda_t$, $\mathbb{V}[\mathbf{Y}_t] = \boldsymbol{\sigma}_t^2 < \infty$*, where expectations and variances are obtained componentwise.*

Furthermore let T *be an almost surely finite* $\mathcal{F}_t$ *stopping time and* $\tilde{T}$ *be an independent copy of* T *(an independent random variable with the same distribution).*

For any given component $\ell \in \mathcal{N}$*, let*

$$S_T^{(\ell)} = \sum_{t=1}^{T} Y_{\ell,t},$$

then

$$\mathbb{E}\left[\left(S_T^{(\ell)}\right)^2\right] \leq 2\mathbb{E}\left[\left(S_{\tilde{T}}^{(\ell)}\right)^2\right] = 2\mathbb{E}\left[\sum_{t=1}^{T}\sigma_{\ell,t}^2\right] + 2\mathbb{E}\left[\left(\sum_{t=1}^{T}\lambda_{\ell,t}\right)^2\right],$$

where $\lambda_{\ell,t}$ *and* $\sigma_{\ell,t}^2$ *are the mean and variance for component* ℓ *at time* t*.*

We are now ready to present the proof of Theorem 1.

Proof of Theorem 1 It will be convenient to take the sequence of switching epochs S_k, the following stopping times $S_k + \chi_k$, and the following slots leading up to the next switching epoch (if any) as a single infinite sequence, $\tau_0 := 0 < \tau_1 < \tau_2 < \ldots$. The sequence $\{\tau_m\}_{m\in\mathbb{N}_0}$ is defined recursively as follows. First $\tau_0 = 0$ as stated and is therefore a $\mathcal{F}_t$ stopping time. The variable τ_{m+1} is defined recursively as

$$\tau_{m+1} = \begin{cases} \tau_m + 1 \text{ if } \tau_m \neq S_k \text{ for all } k \\ S_k + \chi_k \text{ if there exists } k \leq m \text{ such that } \tau_m = S_k \end{cases}.$$

From (2) it follows that if τ_m is a $\mathcal{F}_t$ stopping time, then

$$\{\omega : \tau_m(\omega) \neq S_k(\omega), \forall k\} \in \mathcal{F}_{\tau_m}, \ \{\omega : \tau_m(\omega) = S_k(\omega)\} \in \mathcal{F}_{\tau_m}, \mathcal{F}_{S_k}.$$

Since, also, $S_k + \chi_k$ is a $\mathcal{F}_t$ stopping time, it is readily shown that $\{\tau_m\}_{m\in\mathbb{N}_0}$ is an infinite sequence of $\mathcal{F}_t$ stopping times. By assumption, if $S_k < \infty$, then $\chi_k < \infty$ almost surely so that it also holds that $\tau_m < \infty$ for all $m \in \mathbb{N}_0$ almost surely.

Define $A_{m,k} = \{\omega : S_k = \tau_m\}$, $k \le m$, and $V_m = \left(\cup_{k=0}^{m} A_{m,k}\right)^c$ as the event that τ_m is not equal to any S_k. Let $\mathbb{1}^{(m,k)}$ and $\mathbb{1}^{(m)}$ be the indicator random variables for the events $A_{m,k}$ and V_m, respectively.

If condition (i) holds with equality, the above is the sequence of reconfigurations $\{S_k\}_{k\in\mathbb{N}_0}$. In any case, define

$$\Phi_m \doteq \left(\mathbf{Q}(\tau_m), \mathbf{I}(\tau_m), \mathbb{1}^{(m)}\right), \; m \in \mathbb{N}_0, \tag{14}$$

and note that Φ_m is not necessarily a Markov chain in m as the state variable ϱ is not included. (Of course, if (i) holds with equality the sequence is a Markov chain and $\mathbb{1}^{(m)} = 0, \forall m$ is redundant.)

Now we define,

$$\Delta_{\tau_m} \doteq \mathbb{E}\left[L(\mathbf{Q}(\tau_{m+1}))|\mathcal{F}_{\tau_m}\right] - L(\mathbf{Q}(\tau_m)), m \in \mathbb{N}_0. \tag{15}$$

It follows from Lemma 2 and the definition of τ_m that there exist constants, $0 < c_3 < \infty$, $\eta > 0$, independent of m, k such that

$$\Delta_{\tau_m} \le c_3 - \eta||\mathbf{Q}(\tau_m)||, \; a.s.$$

on occurrence of $A_{m,k}$. This follows since on occurrence of $A_{m,k}$, $\tau_m = S_k < \infty$ and because $\tau_{m+1} = S_k + \chi_k \le S_{k+1}$, which implies the service vector $I(S_k)$ remains fixed prior to stopping at τ_{m+1}. The constant η is obtained on replacing $F\,(||\mathbf{Q}(S_k)||)$ with $F(1) > 0$ from Lemma 2.

Otherwise the event V_m must have occurred, which implies that $\tau_{m+1} = \tau_m + 1$. Hence

$$\Delta_{\tau_m} \le c_4 - \epsilon||\mathbf{Q}(\tau_m)||, \; a.s.$$

on occurrence of V_m. This follows from condition (iv) of Theorem 1 where we may suppose $\tau_m < \infty$. V_m occurs only if $\tau_m = S_k + t$, for some (necessarily unique) stopping time $S_k + t$ satisfying the stated inequality in condition (iv). Since there are only finitely many points in $\mathcal{C}$, there exists a constant $c_4 < \infty$, independent of m, so that the inequality holds.

The above inequalities imply

$$\Delta_{\tau_m} \le (c_3 - \eta||\mathbf{Q}(\tau_m)||) \sum_{k=1}^{m} \mathbb{1}^{(m,k)} + (c_4 - \epsilon||\mathbf{Q}(\tau_m)||)\, \mathbb{1}^{(m)} \tag{16}$$

almost surely, for any $m \in \mathbb{N}_0$.

Set $D = \max\{c_3, c_4\}$ and $\theta = \min\{\eta, \epsilon\}$. Taking expectations and using the fact that the events V_m and $\cup_{k=1}^{m} A_{m,k}$ are (almost surely) exhaustive and mutually exclusive, we derive

$$\frac{1}{M}\sum_{m=0}^{M-1}\mathbb{E}\left[\Delta_m\right] \le D - \frac{\theta}{M}\sum_{m=0}^{M-1}\mathbb{E}\left[||\mathbf{Q}(\tau_m)||\right].$$

Since $\mathbb{E}\left[\Delta_m\right] = \mathbb{E}\left[L(\mathbf{Q}(\tau_{m+1})) - L(\mathbf{Q}(\tau_m))\right]$, and $L(\mathbf{Q}(\tau_0)) = 0$, the sum on the left may be telescoped to obtain

$$0 \le \frac{1}{M}\sum_{m=0}^{M-1}\mathbb{E}\left[\Delta_m\right] = \frac{1}{M}\mathbb{E}\left[L(\mathbf{Q}(\tau_M))\right] \le D - \frac{\theta}{M}\sum_{m=0}^{M-1}\mathbb{E}\left[||\mathbf{Q}(\tau_m)||\right].$$

It thus holds that

$$\limsup_{M\to\infty}\frac{1}{M}\sum_{m=0}^{M-1}\mathbb{E}\left[||\mathbf{Q}(\tau_m)||\right] \le \frac{D}{\theta} < \infty. \tag{17}$$

The following two bounds, which we will use later, are obtained by similar arguments. First substitute the full bound of Lemma 2 into (16). Then after scaling and taking expectations, the sum may be telescoped as before to deduce

$$\begin{aligned}0 \le c_3 - \frac{\eta}{M}\sum_{m=0}^{M-1}\sum_{k=0}^{m}\mathbb{E}\left[F(||\mathbf{Q}(\tau_m)||)||\mathbf{Q}(\tau_m)||;\, A_{m,k}\right]\\ + c_4 - \frac{\epsilon}{M}\sum_{m=0}^{M-1}\mathbb{E}\left[||\mathbf{Q}(\tau_m)||;\, V_m\right].\end{aligned}$$

Since the terms in the sums are nonnegative, the following two bounds are obtained

$$\begin{aligned}\limsup_{M\to\infty}\frac{1}{M}\sum_{m=0}^{M-1}\sum_{k=0}^{m}\mathbb{E}\left[F(||\mathbf{Q}(\tau_m)||)||\mathbf{Q}(\tau_m)||;\, A_{m,k}\right] \le \frac{c_3+c_4}{\eta} < \infty.\\ \limsup_{M\to\infty}\frac{1}{M}\sum_{m=0}^{M-1}\mathbb{E}\left[||\mathbf{Q}(\tau_m))||;\, V_m\right] \le \frac{c_3+c_4}{\epsilon} < \infty.\end{aligned} \tag{18}$$

In a certain sense Eq. (17) shows that the queue length process embedded at the stopping times τ_m is strongly stable. We now proceed to prove that the queue length process at time slots itself is strongly stable.

By construction, $\tau_{m+1} - \tau_m > 1$ only if the event $A_{m,k}$ occurs for some $k \le m$. Consider only intervals $[S_k, S_k + \chi_k)$, since other steps use only one slot. For all $t \in [S_k, S_k + \chi_k)$ we have

$$Q_\ell(t) \le Q_\ell(S_k) + \sum_{u=0}^{t-S_k} A_\ell(S_k+u) \le Q_\ell(S_k) + \sum_{u=0}^{\chi_k-1} A_\ell(S_k+u).$$

Thus,

$$\sum_{s=0}^{\chi_k-1} Q_\ell(S_k+s) \le \chi_k \left(Q_\ell(S_k) + \sum_{u=0}^{\chi_k-1} A_\ell(S_k+u) \right)$$
$$\le \chi_k Q_\ell(S_k) + \chi_k \sum_{u=0}^{\chi_k-1} (1 + A_\ell(S_k+u))$$
$$\le \chi_k Q_\ell(S_k) + \left(\sum_{u=0}^{\chi_k-1} (1 + A_\ell(S_k+u)) \right)^2.$$

Summing over the queues and taking the conditional expectation, we obtain

$$\mathbb{E}\left[\sum_{s=0}^{\chi_k-1} ||\mathbf{Q}(S_k+s)|| |\mathcal{F}_{S_k} \right] \le \mathbb{E}\left[\chi_k|\mathcal{F}_{S_k}\right] ||\mathbf{Q}(S_k)|| + \sum_{\ell=1}^{N} \mathbb{E}\left[\left\{ \sum_{s=0}^{\chi_k-1} (1 + A_\ell(S_k+s)) \right\}^2 |\mathcal{F}_{S_k} \right]. \quad (19)$$

The last term inside the square is a sum of i.i.d. random variables which are also independent of the process up to time S_k.

