

Qualitative Properties of α -Weighted Scheduling Policies

Devavrat Shah
MIT, LIDS
Cambridge, MA 02139
devavrat@mit.edu

John N. Tsitsiklis
MIT, LIDS
Cambridge, MA 02139
jnt@mit.edu

Yuan Zhong
MIT, LIDS
Cambridge, MA 02139
zhyu4118@mit.edu

ABSTRACT

We consider a switched network, a fairly general constrained queueing network model that has been used successfully to model the detailed packet-level dynamics in communication networks, such as input-queued switches and wireless networks. The main operational issue in this model is that of deciding which queues to serve, subject to certain constraints. In this paper, we study qualitative performance properties of the well known α -weighted scheduling policies. The stability, in the sense of positive recurrence, of these policies has been well understood. We establish exponential upper bounds on the tail of the steady-state distribution of the backlog. Along the way, we prove finiteness of the expected steady-state backlog when $\alpha < 1$, a property that was known only for $\alpha \geq 1$. Finally, we analyze the excursions of the maximum backlog over a finite time horizon for $\alpha \geq 1$. As a consequence, for $\alpha \geq 1$, we establish the full state space collapse property [17, 18].

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Packet Switching Networks; G.3 [Probability and Statistics]: Markov Processes, Queueing Theory, Stochastic Processes

General Terms

Algorithms, Performance, Theory

Keywords

Switched Network, Maximum Weight- α , Markov Chain, Exponential Bound, State Space Collapse

1. INTRODUCTION

This paper studies various qualitative stability and performance properties of the so-called α -weighted policies, as applied to a switched network model (cf.[22, 18]). This model is a special case of the “stochastic processing network model”

(cf.[10]), which has become the canonical framework for the study of a large class of networked queueing systems, including systems arising in communications, manufacturing, transportation, financial markets, etc. The primary reason for the popularity of the switched network model is its ability to faithfully model the behavior of a broad spectrum of networks at a fine granularity. Specifically, the switched network model is useful in describing packet-level (“micro”) behavior of medium access in a wireless network and of the input-queued switches that reside inside Internet routers. This model has proved tractable enough to allow for substantial progress in understanding the stability and performance properties of various control policies.

At a high level, the switched network model involves a collection of queues. Work arrives to these queues exogenously or from another queue and gets serviced; it then either leaves the network or gets re-routed to another queue. Service at the queues requires the use of some commonly shared constrained resources. This leads to the problem of *scheduling* the service of packets queued in the switched network. To utilize the network resources efficiently, a properly designed scheduling policy is required. Of particular interest are the popular Maximum Weight or MW- α policies, introduced in [22]. They are the only known simple and universally applicable policies with performance guarantees. In addition, the MW- α policy has served as an important guide for designing implementable algorithms for input-queued switches and wireless medium access (cf.[14, 21, 8, 7, 16]). This motivates the work in this paper, which focuses on certain qualitative properties of MW- α policies.

Related Prior Work.

Because of the significance of the α -weighted policies, there is a large body of research on their properties. We provide here a brief overview of the work that is most relevant to our purposes.

The most basic performance question concerns throughput and stability. Formally, we say that an algorithm is *throughput optimal* or *stable* if the underlying network Markov chain is positive recurrent whenever the system is *underloaded*. For the MW- α policy, under a general enough stochastic model, stability has been established for any $\alpha > 0$ (cf. [22, 15, 6, 1]).

A second, finer, performance question concerns the evaluation of the average backlog in the system, in steady-state. Bounds on the average backlog are usually obtained by considering the same stochastic Lyapunov function that was used to prove stability, and by building on the drift inequal-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGMETRICS'10, June 14–18, 2010, New York, New York, USA.
Copyright 2010 ACM 978-1-4503-0038-4/10/06 ...\$10.00.

ities established in the course of the stability proof; see, e.g., [5]. Using this approach, it is known that the average expected backlog under α -weighted policies is finite, when $\alpha \geq 1$ ([12]). However, such a result is not known when $\alpha \in (0, 1)$.

An important performance analysis method that has emerged over the past few decades focuses on the *heavy traffic* regime, in which the system is loaded near capacity. For the switched model, heavy traffic analysis has revealed some intriguing relations between the policy parameter α and the performance of the system through a phenomenon known as *state space collapse*. In particular, in the heavy traffic limit and for an appropriately scaled version of the system, the state evolves in a much lower-dimensional space (the state space “collapses”). The structure of the collapsed state space provides important information about the system behavior (cf. [11, 17, 18]). Under certain somewhat specific assumptions, a complete heavy traffic analysis of the switched network model has been carried out in [19, 4]. However, for the more general switched network model, only a weaker result is available, involving a so-called multiplicative state space collapse property [18, 17]. State space collapse results are related to understanding certain transient properties of the network, such as the evolution of the queues over a finite time horizon. To the best of our knowledge, a transient analysis of the switched network model is not available.

A somewhat different approach focuses on tail probabilities of the steady-state backlog and the associated large deviation principle (LDP). This approach provides important insights about the overflow probability in the presence of finite buffers. There have been notable works in this direction, for specific instances of the switched network model, e.g., [20]. In a similar setting, the reference [13] has also established a LDP for the MW-1 policy, using Garcia’s extended-contraction principle for quasi-continuous mappings. More recently, [23, 24] has announced a characterization of the precise tail behavior of the $(1 + \alpha)$ -norm of the backlog, under the MW- α policy. However, in these works, the LDP exponent is only given implicitly, as the solution of a complicated, possibly infinite dimensional optimization problem.