We proceed to bound the terms on the right-hand side of (19). Since χ_k is a $\mathcal{F}_{S_k}$ stopping time, we may apply Lemma 3 to obtain

$$\mathbb{E}\left[\left\{ \sum_{s=0}^{\chi_k-1} (1 + A_\ell(S_k+s)) \right\}^2 |\mathcal{F}_{S_k} \right] \le 2\left(A_{\max}^2 - \lambda_\ell^2\right) \mathbb{E}\left[\chi_k|\mathcal{F}_{S_k}\right] + 2\left(1 + 2\lambda_\ell + \lambda_\ell^2\right) \mathbb{E}\left[\chi_k^2|\mathcal{F}_{S_k}\right]. \quad (20)$$

The first term on the right-hand side of (19) and (20) can be upper bounded using (12). The second term on the right-hand side of (20) can be upper bounded using condition (iii). Combining these results and using that the function $F(\cdot)$ is sublinear, we deduce that there exist constants $c_5, c_6 < \infty$ such that

$$\mathbb{E}\left[\sum_{s=0}^{\chi_k-1} ||\mathbf{Q}(S_k+s)|| |\mathcal{F}_{S_k} \right] \le c_5 + c_6 F\left(||\mathbf{Q}(S_k)||\right) ||\mathbf{Q}(S_k)||. \quad (21)$$

We proceed to obtain a bound for $||\mathbf{Q}(t)||$ summed over time slots. By construction of $\{\tau_m\}_{m\in\mathbb{N}_0}$, $T \le \tau_T$ for any $T \in \mathbb{N}_0$. Thus we may write

$$
\begin{aligned}
\sum_{t=0}^{T-1} ||\mathbf{Q}(t)|| &\leq \sum_{t=0}^{\tau_T - 1} ||\mathbf{Q}(t)|| \\
&= \sum_{m=0}^{T-1} \sum_{u=\tau_m}^{\tau_{m+1}-1} ||\mathbf{Q}(u)|| \\
&= \sum_{m=0}^{T-1} \sum_{u=\tau_m}^{\tau_{m+1}-1} ||\mathbf{Q}(u)|| \left(\sum_{k=0}^{m} \mathbb{1}^{(m,k)} + \mathbb{1}^{(m)} \right) \\
&= \sum_{m=0}^{T-1} \left(\sum_{u=S_k}^{S_k+\chi_k - 1} ||\mathbf{Q}(u)|| \mathbb{1}^{(m,k)} + ||\mathbf{Q}(\tau_m)|| \mathbb{1}^{(m)} \right) \\
&= \sum_{m=0}^{T-1} \left(\sum_{s=0}^{\chi_k - 1} ||\mathbf{Q}(S_k + s)|| \mathbb{1}^{(m,k)} + ||\mathbf{Q}(\tau_m)|| \mathbb{1}^{(m)} \right).
\end{aligned}
$$

It is readily seen that the sum on the right-hand side is finite almost surely. Since $\mathbb{1}^{(m,k)} \in \mathcal{F}_{S_k}$, we may take expectations on both sides and use (21) to obtain

$$
\begin{aligned}
\sum_{t=0}^{T-1} \mathbb{E}\left[||\mathbf{Q}(t)|| \right] \leq & \sum_{m=0}^{T-1} \sum_{k=0}^{m} \mathbb{E}\left[c_5 + c_6 F(||\mathbf{Q}(\tau_m)||) ||\mathbf{Q}(\tau_m)||;\, A_{m,k} \right] \\
& + \sum_{m=0}^{T-1} \mathbb{E}\left[||\mathbf{Q}(\tau_m)||;\, V_m \right].
\end{aligned}
$$

Dividing this expression by T, taking the limsup and using the bounds in (18), we may deduce

$$
\limsup_{T \to \infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}\left[||\mathbf{Q}(t)|| \right] < \infty,
$$

which completes the proof. □

Whereas Theorem 1 establishes sufficient conditions for throughput optimality, this does not by itself show that the processes themselves are ergodic, nor that unique stationary distributions exist. The following corollaries are concerned with these questions.

The first corollary establishes ergodic properties for the stochastic sequence Φ_m where it is a Markov chain. It will be a Markov chain if condition (i) of Theorem 1 holds as we mentioned earlier.

Corollary 1 *Suppose Φ_m, see (14), is a Markov chain. Then this chain is recurrent. If the chain is irreducible, then a unique stationary distribution exists, and if $\mathbf{Q}_\infty$ has this distribution, then $\mathbb{E}\left[||\mathbf{Q}_\infty|| \right] < \infty$.*

Recall that a Markov chain is recurrent iff it enters an ergodic state almost surely. The proof then follows from Eq. (17) in conjunction with the next lemma, the proof of which is presented in Appendix 3.

Lemma 4 *Suppose we have a Markov chain* $\{(\mathbf{Q}(t), J(t))\}_{t\in\mathbb{N}_0}$ *with state space lying in* $\mathcal{X} = \mathbb{N}_0^N \times \mathcal{J}$*, where* $\mathcal{J}$ *is a finite set. If it holds that*

$$\limsup_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}\left[||\mathbf{Q}(t)||\right] = E_{\max} < \infty,$$

then this Markov chain is recurrent.

If the chain is 1-Harris recurrent, then there exists a random variable $\mathbf{Q}_\infty$ *with the stationary distribution, and in addition,* $\mathbb{E}\left[||\mathbf{Q}_\infty||\right] < \infty$.

As far as Generalized Max-Weight policies are concerned, we will be able to show the stronger result that the corresponding Markov chains are 1-Harris recurrent. This fact alone does not imply that the Markov chain over time slots (8) is itself 1-Harris recurrent, since the number of time slots between renewals (visits to a given state) may have infinite expectation. The next result, however, shows that the expectation is in fact finite; for a proof see Appendix 4.

Corollary 2 *Suppose* $(\mathbf{Q}(t), \mathbf{I}(t), \mathbf{Q}(S_k(t)), \varrho(t))$ *is a* f_π*-chain satisfying conditions (i)–(iv) of Theorem 1. Suppose* $\left(\mathbf{Q}(\tau_m), \mathbf{I}(\tau_m), \mathbb{1}^{(m)}\right)$ *is a Markov chain. If, in addition, this chain is 1-Harris recurrent, then the chain* $(\mathbf{Q}(t), \mathbf{I}(t), \mathbf{Q}(S_k(t)), \varrho(t))$ *is also 1-Harris recurrent.*

The above corollary shows that if $\left(\mathbf{Q}(\tau_m), \mathbf{I}(\tau_m), \mathbb{1}^{(m)}\right)$ is a 1-Harris chain, then the chain $(\mathbf{Q}(t), \mathbf{I}(t), \mathbf{Q}(S_k(t)), \varrho(t))$ has a stationary distribution, moreover, $\mathbb{E}[||\mathbf{Q}_\infty||] < \infty$ as a consequence of strong stability.

4 Stability proofs

In this section we return to the example policies described in Sect. 3.1, and show that they satisfy the conditions stated in Theorem 1.

4.1 VFMW policy

Theorem 2 *The VFMW policy induces a* f_π *chain satisfying condition (ii), (iii) of Theorem 1 and condition (i) with equality. Hence VFMW is throughput-optimal.*

Proof It was shown in Sect. 3.1.1 that VFMW induces a f_π-chain. Moreover, by definition of the VFMW policy, we have $S_{k+1} = S_k + \chi_k$, $k \in \mathbb{N}_0$, i.e., condition (i) holds with equality. As noted below the statement of Theorem 1, condition (iv) is then redundant, and therefore only conditions (ii) and (iii) need to be verified.

By definition, $\chi_k = \lceil F\left(||\mathbf{Q}(S_k)||\right)\rceil + T_r \geq F\left(||\mathbf{Q}(S_k)||\right) + T_r$, so that condition (ii) holds with $\delta = 0$ and $c_1 = 1$. Also, since $\chi_k \leq F\left(||\mathbf{Q}(S_k)||\right) + T_r + 1$, it follows that $\chi_k^2 \leq (F(||\mathbf{Q}(S_k)||))^2 + 2F(||\mathbf{Q}(S_k)||)(T_r + 1) + T_r^2 + 2T_r + 1$. Now choose a such that $aF(0) \geq 2(T_r + 1)$ and b such that $b(F(0))^2 \geq 2T_r + 1$, then condition (iii) is satisfied for $c_2 = 1 + a + b$. □

Since the sequence τ_m, $m \in \mathbb{N}_0$ clearly coincides with S_k, $k \in \mathbb{N}_0$, Corollary 1 shows that the chain $(\mathbf{Q}(S_k), \mathbf{I}(S_k))$ as induced by the VFMW policy at reconfiguration epochs is recurrent.

Now let $\mathbf{I}_0$ be the unique Max-Weight service vector selected when reconfiguration is started with all queues empty. Since $F(||\mathbf{Q}||) < \infty$, provided $||\mathbf{Q}|| < \infty$, it follows that the state $(\mathbf{0}, \mathbf{I}_0)$ will be entered with positive probability from any given state $(\mathbf{Q}, \mathbf{I})$. This holds because there is a positive probability all queues will be empty after a finite number of reconfigurations, $R < \infty$, starting from the given ergodic state $(\mathbf{Q}, \mathbf{I})$ which has been entered. Indeed, there is positive probability of no arrivals until $(\mathbf{0}, \mathbf{I}_0)$ as it has been supposed that the arrival processes are independent across queues and over time and $\mathbb{P}\{A_\ell = 0\} > 0$ for all $\ell \in \mathcal{N}$. It follows that the chain is irreducible (confined to those states which communicate with $(\mathbf{0}, \mathbf{I}_0)$). Corollary 1 implies that the chain $(\mathbf{Q}(S_k), \mathbf{I}(S_k))$ induced by VFMW has a unique stationary distribution.

In addition, Corollary 2 now implies that the chain $(\mathbf{Q}(t), \mathbf{I}(t), \upsilon(t), \phi(t))$ (see (9)) is 1-Harris recurrent, so that the stationary distribution of the queue lengths over time slots exists with bounded first moments.

4.2 BB policy

Theorem 3 *The BB policy induces a f_π chain satisfying conditions (ii), (iii) of Theorem 1 and condition (i) with equality. Hence BB is throughput-optimal.*

Proof It was shown in Sect. 3.1.2 that BB induces a f_π-chain. Moreover, under the BB policy, we also have $S_{k+1} = S_k + \chi_k$, $k \in \mathbb{N}_0$, with $\chi_k < \infty$ almost surely if $S_k < \infty$ as shown earlier, so that condition (i) holds with equality. Thus only conditions (ii) and (iii) need to be verified. This is done in Lemmas 5 and 8, respectively, which are presented below. □

Since Theorem 3 holds, Corollary 1 now shows that the chain $(\mathbf{Q}(S_k), \mathbf{I}(S_k))$ as generated by the BB policy at reconfiguration epochs is recurrent. Irreducibility of this chain can be established by a straightforward but tedious argument which we do not present. This establishes that there is a queue state $\mathbf{Q}_1$ and corresponding Max-Weight schedule $\mathbf{I}_1$ such that entry into $(\mathbf{Q}_1, \mathbf{I}_1)$ occurs infinitely often.

Corollary 2 then implies that the chain $(\mathbf{Q}(t), \mathbf{I}(t), \mathbf{Q}(S_k), \phi(t))$ at time slots is itself 1-Harris recurrent. It follows that the stationary distribution of queue lengths over time slots exists with bounded first moments.

We now present Lemmas 5 and 8. Lemma 5 provides a lower bound for $\mathbb{E}\left[\chi_k|\mathcal{F}_{S_k}\right]$, establishing condition (ii) of Theorem 1, and Lemma 8 an upper bound for $\mathbb{E}\left[\chi_k^2|\mathcal{F}_{S_k}\right]$, establishing condition (iii). The proofs of these two lemmas are given in Appendices 5 and 8, respectively.

Lemma 5 (Lower bound on conditional first moment of χ_k) *There is a fixed constant $c_1 < \infty$ such that, on occurrence of $S_k < \infty$,*

$$\mathbb{E}\left[\chi_k | \mathcal{F}_{S_k}\right] \geq c_1(1-\delta(\mathbf{Q}(S_k)))F\left(||\mathbf{Q}(S_k)||\right), \ a.s. \ \forall k \in \mathbb{N}_0,$$

where $\delta(\cdot)$ is a nonnegative decreasing function with $\lim_{x\to\infty}\delta(x) = 0$.

Before establishing a (conditional) upper bound on $\mathbb{E}\left[\chi_k^2 | \mathcal{F}_{S_k}\right]$, we state two auxiliary lemmas which are proved in Appendices 6 and 7, respectively.

Lemma 6 *Fix an arrival rate vector λ in the interior of the stability region Λ, $\theta_{BB} > 0$, together with a sublinear function $F(\cdot)$. Given a queue vector $\mathbf{Q}$, let $\mathbf{I}$ be any corresponding Max-Weight vector. Then there exists a compact set $\mathcal{C}$ such that for all $\mathbf{Q} \notin \mathcal{C}$ there is an $\ell \in \mathcal{N}$ such that*

$$Q_\ell \geq \lceil \theta_{BB} F(||\mathbf{Q}||) \rceil,$$

and under the Max-Weight service vector $\mathbf{I}$, Q_ℓ has strictly negative drift, i.e., $I_\ell > \lambda_\ell$.

Lemma 6 shows that outside a compact set there is always a queue $\ell \in \mathcal{N}$ with a large size compared to $F(||\mathbf{Q}||)$ that has negative drift under the Max-Weight service vector. Thus, even if all the other queue lengths were to remain unchanged, this queue will drift toward empty, causing a reconfiguration by itself. Actual changes in any other queue lengths will only serve to cause this event to occur sooner, rendering the time for Q_ℓ to empty an upper bound on χ_k.

The following lemma is a general result for a single-server queue with batch service of a fixed number of customers D per time slot and i.i.d. arrivals $A(t)$ in slot $t \in \mathbb{N}_0$ with finite second moments, $\mathbb{E}\left[A^2(t)\right] < \infty$. Suppose also that the queue has strictly negative drift, i.e., $D > \lambda = \mathbb{E}\left[A(t)\right]$. With service and arrivals proceeding as described in Sect. 2 and with the natural filtration $\mathcal{H}_t \doteq \sigma(A(0), \ldots, A(t))$, suppose that $Q(0) > D$ and let T be the number of slots until $Q(T) < Q(0)$. Then we have the following lemma.

Lemma 7 *Given the above single-server queue with batch service and initial state $Q(0) > D$ satisfying the recursion*

$$Q(t+1) = [Q(t) - D]_+ + A(t), \ \forall t \in \mathbb{N}_0,$$

where $A(t) \in \mathbb{N}_0$ are i.i.d. with finite second moments and with the natural filtration $\mathcal{H}_t$ as given above. Also, suppose the drift is strictly negative, $\lambda < D$, where λ is defined above. Then T is a $\mathcal{H}_t$ stopping time with finite first and second moments.

Lemma 7 implies that the amount of time for the length of queue ℓ^* to reduce by a given amount $M_0 \geq 1$ also has finite variance, provided the initial queue length is sufficiently large.

The next lemma provides an upper bound for $\mathbb{E}\left[\chi_k^2 | \mathcal{F}_{S_k}\right]$, establishing condition (iii) of Theorem 1.

Lemma 8 (Upper bound on conditional second moment of χ_k) *There is a fixed constant $c_2 < \infty$ such that*

$$\mathbb{E}\left[\chi_k^2 | \mathcal{F}_{S_k}\right] \leq T_r^2 + c_2(F\,(||\mathbf{Q}(S_k)||))^2, \tag{22}$$

for all $\mathbf{Q}(S_k) \in \mathbb{N}_0^N$.

4.3 SCB policy

In addressing the SCB policy below, we will suppose without loss of generality that $\lambda > 0$. In addition, we will assume that there is no *dominating* service vector. A service vector $\mathbf{I}^D$ is said to be strictly dominating if it holds that $\mathbf{I}^D > \mathbf{I}_1$ for all $\mathbf{I}_1 \in \mathcal{I}, \mathbf{I}_1 \neq \mathbf{I}^D$. The vector $\mathbf{I}^D$ is said to be dominating if the inequality is not strict, but $\mathbf{I}^D$ is always selected where there are ties. By an elementary argument it can then be shown that SCB will reconfigure infinitely often provided that $\lambda \in \Lambda_0^o$.

If there is a dominating service vector, then it is unique and always Max-Weight. It is straightforward to show that almost surely the system will reconfigure at some step to adopt and maintain this service vector for such arrival rate vectors in the interior of Λ_0. It follows that the SCB policy is throughput-optimal in this case.

Theorem 4 *Suppose there is no dominating service vector. Then the SCB policy satisfies conditions (i)–(iv) of Theorem 1, and hence is throughput-optimal.*

Proof Since there is no dominating service vector, a straightforward argument which we omit shows that $S_k < \infty$ for all $k \in \mathbb{N}_0$. Denote by $\nu_k = S_{k+1} - S_k \geq T_r$ the duration of the k^{th} service interval. In order to establish conditions (i)–(iii), we will exhibit a $\mathcal{F}_{S_k}$ stopping time χ_k satisfying $\chi_k(\omega) \leq \nu_k(\omega)$ for all ω such that $S_k(\omega) < \infty$. We also will verify conditions (ii) and (iii) for χ_k.

We first prove that there exists a compact set $\mathcal{A}$ such that $\chi_k^{BB}(\omega) \leq \nu_k(\omega)$ for a given BB policy, provided $\mathbf{Q}(S_k) \notin \mathcal{A}$, with χ_k^{BB} denoting the stopping time under the BB policy, i.e., the slot at which it would reconfigure.

To show the above, let $\boldsymbol{\Delta}_\mathbf{Q}$ be the vector change in queue lengths at the time $S_{k+1} = S_k + \nu_k < \infty$ that the SCB policy reconfigures. Let $\mathbf{I}^*$ and $\mathbf{I}^o$ be the Max-Weight service vectors at time slots S_k and S_{k+1}, respectively, and fix a sample path. Then

$$\boldsymbol{\Delta}_\mathbf{Q} \cdot (\mathbf{I}^o - \mathbf{I}^*) \geq F(||\mathbf{Q}(S_{k+1})||).$$

It follows from the Cauchy–Schwarz inequality and Taylor's theorem that

$$\begin{aligned} ||\boldsymbol{\Delta}_\mathbf{Q}|| \times ||\mathbf{I}^o - \mathbf{I}^*|| &\geq F(||\mathbf{Q}(S_{k+1})||) \\ ||\boldsymbol{\Delta}_\mathbf{Q}|| &\geq \eta F(|||\mathbf{Q}(S_k)|| - ||\boldsymbol{\Delta}_\mathbf{Q}|||) \\ &= \eta\left[F(||\mathbf{Q}(S_k)||) - ||\boldsymbol{\Delta}_\mathbf{Q}||F'(y_{\boldsymbol{\Delta}_\mathbf{Q}})\right], \end{aligned}$$

with $y_{\boldsymbol{\Delta}_\mathbf{Q}} \in [||\mathbf{Q}(S_k)|| - \boldsymbol{\Delta}_\mathbf{Q}, ||\mathbf{Q}(S_k)||]$. The first inequality holds with $\eta = 1/(2N\mu_{\max})$ because $1 > \eta||\mathbf{I}^o - \mathbf{I}^*||$ and the last equality is obtained under the supposition that $||\boldsymbol{\Delta}_\mathbf{Q}|| \le ||\mathbf{Q}(S_k)||$. Invoking Definition 5, we deduce

$$||\boldsymbol{\Delta}_\mathbf{Q}|| \ge \frac{\eta}{1+\eta B} F\left(||\mathbf{Q}(S_k)||\right),$$

with $B = \max_{y \ge 0} F'(y) < \infty$. If we define $\theta_{BB} = \frac{\eta}{1+\eta B}$, then this inequality implies that the corresponding BB policy must have reconfigured already by time S_{k+1}, provided $||\mathbf{Q}(S_k)|| > A_1$ for some constant $A_1 > 0$. Clearly the same conclusion holds for the same BB policy if it was the case that $||\boldsymbol{\Delta}_\mathbf{Q}|| > ||\mathbf{Q}(S_k)||$, provided $||\mathbf{Q}(S_k)|| > A_2$ for some constant A_2, so that $F(y) < y$ for all $y > A_2$. We have thus shown that if we take $\mathcal{A} \doteq \{\mathbf{Q} : ||\mathbf{Q}|| \le \max\{A_1, A_2\}\}$, then we obtain the required compact set.

We now set $\chi_k = T_r$ if $\mathbf{Q}(S_k) \in \mathcal{A}$ and $\chi_k = \chi_k^{BB}$ if $\mathbf{Q}(S_k) \notin \mathcal{A}$. The above property then implies that condition (i) is satisfied. It follows directly from Theorem 3 that conditions (ii) and (iii) are satisfied when $\mathbf{Q}(S_k) \notin \mathcal{A}$ for some suitable choice of constants c_1, c_2 and function $\delta(\cdot)$. Since the set $\mathcal{A}$ is compact and therefore finite we may take $\delta = 1$ on $\mathcal{A}$ to meet (ii), and (iii) is satisfied by definition of χ_k in $\mathcal{A}$.

It remains to be established that condition (iv) is satisfied. First we define the single-step drift,

$$\Xi_k(t) = \mathbb{E}\left[L(\mathbf{Q}(S_k + t + 1)) - L(\mathbf{Q}(S_k + t))|\mathcal{F}_{S_k}\right],$$

which we will investigate only when $S_k < \infty$. Recalling the definition of χ_k as given above, suppose that $t \in [\chi_k, \nu_k)$. It follows that if $\mathbf{I}^o$ is the service vector selected at time S_k, then for any service vector $\mathbf{I} \in \mathcal{I}$ it holds that

$$\mathbf{Q}(S_k + t) \cdot \mathbf{I} < F(||\mathbf{Q}(S_k + t)||) + \mathbf{I}^o \cdot \mathbf{Q}(S_k + t). \tag{23}$$

If this were not the case, then by definition a reconfiguration would occur at time $S_k + t$, contradicting the supposition that $t \in [\chi_k, \nu_k)$.

However, (23) shows that while $\mathbf{I}^o$ is no longer necessarily the Max-Weight service vector at time $S_k + t$, it is still in some sense close to being so. We will now obtain an estimate for the drift under $\mathbf{I}^o$. For any arrival rate vector λ that is in the interior of the stability region Λ, there exist real numbers $\alpha^1, \ldots, \alpha^{|\mathcal{I}|}$ such that $\alpha^j \ge 0$ for all $j = 1, \ldots, |\mathcal{I}|$, $\sum_{j=1}^{|\mathcal{I}|} \alpha^j = 1 - \epsilon$ for some $\epsilon > 0$ and $\lambda = \sum_{j=1}^{|\mathcal{I}|} \alpha^j \mathbf{I}^j$. Define $K = N(A_{\max}^2 + \mu_{\max}^2)$ and denote by $\mathbf{I}^*(\mathbf{Q}(t))$ the Max-Weight service vector for $\mathbf{Q}(t)$; in the case of ties let it be decided as explained earlier.

Then

$$\Xi_k(t) \le 2K + 2\mathbf{Q}(S_k + t) \cdot \left(\sum_{j=1}^{|\mathcal{I}|} \alpha^j \mathbf{I}^j\right) - 2\left(\mathbf{Q}(S_k + t) \cdot \mathbf{I}^*(\mathbf{Q}(S_k + t))\right.$$
$$\left. - F(||\mathbf{Q}(S_k + t)||)\right)$$

$$
\begin{aligned}
&= 2\left(K + F(||\mathbf{Q}(S_k + t)||)\right) + 2\mathbf{Q}(S_k + t) \cdot \mathbf{I}^*(\mathbf{Q}(S_k + t))(1 - \epsilon) \\
&\quad -2\mathbf{Q}(S_k + t) \cdot \mathbf{I}^*(\mathbf{Q}(S_k + t)) \\
&= 2\left(K + F(||\mathbf{Q}(S_k + t)||)\right) - 2\epsilon \mathbf{Q}(S_k + t) \cdot \mathbf{I}^*(\mathbf{Q}(S_k + t)) \\
&\le 2\left(K + F(||\mathbf{Q}(S_k + t)||)\right) - \frac{2\epsilon}{N}||\mathbf{Q}(S_k + t)|| \\
&\le -\frac{\epsilon}{N}||\mathbf{Q}(S_k + t)||,
\end{aligned}
$$

as long as $\mathbf{Q}(S_k + t)$ is outside a compact set $\mathcal{C}_\epsilon$.

The first inequality follows from taking expectations after squaring both sides and finally applying (23), choosing $\mathbf{I}$ to be $\mathbf{I}^*$ on the left-hand side and substituting for λ. The second inequality follows by replacing each service vector with the Max-Weight service vector and then substituting. The third equality is immediate. The fourth inequality follows from the choice of $\mathcal{C}_\epsilon$ and the inequality

$$\mathbf{Q} \cdot \mathbf{I}^* \ge Q_{\max} \ge ||\mathbf{Q}||/N,$$

where the first inequality holds since for at least one service vector $\mathbf{I} \in \mathcal{I}$ a queue with maximum queue length is served, and the second inequality follows immediately.

This shows that condition (iv) is satisfied on taking the compact set $\mathcal{C}_\epsilon$ and with ϵ determined by the arrival rate vector. □

It is readily seen that $\Phi^{(SCB)}$ is a Markov chain, and Corollary 1 implies that the chain is recurrent. (As already mentioned, we neglect the case where there are only finitely many reconfigurations, as this implies that there is a unique Max-Weight service vector over all queue lengths.)

As with the BB policy, it can be shown that there is a given state $(\mathbf{Q}, \mathbf{I}, 0)$ which is both ergodic and which is entered almost surely. It follows that Φ is 1-Harris recurrent and therefore by Corollary 2 the full chain $(\mathbf{Q}(t), \mathbf{I}(t), \varrho(t))$ over time slots is also 1-Harris recurrent.

5 Performance considerations

As indicated by Theorems 1-4, the VFMW, SCB, and BB policies achieve throughput optimality for any sublinear function $F(\cdot)$. The question thus naturally arises whether some candidate functions $F(\cdot)$ may provide better performance than others in terms of mean delays or mean queue lengths for example. While a detailed investigation is beyond the scope of the present paper, we now briefly examine this question. First of all, define

$$\rho = \inf\{\gamma : \lambda \in \gamma \Lambda_0\}$$

as the 'load' of the system, so that the condition that the arrival rate vector λ lies in the interior of Λ_0 may be equivalently written as $\rho < 1$.

Now observe that in order for the system to be stable, at least a fraction of the time ρ one of the service vectors must be invoked, and hence at most a fraction of the time

$1 - \rho$ can be spent on reconfiguration. This can be expressed in terms of the expected length of a service interval as

$$\frac{T_r}{\mathbb{E}[\chi]} \leq 1 - \rho,$$

with $\chi \geq T_r$ a random variable with the limiting distribution of χ_k as $k \to \infty$, yielding

$$\mathbb{E}[\chi] \geq \frac{T_r}{1 - \rho}.$$

By definition of the VFMW policy, and by the fact that the BB and SCB policies satisfy condition (ii) of Theorem 1, the expected length of a service period can in turn be related to the queue lengths at the beginning of a service period by

$$\mathbb{E}[\chi] \leq T_r + \beta \mathbb{E}[\lceil F(||\mathbf{Q}||) \rceil] \leq T_r + \beta \mathbb{E}[F(||\mathbf{Q}||)] + 1,$$

with $\beta = 1$ for the VFMW policy, $\beta = \sqrt{c_2}$ for the BB and SCB policies, and $\mathbf{Q}$ denoting a random vector with the limiting distribution of $\mathbf{Q}(S_k)$ as $k \to \infty$. In the case where $F(\cdot)$ is concave, so that the inverse $F^{-1}(\cdot)$ is convex, Jensen's inequality yields

$$\mathbb{E}[||\mathbf{Q}||] \geq F^{-1}(\mathbb{E}[F(||\mathbf{Q}||)]).$$

Combining the above inequalities, we obtain

$$\sum_{l=1}^{N} \mathbb{E}[Q_l] \geq F^{-1}\left(\frac{1}{\beta}\left(\frac{\rho T_r}{1 - \rho} - 1\right)\right). \tag{24}$$

In particular, in the case $F(s) = s^\alpha$, $\alpha \in (0, 1)$,

$$\sum_{l=1}^{N} \mathbb{E}[Q_l] \geq \left(\frac{1}{\beta}\left(\frac{\rho T_r}{1 - \rho} - 1\right)\right)^{1/\alpha}. \tag{25}$$

The above lower bound suggests that the higher the value of α, i.e., the closer to 1, the smaller the total expected queue length at the beginning of a service interval. This may further suggest that setting the value of α equal to 1 may yield an even smaller total expected queue length. Observe, however, that Lemma 2 does not cover that case, and in fact the proof of that lemma no longer applies then. This is not just an artifact of the proof, nor does it imply that such a choice would necessarily fail to achieve maximum stability, but it does reflect a fundamental issue associated with the quadratic Lyapunov function.

In order to illustrate that, let us briefly revisit the two-queue scenario described in Sect. 2.4, and assume that the VFMW policy is used with $F(s) = \theta s$. Suppose that the queue lengths at the beginning of a service interval are $(Q_1, Q_2) = (M, M + 1)$, so that the activity vector (0, 1) is selected, and invoked for a period of time $\theta(2M+1)$. At the end of the service interval, the expected queue lengths are

$$\mathbb{E}\left[Q_1'\right] = M + \lambda_1\chi = M + \lambda_1(T_r + \theta(2M+1)) = (1+2\lambda_1\theta)M + \lambda_1(T_r+\theta),$$

and

$$\begin{aligned}\mathbb{E}\left[Q_2'\right] &\geq M + 1 + \lambda_2\chi - (\chi - T_r) = M + 1 + \lambda_2 T_r + (\lambda_2 - 1)(\chi - T_r)\\ &= M + 1 + \lambda_2 T_r + (\lambda_2 - 1)\theta(2M+1)\\ &= (1 - 2(1-\lambda_2)\theta)M + \lambda_2(T_r + \theta) - \theta + 1.\end{aligned}$$

Thus, at the end of the service interval, the expected value of the Lyapunov function is

$$\begin{aligned}\mathbb{E}\left[(Q_1')^2\right] + \mathbb{E}\left[(Q_2')^2\right] &\geq (\mathbb{E}\left[Q_1'\right])^2 + (\mathbb{E}\left[Q_2'\right])^2\\ &\geq [(1+2\lambda_1\theta)^2 + (1-2(1-\lambda_2)\theta)^2]M^2 - \mathrm{o}(M^2)\\ &= \Big(2 + 4\theta\big[\lambda_1 + \lambda_2 - 1 + \theta\big(\lambda_1^2 + \big(1-\lambda_2\big)^2\big)\big]\Big)M^2 - \mathrm{o}(M^2).\end{aligned}$$

Note that for any given value of $\theta > 0$, the latter expression is larger than $M^2 + (M+1)^2$ for M sufficiently large and $\lambda_1 + \lambda_2$ sufficiently close to 1. This implies that a quadratic Lyapunov function *cannot* have negative drift, unless the value of θ is sufficiently small compared to $1 - \lambda_1 - \lambda_2$.

Remark 1 The above issue is also illuminated when we consider fluid limits where the system dynamics are scaled both in space and time. Under mild assumptions, a Markov chain is positive-recurrent when its fluid limit reaches zero in finite time for any initial state. We claim (without proof) that for linear $F(\cdot)$ functions, the fluid limits will follow piecewise linear trajectories. For ρ sufficiently close to 1, the $\mathcal{L}_2$ norm may increase along some of the segments, so that a quadratic Lyapunov function cannot be used to show that the fluid limit will reach zero in finite time.

Remark 2 The fact that quadratic Lyapunov functions do not allow for linear functions $F(\cdot)$ also points to an issue in the work of Hung and Chang [16] who consider a so-called dynamic cone policy. The latter policy causes the system to use a service vector for a linear amount of time (with respect to the current queue lengths). The paper relies on a quadratic Lyapunov function in order to claim that the dynamic cone policy provides maximum stability for suitably constructed cones. The negative drift is, however, argued by assuming a uniform constant G in the proof of Theorem 2 in [16], which is based on Lemma 1, where the upper bound B for the duration of an interswitching interval in fact *does not hold uniformly*, but depends on the arrival rate vector λ as well as the initial queue length vector. Indeed, $B \geq c/\Delta$, where c depends linearly on the initial queue length vector and $\Delta = \min_{\mathbf{I}\in\mathcal{I}:\max_{\ell\in\mathcal{N}}(I_\ell - \lambda_\ell)>0} \min_{\ell\in\mathcal{N}:I_\ell>\lambda_\ell}(I_\ell - \lambda_\ell)$, which could be arbitrarily close to zero for arrival rate vectors close to one of the maximal feasible service vectors. (Note that if the interswitching interval were uniformly bounded by B, this would in fact mean that the fraction of time spent on reconfiguration would be uniformly bounded away from zero by $T_r/(B+T_r)$. This in turn would preclude stability when the load is sufficiently close to unity, and hence directly rule out maximum stability.)

Since the upper bound B *does not hold uniformly*, the coefficient G in the proof of Theorem 2 in [16] is *not a uniform bound* either. Indeed, the coefficient G includes the term in Eq. (26) in [16] which grows quadratically with B and hence quadratically with the initial queue length vector. As a result, the coefficient G cannot possibly be offset by the negative linear term in Eq. (34) in [16]. The claim that the quadratic Lyapunov function has negative drift therefore lacks basis (and in fact *cannot* be valid as we observed above without further assumptions on the load vector). The dynamic cone policy may still provide maximum stability, but a different proof method with a nonquadratic Lyapunov function is required to establish that.

In order to further illuminate the above issue, it is useful to revisit the two-queue example considered in Sect. 2.4. Under the dynamic cone policy we then switch from serving queue i to serving queue $3-i$ when Q_{3-i} exceeds $r_{3-i}Q_i$, with $r_i > 1$, $i = 1, 2$. When $\lambda_1 + \lambda_2$ is sufficiently close to unity, the queue lengths (Q_1, Q_2) will essentially move between the rays $Q_1 = r_1 Q_2$ and $Q_2 = r_2 Q_1$ along trajectories that are virtually perpendicular to the diagonal $Q_1 = Q_2$. Clearly, when $r_1 \neq r_2$, a quadratic Lyaponov function can then only generate negative drift in one direction, but not in the other. When $r_1 = r_2 = r$, a quadratic Lyapunov function will have negative drift in both directions. (Note that the dynamic cone policy then corresponds to the SCB policy with $F(x) = fx$ and $r = (1+f)/(1-f)$.) However, even when the Lyapunov function does have negative drift, substantially sharper bounds than Eqs. (26) and (33) in [16] are necessary to show that. Of course, there is certainly no reason in this two-queue scenario to doubt that the dynamic cone policy achieves maximum stability, but the proof technique in [16] is fundamentally ill-suited to prove that.

6 Simulation results

The first simulation results are for a two-queue network with mean number of arriving packets per slot (0.6, 0.2) and both distributed as a Pareto random variable A so that

$$\mathbb{P}\{A \geq k\} = \frac{\eta^{\beta-1}}{k^{\beta-1}},\ k \geq 1$$

and $\beta = 4$ so that the first two moments exist. The function $F(\cdot)$ was chosen to be y^α (which is 0 at $y = 0$ but this does not affect what follows). Our results are for the mean total number of packets in queue Q_T, and are presented for the VFMW policy and the SCB policy with $T_r = 3$ and a range of choices of the parameter α (Fig. 2).

The results indicate that the mean total number in queue gets smaller as α approaches 1 from below. In other words, the number in queue is smallest for the least amount of switching. Our results also show that the mean total queue length under the VFMW policy is larger than the lower bound (25), as it should be. SCB outperforms this bound, particularly at lower values of α.

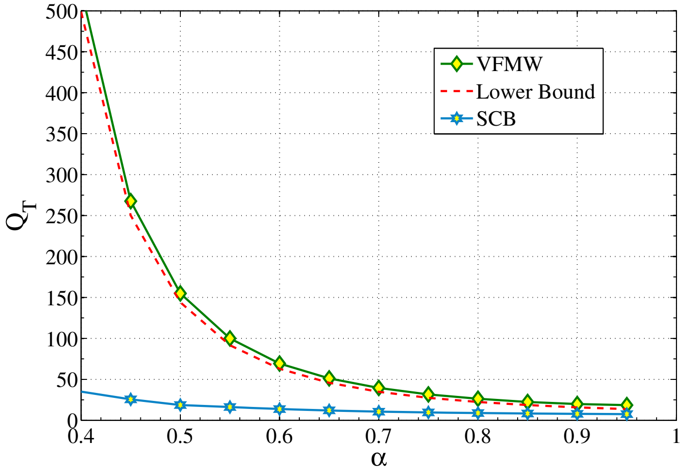

Fig. 2 Total mean queue length Q_T vs α for VFMW and SCB policies

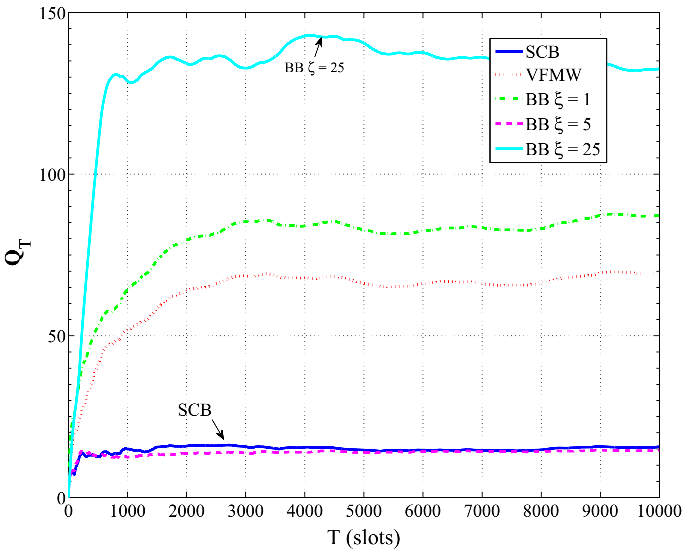

Fig. 3 Sample time average total queue length

Figure 3 shows results for Q_T as a function of T slots elapsed. The values of the Pareto parameters η_1, η_2, η_3 were chosen so as to obtain a mean number of arrivals per slot equal to (0.2, 0.7, 0.6).

Results were obtained for the VFMW and SCB policies, and show that the mean total queue length under the SCB policy is significantly smaller than that for the VFMW policy.

Additional results were obtained for the BB policy, and showed that the performance of the BB policy is highly sensitive to the choice of θ_{BB}. Small queue lengths were

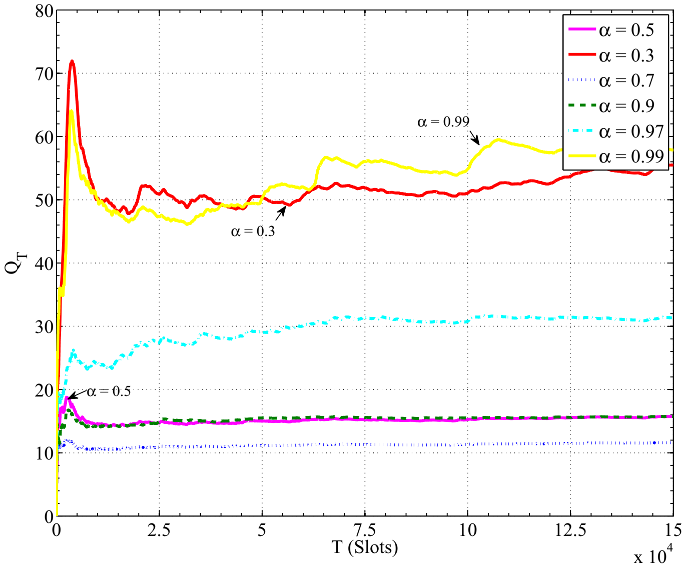

Fig. 4 Sample time average for SCB various α parameters

found with $\theta_{BB} = 5$. However, both $\theta_{BB} = 1$ and $\theta_{BB} = 25$ produced much poorer results. In the case of $\theta_{BB} = 1$ this is because of too much switching, whereas it appears that $\theta_{BB} = 25$ does not switch enough.

The following results, depicted in Fig. 4, were obtained for functions $F(\mathbf{Q}) = (||\mathbf{Q}||)^{\alpha}$ for various values of α. The results show that intermediate values of α yield significantly better performance. The average delays for $\alpha = 0.3, 0.9$ are about five times longer than for $\alpha = 0.5, 0.7$.

7 Conclusions

We investigated scheduling in networks with interference constraints and reconfiguration delays, and showed that the ordinary Max-Weight policy may fail to achieve throughput optimality in the case of nonzero reconfiguration delays. Motivated by the latter issue, we proposed a class of adaptive scheduling policies which persist with the current service vector until a certain stopping criterion is reached, before switching to the next service vector.

While earlier proposed VFMW policies belong to this class of policies, we also presented BB and SCB policies that are more responsive to bursty arrivals and queue dynamics. We developed novel Lyapunov drift techniques to prove that this class of policies under certain conditions achieve throughput optimality by dynamically adapting the durations of the service intervals. In particular, we proved that if a policy has a negative drift throughout a service interval, and if another policy with an earlier switching decision is throughput-optimal, then the former policy is also throughput-optimal.

The VFMW policies persist with a service vector for an amount of time that depends on the initial queue lengths, and thus dynamically adapt the service interval to changing queue lengths. The BB and SCB policies are more responsive to sudden changes in queue states than the VFMW policy, as they do not fix the service interval in advance, but make switching decisions based on the evolution of the queues. As a result, the BB and in particular the SCB policy make infrequent reconfiguration decisions for large queue lengths, while enabling frequent reconfiguration for small queue lengths, thus producing better delay performance than the VFMW policy.

Since the proposed scheduling algorithms always select the Max-Weight service vector when a reconfiguration takes place, the complexity is similar, except that the computation is only required once per service interval, rather than in each time slot as in the ordinary Max-Weight algorithm. In other words, the computational burden can be amortized over a much longer time period, especially in high-load regimes where the queue lengths are large and the service intervals are correspondingly long. It is worth observing that even in standard scenarios without any reconfiguration delays, the above observation can be exploited to significantly reduce the frequency of the computation of the Max-Weight service vector.

In future work, we intend to develop low-complexity distributed policies for systems with nonzero reconfiguration delay that additionally achieve asymptotically optimal delay performance. The problem of joint scheduling and routing in multihop networks with interference constraints and reconfiguration delays provides a further relevant direction for future research.

Acknowledgments Philip Whiting would like to acknowledge the generous funding provided by MacQuarie University in respect of the Vice-Chancellor's Innovation Fellowship Fund which partially supported this research.

Appendix 1: Proof of Lemma 1

Let the state of the process be given by $(\mathbf{Q}(t), I(t), \varrho(t))$ where $\mathbf{Q}(t)$ is the queue length vector at the start of the slot t, $I(t)$ indicates which queue is being served, and $\varrho(t)$ denotes the number of reconfiguration slots to go (if any). Under the ordinary Max-Weight policy, if $Q_1(t) > Q_2(t)$ and $I(t) = 2$, then reconfiguration is invoked so that $I(t+1) = 1$, $\varrho(t+1) = 1$, and service is not applied in slot t. The opposite is done if $Q_2(t) > Q_1(t)$. Otherwise the current service vector is applied, followed by any queue arrivals, as explained in Sect. 2.1. (Note that the state variable $\varrho(t)$ is redundant.) With the state $(\mathbf{Q}(t), I(t), \varrho(t))$ the process forms a time-homogeneous Markov chain with time index $t \in \mathbb{N}_0$.

Let $\mathcal{D} \doteq \{\mathbf{Q} : |Q_1 - Q_2| \leq 3\}$; elementary considerations show that the main diagonal $Q_1 = Q_2$ is reached with probability 1 for $p \in (0, 1/2)$. Once the main diagonal is reached the process will be confined to $\mathcal{D}$. Restricting to slots where reconfiguration is not invoked, a slot t is said to be a *freeze* slot if $\mathbf{Q}(t) = \mathbf{Q}(t+1)$ and is said to be *progressive* otherwise. Finally, further consideration shows that reconfiguration occurs infinitely often a.s. with at most four progressive slots occurring between two slots where reconfiguration is invoked, again once the main diagonal is reached.

To show the Markov chain is transient (instability), let S_k be the slot in which the kth reconfiguration takes place and let $M_k = Q_1(t_k) + Q_2(t_k)$ be the total queue length at time S_k, $k = 0, 1, 2, \ldots$. Consider the interval between one reconfiguration and the next, where we will examine the mean drift of M_k. Since a freeze slot occurs with probability $p(1-p)$, the number of freeze slots until the occurrence of a progressive one is geometrically distributed with mean $\frac{p(1-p)}{1-p+p^2}$. Furthermore, the total drift along the main diagonal per service slot is $-(1-2p)$, which is negative (toward the origin) for $p < 1/2$.

From the above it follows that the accumulated drift during a service interval is no less than

$$-4(1-2p) \times \left(1 + \frac{p(1-p)}{1-p+p^2}\right) = \frac{-4(1-2p)}{1-p+p^2}.$$

During a reconfiguration slot, the positive drift (away from the origin) is $2p$. Thus, the total accumulated drift along the main diagonal must be positive when

$$2p - \frac{4(1-2p)}{1-p+p^2} > 0,$$

which amounts to $p > 0.42049$.

We have thus shown that there is a $p < 1/2$ such that the sequence M_k satisfies

$$\mathbb{E}\left[M_{k+1} | M_k\right] \geq M_k + \theta,$$

where $\theta > 0$ is a fixed constant. We now show that $M_k \to \infty$ as $k \to \infty$ almost surely, which in turn implies $Q_\ell(t) \to \infty$ as $t \to \infty$, $\ell = 1, 2$. Introduce $Z_k = M_{k+1} - M_k$, $k = 0, 1, 2, \ldots$, and observe that $Z_k \leq 6$. Additionally, define $R_k = \frac{1}{M_k+1}$, and observe

$$\begin{aligned}\mathbb{E}\left[\frac{1}{M_{k+1}+1} | \mathcal{F}_k\right] &= \frac{1}{M_k+1} - \mathbb{E}\left[\frac{Z_k}{(M_k+1)(M_k+1+Z_k)}\right] \\ &\leq \frac{1}{M_k+1} - \frac{\theta}{(M_k+1)(M_k+1+6)},\end{aligned}$$

from which it follows that R_k is a nonnegative supermartingale where $\mathcal{F}_k$ is an appropriate sigma algebra at time slot t_k. Hence $\lim_{k\to\infty} R_k \leq 1$ exists almost surely, and so does $\lim_{k\to\infty} M_k$. Finally,

$$\mathbb{E}\left[R_\infty\right] \leq \liminf_{k\to\infty} \mathbb{E}\left[R_k\right] \leq 1 - \sum_{k=0}^{\infty} \mathbb{E}\left[\frac{\theta}{(M_k+1)(M_k+7)}\right],$$

which implies

$$\liminf_{k\to\infty} \mathbb{E}\left[\left(\frac{\theta}{(M_k+7)}\right)^2\right] \leq \liminf_{k\to\infty} \mathbb{E}\left[\frac{\theta}{(M_k+1)(M_k+7)}\right] = 0.$$

This shows that $\liminf_{k\to\infty}(M_k+7)^{-2}=0$ almost surely by Fatou's lemma, and since actually there is a limit, we have $\lim_{k\to\infty} M_k=\infty$ almost surely. This then implies $\lim_{k\to\infty} Q_\ell(t_k)=\infty$ almost surely, since $|Q_1-Q_2|\le 3$. Moreover, if $t\in[t_k,t_{k+1})$, then $|Q_\ell(t)-Q_\ell(t_k)|\le 4$, $\ell=1,2$, so that $Q_\ell(t)\to\infty$ as $t\to\infty$ almost surely. Since $\liminf_{t\to\infty}\mathbb{E}\left[Q_\ell(t)\right]\ge\mathbb{E}\left[\lim_{t\to\infty}Q_\ell(t)\right]=\infty$, the system is not strongly stable. □

Appendix 2: Proof of Lemma 2

On occurrence of $S_k<\infty$, condition (iii) in Theorem 1 implies that $\mathbb{E}\left[\chi_k|\mathcal{F}_{S_k}\right]\le T_r+\sqrt{c_2}F(\mathbf{Q}(S_k))<\infty$, so that χ_k is finite almost surely. Using the Lindley recursion in (1) for time slots $t=S_k,\ldots,S_k+\chi_k-1$, we obtain

$$Q_\ell(S_k+\chi_k)\le\max\left\{Q_\ell(S_k)-\sum_{s=T_r-1}^{\chi_k-1}I_\ell(S_k+s),0\right\}+\sum_{s=0}^{\chi_k-1}A_\ell(S_k+s). \quad (26)$$

Note that if $\sum_{s=T_r-1}^{\chi_k-1}I_\ell(S_k+s)$, representing the total service opportunity given to queue ℓ during the k^{th} service interval, is smaller than $Q_\ell(S_k)$, then the inequality in (26) in fact holds with equality. Otherwise, the first term is 0 and (26) holds with inequality, because some of the arrivals during the service interval might depart before the end of the service interval.

Squaring both sides of (26), using $\max\{0,x\}^2\le x^2$ and $I_\ell(t)\le\mu_{\max}$ for all $\ell=1,\ldots,N$, $t=S_k+T_r-1,\ldots,S_k+\chi_k-1$, we obtain

$$\begin{aligned}Q_\ell(S_k+\chi_k)^2-Q_\ell(S_k)^2&\le\chi_k^2\mu_{\max}^2+\left(\sum_{s=0}^{\chi_k-1}A_\ell(S_k+s)\right)^2\\&\quad-2Q_\ell(S_k)\left(\sum_{s=T_r-1}^{\chi_k-1}I_\ell(S_k+s)-\sum_{s=0}^{\chi_k-1}A_\ell(S_k+s)\right). \quad (27)\end{aligned}$$

Summing (27) over the queues, taking conditional expectations, and then using Wald's equality (or optional stopping), we derive

$$\begin{aligned}\Delta_{S_k}&\le N\mu_{\max}^2\mathbb{E}\left[\chi_k^2|\mathcal{F}_{S_k}\right]+\sum_{\ell=1}^N\mathbb{E}\left[\left(\sum_{s=0}^{\chi_k-1}A_\ell(S_k+s)\right)^2|\mathcal{F}_{S_k}\right] \quad (28)\\&\quad+2\mathbb{E}\left[\chi_k|\mathcal{F}_{S_k}\right]\sum_{\ell=1}^N\lambda_\ell Q_\ell(S_k)-2\sum_{\ell=1}^N Q_\ell(S_k)\mathbb{E}\left[\sum_{s=T_r-1}^{\chi_k-1}I_\ell(S_k+s)|\mathcal{F}_{S_k}\right],\end{aligned}$$

where we used the fact that the arrival processes are i.i.d. over time and independent of the queue lengths.

Now observe that for any arrival rate vector $\lambda \in \Lambda_0^0$, there exist real numbers $\beta^1, \ldots, \beta^{|\mathcal{I}|}$ such that $\beta^j \geq 0$ for all $j = 1, \ldots, |\mathcal{I}|$, $\sum_{j=1}^{|\mathcal{I}|} \beta^j = 1 - \epsilon$ for some $\epsilon > 0$ and $\lambda = \sum_{j=1}^{|\mathcal{I}|} \beta^j \mathbf{I}^j$.

Substituting the latter expression in (28) and using conditions (ii) and (iii) of Theorem 1, we obtain

$$\begin{aligned}
\Delta_{S_k} &\leq N\mu_{\max}^2(T_r^2 + c_2(F(||\mathbf{Q}(S_k)||))^2) + \sum_{\ell=1}^{N} \mathbb{E}\left[\left(\sum_{s=0}^{\chi_k-1} A_\ell(S_k+s)\right)^2 |\mathcal{F}_{S_k}\right] \\
&\quad + 2\mathbb{E}\left[\chi_k|\mathcal{F}_{S_k}\right]\mathbf{Q}(S_k)\cdot\sum_{j=1}^{|\mathcal{I}|}\beta^j\mathbf{I}^j - 2\mathbb{E}\left[\chi_k - T_r|\mathcal{F}_{S_k}\right]\mathbf{Q}(S_k)\cdot\mathbf{I}^*(S_k) \\
&\leq N\mu_{\max}^2(T_r^2 + c_2(F(||\mathbf{Q}(S_k)||))^2) + \sum_{\ell=1}^{N} \mathbb{E}\left[\left(\sum_{\tau=0}^{\chi_k-1} A_\ell(S_k+\tau)\right)^2 |\mathcal{F}_{S_k}\right] \\
&\quad -2c_1\epsilon(1-\delta(||\mathbf{Q}(S_k)||))F(||\mathbf{Q}(S_k)||)\,\mathbf{Q}(S_k)\cdot\mathbf{I}^*(S_k) + 2T_r\mathbf{Q}(S_k)\cdot\mathbf{I}^*(S_k),
\end{aligned} \tag{29}$$

where in the last inequality we used the fact that $\mathbf{Q}(S_k)\cdot\mathbf{I}^*(S_k) \geq \mathbf{Q}(S_k)\cdot\mathbf{I}$ for all $\mathbf{I} \in \mathcal{I}$ by definition of the Max-Weight service vector.

Applying Lemma 3 to the second term in (29), we derive

$$\begin{aligned}
&\mathbb{E}\left[\left(\sum_{s=0}^{\chi_k-1} A_\ell(S_k+s)\right)^2 |\mathcal{F}_{S_k}\right] \\
&\quad\leq 2\mathbb{E}\left[\sum_{s=0}^{\chi_k-1} A_{\max}^2 - \lambda_\ell^2|\mathcal{F}_{S_k}\right] + 2\mathbb{E}\left[\left(\sum_{s=0}^{\chi_k-1}\lambda_\ell\right)^2 |\mathcal{F}_{S_k}\right] \\
&\quad= 2\left(A_{\max}^2 - \lambda_\ell^2\right)\mathbb{E}\left[\chi_k|\mathbf{Q}(S_k)\right] + 2\lambda_\ell^2\mathbb{E}\left[\chi_k^2|\mathcal{F}_{S_k}\right].
\end{aligned}$$

Using conditions (ii) and (iii) of Theorem 1, we obtain

$$\mathbb{E}\left[\left(\sum_{s=0}^{\chi_k-1} A_\ell(S_k+s)\right)^2 |\mathcal{F}_{S_k}\right] \leq d_1(F(||\mathbf{Q}(S_k)||))^2,$$

for some constant $d_1 < \infty$. Substituting the latter inequality in (29), we derive

$$\Delta_{S_k} \leq N\mu_{\max}^2 T_r^2 + F(||\mathbf{Q}(S_k)||) \times\left((c_2 N\mu_{\max}^2 + d_1)F(||\mathbf{Q}(S_k)||) - 2\left(c_1(1-\delta(||\mathbf{Q}(S_k)||))\epsilon - \frac{T_r}{F(||\mathbf{Q}(S_k)||)}\right) \times\mathbf{Q}(S_k)\cdot\mathbf{I}^*(S_k)\right).$$

Since $F(||\mathbf{Q}(S_k)||)$ is a monotonically increasing function of $||\mathbf{Q}(S_k)||$, and $\delta(||\mathbf{Q}(S_k)||)$ is a monotonically decreasing function of $||\mathbf{Q}(S_k)||$, there exists a constant d_2 such that if $||\mathbf{Q}(S_k)|| > d_2$, then $c_1(1-\delta(||\mathbf{Q}(S_k)||))\epsilon - \frac{T_r}{F(||\mathbf{Q}(S_k)||)} > \delta_1 > 0$, yielding

$$\Delta_{S_k} \leq N\mu_{\max}^2 T_r^2 + (c_2 N\mu_{\max}^2 + d_1)(F(||\mathbf{Q}(S_k)||))^2 - 2\delta_1 F(||\mathbf{Q}(S_k)||)\,\mathbf{Q}(S_k)\cdot\mathbf{I}^*(S_k).$$

Hence, for $||\mathbf{Q}(S_k)|| > d_2$, we use $\mathbf{Q}(S_k)\cdot\mathbf{I}^*(S_k) \geq \frac{1}{N}\sum_{\ell=1}^{N} Q_\ell(t_k)$ to arrive at

$$\Delta_{S_k} \leq N\mu_{\max}^2 T_r + (c_2 N\mu_{\max}^2 + d_1)(F(||\mathbf{Q}(S_k)||))^2 - \frac{2\delta_1}{N} F(||\mathbf{Q}(S_k)||)\,||\mathbf{Q}(S_k)||.$$

Since $F(\cdot)$ is a sublinear function, it follows that there exist fixed constants $c_3 < \infty$, $\eta = \delta_1/N$ such that

$$\Delta_{S_k} \leq c_3 - \eta F(||\mathbf{Q}(S_k)||)\,||\mathbf{Q}(S_k)||.$$

This completes the proof. □

Appendix 3: Proof of Lemma 4

Divide the state space into the set $\mathcal{E}$ containing the ergodic states and its complement $\mathcal{T}$ consisting of the null and transient states. By Markov's inequality, given $\varepsilon > 0$, there exists C_ε such that

$$\mathbb{P}\{||\mathbf{Q}(t)|| \leq C_\varepsilon\} \geq 1-\varepsilon \qquad \text{for all } t \in \mathbb{N}_0.$$

Let E_t be the event that $(\mathbf{Q}(t), J(t)) \in \mathcal{E}$. Since there are only finitely many states in $\mathcal{T} \cap \{||\mathbf{Q}|| \leq C_\varepsilon\}$, we may apply Theorem 5 in [13] on page 389, and the reverse Fatou lemma to obtain

$$\mathbb{P}\{E_t\ i.o.\} \geq \limsup_{t\to\infty} \mathbb{P}\{||\mathbf{Q}(t)|| \leq C_\varepsilon\} \geq 1-\varepsilon.$$

Since $\varepsilon > 0$ is arbitrary, the above implies that an ergodic state is entered with probability 1.

The existence of a random variable Q_∞ with the stationary distribution follows directly from irreducibility and [14, Sect. XI.8, Theorem 1, p. 379]. Finally,

$$\mathbb{E}\left[||\mathbf{Q}_\infty||\right] \le \liminf \mathbb{E}\left[||\mathbf{Q}(t)||\right] \le E_{\max} < \infty$$

as in Theorem 5.3 in [3, p. 32], as the distribution of $\mathbf{Q}(t)$ converges weakly to a unique stationary distribution. □

Appendix 4: Proof of Corollary 2

For the sake of the proof, we may as well suppose that the sequence $S_k < \infty$ almost surely. Otherwise, the system reconfigures only finitely many times as it is 1-Harris recurrent, so that the bound given in (31) holds with $K_T = 1$. Hence we may suppose that there is a reconfiguration state $\mathbf{s} = (\mathbf{Q}, \mathbf{I}, \mathbf{Q}, \varrho_0)$ which occurs i.o. almost surely, by assumption. Call an occurrence of $\mathbf{s}$ a renewal. Now define T_M to be the number of slots used in taking M steps of the process $\left(\mathbf{Q}(\tau_m), \mathbf{I}(\tau_m), \mathbb{1}^{(m)}\right)$ starting from $m = 0$,

$$T_M = \sum_{m=0}^{M-1} \sum_{k=0}^{m} \chi_k \mathbb{1}^{(m,k)} + \sum_{m=0}^{M-1} \mathbb{1}^{(m)} \ge M,$$

because multiple slots are used after each reconfiguration.

It is obvious that

$$\sum_{k=0}^{m} \chi_k \mathbb{1}^{(m,k)} \le \sum_{k=0}^{m} \chi_k ||\mathbf{Q}(S_k)|| \mathbb{1}^{(m,k)} + \sum_{k=0}^{m} \chi_k \mathbb{1}^{(m,k)} \mathbb{I}\left[\mathbf{Q}(\tau_m) = 0\right]. \tag{30}$$

Taking expectations and bounding the last term in (30), we obtain

$$\begin{aligned}\sum_{k=0}^{m} \mathbb{E}\left[\chi_k; A_{m,k}, \mathbf{Q}(\tau_m) = 0\right] &= \sum_{k=0}^{m} \mathbb{E}\left[\chi_k; A_{m,k}, \mathbf{Q}(S_k) = 0\right] \\ &\le \sum_{k=0}^{m} \mathbb{E}\left[T_r + \sqrt{c_2} F(0); A_{m,k}\right] \\ &\le T_r + \sqrt{c_2} F(0) =: \mu_0.\end{aligned}$$

The equality holds since $A_{m,k} \in \mathcal{F}_{S_k}$ and since $\mathbf{Q}(S_k) = 0$ must have occurred. The first inequality follows on applying (12). The last inequality holds since the events $A_{m,k}$ are disjoint.

Taking expectations and turning to the first term in the right-hand side of (30), we may again apply (12), as the event $A_{m,k} \in \mathcal{F}_{S_k}$, to obtain

$$\mathbb{E}\left[||\mathbf{Q}(S_k)||\chi_k; A_{m,k}\right] \le c\mathbb{E}\left[F(||\mathbf{Q}(\tau_m)||)||\mathbf{Q}(\tau_m)||; A_{m,k}\right]$$

for some finite positive constant $c < \infty$. Combining the above two inequalities, it follows that

$$\limsup_{M\to\infty} \frac{1}{M}\mathbb{E}\left[T_M\right] \leq 1 + \mu_0 + \limsup_{M\to\infty} \frac{c}{M} \sum_{m=0}^{M-1} \sum_{k=0}^{m} \mathbb{E}\left[F(||\mathbf{Q}(\tau_m))||)||\mathbf{Q}(\tau_m)||;\, A_{m,k}\right].$$

Equation (18) now shows that

$$\limsup_{M\to\infty} \frac{1}{M}\mathbb{E}\left[T_M\right] = K_T < \infty. \tag{31}$$

Now let U_r be the number of slots between renewal r and renewal $r+1$ with a delay of V slots until the first renewal. The assumption of 1-Harris recurrence implies $V < \infty$. Let $\mu_R = \mathbb{E}\left[U_1\right] \leq \infty$ be the expected number of slots between renewals.

Let $R(M)$ be the number of renewals over the first M steps, so that

$$\sum_{r=1}^{R(M)-1} U_r \leq T_M. \tag{32}$$

Since $\mathbf{s}$ is an ergodic state,

$$\lim_{M\to\infty} \frac{R(M)}{M} = \phi_{\mathbf{s}},$$

with $\phi_{\mathbf{s}} \in (0, 1]$ by the renewal theorem. Moreover, by the extended Strong Law of Large Numbers,

$$\lim_{M\to\infty} \frac{1}{R(M)} \sum_{r=1}^{R(M)} U_r = \mu_R \text{ almost surely.}$$

Rewriting (32), we obtain

$$\frac{R(M)}{M} \times \frac{1}{R(M)} \sum_{r=1}^{R(M)} U_r \leq \frac{T_M}{M}.$$

Now let $M \to \infty$ along any sequence so that the term on the right-hand side converges to $\liminf_{M\to\infty} \frac{T_M}{M} \geq 1$, then

$$\phi_{\mathbf{s}}\mu_R \leq \liminf_{M\to\infty} \frac{T_M}{M}.$$

Using Fatou's lemma, we deduce

$$\mathbb{E}\left[\liminf_{M\to\infty} \frac{T_M}{M}\right] \leq \liminf_{M\to\infty} \mathbb{E}\left[\frac{T_M}{M}\right] \leq \limsup_{M\to\infty} \mathbb{E}\left[\frac{T_M}{M}\right] = K_T < \infty.$$

It follows that $\liminf_{M\to\infty} \frac{T_M}{M}$ is almost surely finite, and therefore $\mu_R < \infty$.

The 1-Harris recurrence of the chain $(\mathbf{Q}(t), \mathbf{I}(t), \mathbf{Q}(S_k(t)), \varrho(t))$ follows immediately as $\mathbf{s}$ has been shown to be an ergodic state within that chain. This completes the proof. □

Appendix 5: Proof of Lemma 5

To obtain a lower bound for χ_k, we consider the arrivals and departures separately, since the triangle inequality shows that the change in queue length cannot exceed their sum, at any stage. (Thus the sum can only exceed the BB threshold at χ_k or an earlier slot.) Moreover, we may count virtual departures (when queues are empty) because we are only concerned with a lower bound.

Define T_E to be the first time slot when the sum of arrivals and virtual departures, $\Sigma_j,\ j \geq T_r$,

$$\Sigma_j \doteq \sum_{s=0}^{j-1} ||\mathbf{A}(S_k+s)|| + \sum_{s=T_r}^{j} ||\mathbf{I}(S_k+s)||$$

reaches the threshold level $\theta_{BB} F\,(||\mathbf{Q}(S_k)||)$ after reconfiguration at S_k. $T_E \geq T_r$ is therefore a $\mathcal{F}_{S_k}$ stopping time. Let $M_j \doteq \Sigma_j - \mathbb{E}\left[\Sigma_j | \mathcal{F}_{S_k+j-1}\right],\ j \geq 1$, with $M_0 = 0$, so that M_j is a $\mathcal{F}_{S_k+j}$ martingale, null at 0, and bounded in $\mathcal{L}^1$. Clearly T_E is almost surely finite and has finite expectation, since all the summands are positive and have strictly positive expectation. It follows that condition (16) of Corollary 5 in [10, p. 243], holds. By the definition of T_E, at stopping it holds that

$$\sum_{s=0}^{T_E-1} ||\mathbf{A}(S_k+s)|| + \sum_{s=T_r}^{T_E} ||\mathbf{I}(S_k+s)|| \geq \theta_{BB} F\,(||\mathbf{Q}(S_k)||),$$

and from optional stopping we obtain

$$\mathbb{E}\left[\sum_{s=0}^{T_E-1} ||\mathbf{A}(S_k+s)|| + \sum_{s=T_r}^{T_E} ||\mathbf{I}(S_k+s)||\right] = \mathbb{E}\,[T_E] \sum_{\ell=1}^{N} \lambda_\ell + (\mathbb{E}\,[T_E] - T_r + 1) \sum_{\ell=1}^{N} \mathbf{I}_\ell.$$

Therefore,

$$\mathbb{E}\left[\chi_k | \mathcal{F}_{S_k}\right] \geq \mathbb{E}\left[T_E | \mathcal{F}_{S_k}\right] \geq \frac{\theta_{BB} F\,(||\mathbf{Q}(S_k)||)}{\sum_{\ell=1}^{N} (\lambda_\ell + \mathbf{I}_\ell)},$$

which yields the stated lower bound on $\mathbb{E}\left[\chi_k|\mathcal{F}_{S_k}\right]$ with $\delta = 0$ and

$$c_1 = \frac{\theta_{BB}}{||\lambda|| + N\mu_{\max}}.$$

□

Appendix 6: Proof of Lemma 6

For any arrival rate vector λ in the interior of the stability region Λ, there exists an $\epsilon > 0$ such that $\lambda + \epsilon \mathbf{1} \in \Lambda$. Therefore, the Max-Weight service vector $\mathbf{I}^*(S_k) = [\mathbf{I}_1^*, \ldots, \mathbf{I}_N^*]$ satisfies

$$\sum_{\ell=1}^{N} Q_\ell(S_k)\mathbf{I}_\ell^* \geq \mathbf{Q}(S_k) \cdot \lambda + \epsilon||\mathbf{Q}(S_k)||,$$

yielding

$$\sum_{\ell=1}^{N} Q_\ell(S_k)(\mathbf{I}_\ell^* - \lambda_\ell) \geq \epsilon \sum_{\ell=1}^{N} Q_\ell(S_k).$$

Since some queues might have negative contributions to the sum on the left-hand side, we have

$$\sum_{\ell:\mathbf{I}_\ell^* > \lambda_\ell} Q_\ell(S_k)(\mathbf{I}_\ell^* - \lambda_\ell) \geq \epsilon \sum_{\ell=1}^{N} Q_\ell(S_k).$$

Let ℓ_0 be the queue with the maximum contribution to the sum on the left-hand side. We have

$$Q_{\ell_0}(S_k) \geq \frac{\epsilon||\mathbf{Q}(S_k)||}{N\left(I_{\ell_0}^* - \lambda_{\ell_0}\right)}.$$

Since there are only finitely many possibilities for the strictly positive denominator, there is an $A_\epsilon > 0$ and an ℓ_0 such that

$$Q_{\ell_0}(S_k) \geq A_\epsilon||\mathbf{Q}(S_k)||$$

for every $\mathbf{Q}(S_k) \in \mathbb{N}_0^N$. Since the function $F(\cdot)$ is sublinear, it follows that we may take

$$\mathcal{C} = \{||\mathbf{Q}|| \leq D\}$$

for some finite positive constant D, so that $Q_{\ell_0} \geq \lceil\theta_{BB}F(||\mathbf{Q}||)\rceil$. Since it is also the case that the drift of queue ℓ_0 is strictly negative, the proof is complete. □

Appendix 7: Proof of Lemma 7

Let $D = I_{\ell^*} > \lambda_{\ell^*} = a$ be the number of packets that can be served at queue ℓ^* under the Max-Weight service vector in each time slot. (For compactness, we drop the queue index ℓ^* from the notation in the remainder of the proof.) Thus the queue length evolves according to the usual Lindley recursion

$$Q(t+1) = \max\{Q(t) - D + A(t), 0\}.$$

Let T_1 be the random amount of time for the queue length to reduce by at least one packet, assuming that the initial number of packets is D or larger. Thus T_1 is the so-called descending ladder index of the associated random walk. It is well known [14], pages 396–397, that $\mathbb{E}[T_1] < \infty$, since the random walk has strictly negative drift. If the number of arrivals per slot has finite variance $\sigma_A^2 < \infty$, it is reasonable to suppose that this implies $\sigma_{T_1}^2 = \mathbb{V}[T_1] < \infty$, and we now proceed to show that this is indeed the case.

Suppose we start a busy period with the service of $Q(1) = D$ packets. Arrivals may take place during this time slot, and so the busy period will continue under the recursion

$$Q(n+1) = [Q(n) - D + A(n)]^+, \quad n = 1, 2, \ldots,$$

and we stop at step $T_1 = n \in \mathbb{N}$ as soon as $[Q(n) - D + A(n)]^+ \in \{0, 1, \ldots, D-1\}$, which is the first time the number of packets falls by at least one. Further define $S(n) = \sum_{k=1}^{n} A(k)$ and $U(n) = S(n) - na$, which we may take as null for $n = 0$. By construction, $U(n)$ is a $\mathcal{L}_2$ martingale. At the stopping time T_1, the following equality holds:

$$S(T_1) = DT_1 - \varepsilon_{T_1},$$

where $\varepsilon_{T_1} \in \{1, \ldots, D\}$. This holds since full use has been made of the service in each time slot except for the slot following T_1.

Although not needed, it is the case that optional stopping holds for the martingale $U(n)$ with respect to T_1. This follows since $\mathbb{E}[|A(1) - a|] = M < \infty$ and the sequence is independent. Finally, since $\mathbb{E}[T_1] < \infty$, we find that

$$\mathbb{E}\left[\sum_{k=1}^{T_1} \mathbb{E}\left[|A(k) - a||\mathcal{F}_{k-1}\right]\right] = M\mathbb{E}[T_1] < \infty.$$

Hence, from [10, Chapter 7.4, Corollary 5, p. 243], we have $\mathbb{E}[U(T_1)] = 0$ which shows that

$$(D - a)\mathbb{E}[T_1] = \mathbb{E}\left[\varepsilon_{T_1}\right],$$

yielding $\mathbb{E}[T_1] = 1/(1-a)$ in the special case $D = 1$.

We also obtain from Theorem 7 of [10, p. 245], that

$$\mathbb{E}\left[U^2(T_1)\right] \leq \sigma_A^2 \mathbb{E}\left[T_1\right] = \mathbb{E}\left[\sum_{k=1}^{T_1}(A(k)-a)^2\right],$$

since the $A(k)$ are independent. Substituting, we find

$$(D-a)^2\mathbb{E}\left[T_1^2\right] - 2(D-a)\mathbb{E}\left[T_1\varepsilon_{T_1}\right] + \mathbb{E}\left[\varepsilon_{T_1}^2\right] \leq \sigma_A^2\mathbb{E}\left[T_1\right],$$

which shows that $\mathbb{V}\left[T_1\right] < \infty$. □

Appendix 8: Proof of Lemma 8

For compactness, define $M = \theta_{BB}F(||\mathbf{Q}(S_k)||)$, where we suppose that $M > 0$. Next suppose that we are outside the compact set $\mathcal{C}$ shown to exist in Lemma 6, but where the inequality on the right is multiplied by 3. Denote by ℓ^* the index of the queue as given in the lemma. It follows that $Q_{\ell^*} \geq 3M$. Let Y be the amount of time for Q_{ℓ^*} to decrease by $2M$, and as we have seen in Lemma 7 such a decrease is possible. The value $2M$ is considered, since Q_{ℓ^*} may increase by up to M packets without causing an immediate reconfiguration after T_r time slots. The behavior of Q_{ℓ^*} now follows a discrete-time random walk on $\mathbb{N}_0$. Clearly, $\chi_k \leq Y + T_r$, and Y is finite almost surely since Q_{ℓ^*} has strictly negative drift.

Let T_{ℓ^*} be the amount of time for Q_{ℓ^*} to reduce by at least one packet. Lemma 7 gives $\mathbb{E}\left[T_{\ell^*}^2\right] < \infty$.

The random variable Y is stochastically smaller than the sum of $2M$ i.i.d. copies $\tau_1, \ldots, \tau_{2M}$ of T_{ℓ^*}, since each descent is by at least one. The Cauchy–Schwarz inequality [34, p. 62], then implies

$$\mathbb{E}\left[Y^2\right] \leq \mathbb{E}\left[\left(\sum_{k=1}^{2M}\tau_k\right)^2\right] \leq 4M^2\mathbb{E}\left[T_{\ell^*}^2\right] < \infty.$$

Since $F(y) \uparrow \infty$ as $y \to \infty$ and given the definition of M, it follows that there is a constant c_2 such that (22) holds for all $\mathbf{Q}(S_k) \notin \mathcal{C}$.

If $\mathbf{Q}(S_k) \in \mathcal{C}$, then there exists a finite constant $A \geq \max_{\mathbf{Q}\in\mathcal{C}} F(||\mathbf{Q}||)$. Moreover, the set $\mathcal{C}_A$ of points which are within A of some point in $\mathcal{C}$ is also compact. Let $B = \sup_{\mathbf{Q}\in\mathcal{C}_A} ||\mathbf{Q}||$. Then if there are more than B arrivals in a given time slot, the stopping criterion must be triggered, if it has not been reached already. The time until B arrivals have occurred has a geometric distribution with finite first and second moments, since the distribution of the number of arrivals has unbounded support. It follows that (22) holds for all $\mathbf{Q} \in \mathbb{N}_0^N$ for c_2 sufficiently large. □

References

1. Akyildiz, I.F., Wang, X.: Wireless Mesh Networks. Wiley, New York (2009)
2. Armony, M., Bambos, N.: Queueing dynamics and maximal throughput scheduling in switched processing systems. Queueing Syst. **44**(3), 209–252 (2003)
3. Billingsley, P.: Convergence of Probability Measures. Wiley, New York (1968)
4. Blake, L., Long, M.: Antennas: Fundamentals, Design, Measurement. SciTech, Daryaganj (2009)
5. Borst, S.C., Boxma, O.J.: Polling models with and without switchover times. Oper. Res. **45**(4), 536–543 (1997)
6. Van den Broek, M.X., Van Leeuwaarden, J.S.H., Adan, I.J.B.F., Boxma, O.J.: Bounds and approximations for the fixed-cycle traffic-light queue. Transp. Sci. **40**(4), 484–496 (2006)
7. Brzezinski, A., Modiano, E.: Dynamic reconfiguration and routing algorithms for IP-over-WDM networks with stochastic traffic. IEEE J. Lightwave Technol. **23**(10), 3188–3205 (2005)
8. Çelik, G.D., Le, L.B., Modiano, E.: Scheduling in parallel queues with randomly varying connectivity and switchover delay. In: Proceedings of IEEE INFOCOM 2011 (Mini Conference) (April 2011)
9. Çelik, G.D., Borst, S.C., Whiting, P.A., Modiano, E.: Variable frame based Max-Weight algorithms for networks with switchover delay. In: Proceedings of IEEE ISIT 2011, St. Petersburg (August 2011)
10. Chow, Y., Teicher, H.: Probability Theory (Independence, Interchangeability and Martingales), 2nd edn. Springer, New York (1980)
11. Durrett, R.: Probability: Theory and Examples, 2nd edn. Duxbury Press, Belmont (1996)
12. Eryilmaz, A., Ozdaglar, A., Modiano, E.: Polynomial complexity algorithms for full utilization of multi-hop wireless networks. In: Proceedings of IEEE INFOCOM ’07 (May 2007)
13. Feller, W.: An Introduction to Probability Theory and Its Applications, vol. 1. Wiley, New York (1953)
14. Feller, W.: An Introduction to Probability Theory and Its Applications, vol. 2. Wiley, New York (1970)
15. Georgiadis, L., Neely, M., Tassiulas, L.: Resource Allocation and Cross-Layer Control in Wireless Networks. Now Publishers, Hanover (2006)
16. Hung, Y.-C., Chang, C.-C.: Dynamic scheduling for switched processing systems with substantial service-mode switching times. Queueing Syst. **60**, 87–109 (2008)
17. Meyn, S.P., Tweedie, R.L.: Markov Chains and Stochastic Stability. Cambridge University Press, Cambridge (2009)
18. Modiano, E., Barry, R.: A novel medium access control protocol for WDM-based LAN’s and access networks using a Master/Slave scheduler. IEEE J. Lightwave Technol. **18**(4), 461–468 (2000)
19. Modiano, E., Shah, D., Zussman, G.: Maximizing throughput in wireless networks via gossip. In: Proceedings of ACM SIGMETRICS/Performance ’06, June (2006)
20. Neely, M.J.: Stochastic Network Optimization with Application to Communication and Queueing Systems. Morgan and Claypool, San Rafael (2010)
21. Neely, M.J., Modiano, E., Rohrs, C.E.: Tradeoffs in delay guarantees and computation complexity in $n \times n$ packet switches. In: Proceedings of CISS ’02, Princeton, NJ (March 2002)
22. Neely, M.J., Modiano, E., Rohrs, C.E.: Power allocation and routing in multi-beam satellites with time-varying channels. IEEE Trans. Netw. **11**(1), 138–152 (2003)
23. Neely, M.J., Modiano, E., Rohrs, C.E.: Dynamic power allocation and routing for time varying wireless networks. IEEE J. Sel. Areas Commun. **23**(1), 89–103 (2005)
24. Le Ny, J., Dahleh, M., Feron, E.: Multi-UAV dynamic routing with partial observations using restless bandit allocation indices. In: Proceedings of ACC ’08 (June 2008)
25. de la Pena, V.H., Govindarajulu, Z.: A note on second moment of a randomly stopped sum of independent variables. Stat. Prob. Lett. **14**(4), 275–281 (1992)
26. Shah, D., Wischik, D.J.: Optimal scheduling algorithms for input-queued switches. In: Proceedings of IEEE INFOCOM ’06 (March 2006)
27. Stolyar, A.L.: Maxweight scheduling in a generalized switch: state space collapse and workload minimization in heavy traffic. Ann. Appl. Prob. **14**(1), 1–53 (2004)
28. Takagi, H.: Analysis of Polling Systems. MIT Press, Cambridge (1986)
29. Tassiulas, L.: Adaptive back-pressure congestion control based on local information. IEEE Trans. Autom. Control **40**(2), 236–250 (1995)
30. 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)

31. Tassiulas, L., Ephremides, A.: Dynamic server allocation to parallel queues with randomly varying connectivity. IEEE Trans. Inf. Theory **39**(2), 466–478 (1993)
32. Tassiulas, L., Papavassiliou, S.: Optimal anticipative scheduling with asynchronous transmission opportunities. IEEE Trans. Autom. Control **40**(12), 2052–2062 (1995)
33. Tolkachev, A., Denisenko, V., Shishlov, A., Shubov, A.: High gain antenna systems for millimeter wave radars with combined electronical and mechanical beam steering. In: Proceedings of IEEE Symposium on Phased Array Systems and Technology (October 2006)
34. Williams, D.: Probability with Martingales. Cambridge University Press, Cambridge (1989)
35. Wunderlich, R., Elhanany, I., Urbanik, T.: A stable longest queue first signal scheduling algorithm for an isolated intersection. In: IEEE Proceedings of Vehicular Electronics and Safety, 2007. ICVES (2007)
36. Yechiali, U.: Analysis and control of polling systems. Performance Evaluation of Computer and Communication Systems, vol. 729, pp. 630–650. Springer, Berlin (1993)
37. Ying, L., Shakkottai, S.: On throughput-optimality with delayed network-state information. In: Proc. ITA ’08, (January 2008)
38. Yun, M., Zhou, Y., Arora, A., Choi, H.: Channel-assignment and scheduling in wireless mesh networks considering switching overhead. In: Proc. IEEE ICC ’09, (June 2009)