Our Contributions.

We establish various qualitative performance bounds for α -weighted policies, under the switched network model. In the stationary regime, we establish finiteness of the expected backlog, and an exponential upper bound on the steady-state tail probabilities of the backlog. In the transient regime, we establish a maximal inequality on the queue-size process, and the strong state space collapse property under α -weighted policies, when $\alpha \geq 1$. Our analysis is based on drift inequalities on suitable Lyapunov functions. Our methods, however, depart from prior work because they rely on different classes of Lyapunov functions, and also involve some new techniques.

In more detail, we begin by establishing the finiteness of the steady-state expected backlog under the MW- α policy, for any $\alpha \in (0, 1)$. Instead of the traditional Lyapunov function $\|\cdot\|_{\alpha+1}^{\alpha+1}$, we rely on a Lyapunov function which is a suitably smoothed version of $\|\cdot\|_{\alpha+1}^2$.

We continue by deriving a drift inequality for a “norm” or “norm-like” Lyapunov function, namely, $\|\cdot\|_{\alpha+1}$ or a suitably smoothed version. Using the drift inequality, we establish exponential tail bounds for the steady-state backlog

distribution under the MW- α policy, for any $\alpha \in (0, \infty)$. Our method builds on certain results from [3] that allow us to translate drift inequalities into closed-form tail bounds; it yields an explicit bound on the tail exponent, in terms of the system load and the total number of queues. This is in contrast with the earlier work in [20, 24]. That work provides an exact but implicit characterization of the tail exponents, in terms of a complicated optimization problem, and provides no immediate insights on the dependence of the tail exponents on the system parameters, such as the load and the number of queues. Furthermore, in contrast to the sophisticated mathematical techniques used in [20, 24], our explicit bounds are obtained through elementary methods. For some additional perspective, we also consider a special case and compare our upper bound with available lower bounds. This comparison is reported in the Appendix of the full version of this paper [25].

Finally, we provide a transient analysis under MW- α policies, for the case where $\alpha \geq 1$. We use a Lyapunov drift inequality to obtain a bound on the probability that the maximal backlog over a given finite time interval exceeds a certain threshold. This bound leads to the resolution of the strong state space collapse conjecture for the switched network model when $\alpha \geq 1$. This strengthens the multiplicative state space collapse results in [17, 18].

Organization of the Paper.

The rest of the paper is organized as follows. In Section 2, we define the notation we will employ, and describe the switched network model. In Section 3, we provide formal statements of our main results. In Section 4, we establish a drift inequality for a suitable Lyapunov function, which will be key to the proof of the exponential upper bound on tail probabilities. In Section 5, we prove the finiteness of steady-state expected backlog when $\alpha \in (0, 1)$. We prove the exponential upper bound in Section 6. The transient analysis is presented in Section 7. We start with a general lemma, and specialize it to obtain a maximal inequality under the MW- α policy, for $\alpha \geq 1$. We then apply the latter inequality to prove the full state space collapse result for $\alpha \geq 1$. We conclude the paper with a brief discussion in Section 8.

2. MODEL AND NOTATION

2.1 Notation

We introduce here the notation that will be employed throughout the paper. We denote the real vector space of dimension M by \mathbb{R}^M and the set of nonnegative M -tuples by \mathbb{R}_+^M . We write \mathbb{R} for \mathbb{R}^1 , and \mathbb{R}_+ for \mathbb{R}_+^1 . We let \mathbb{Z} be the set of integers, \mathbb{Z}_+ the set of nonnegative integers, and \mathbb{N} the set of positive integers.

For any vector $\mathbf{x} \in \mathbb{R}^M$, and any $\alpha > 0$, we define

$$\|\mathbf{x}\|_\alpha = \left(\sum_{i=1}^M |x_i|^\alpha \right)^{1/\alpha}.$$

For any two vectors $\mathbf{x} = (x_i)_{i=1}^M$ and $\mathbf{y} = (y_i)_{i=1}^M$ of the same dimension, we let $\mathbf{x} \cdot \mathbf{y} = \sum_{i=1}^M x_i y_i$ be the dot product of \mathbf{x} and \mathbf{y} . For two real numbers x and y , we let $x \vee y = \max\{x, y\}$. We also let $[x]^+ = x \vee 0$. We introduce the Kronecker delta symbol δ_{ij} , defined as $\delta_{ij} = 1$ if $i = j$, and

$\delta_{ij} = 0$ if $i \neq j$. We let $\mathbf{e}_i = (\delta_{ij})_{j=1}^M$ be the i -th unit vector in \mathbb{R}^M , and $\mathbf{1}$ the vector of all ones. For a set S , we denote its cardinality by $|S|$, and its indicator function by \mathbb{I}_S . For a matrix A , we let A^T denote its transpose. We will also use the abbreviations ‘‘RHS/LHS’’ for ‘‘right/left-hand side,’’ and ‘‘iff’’ for ‘‘if and only if.’’

2.2 Switched Network Model

The Model.

We adopt the model in [18], while restricting to the case of *single-hop* networks, for ease of exposition. However, our results naturally extend to *multi-hop* models, under the ‘‘back-pressure’’ variant of the MW- α policy.

Consider a collection of M queues. Let time be discrete: timeslot $\tau \in \{0, 1, \dots\}$ runs from time τ to $\tau + 1$. Let $Q_i(\tau)$ denote the (nonnegative integer) length of queue $i \in \{1, 2, \dots, M\}$ at the beginning of timeslot τ , and let $\mathbf{Q}(\tau)$ be the vector $(Q_i(\tau))_{i=1}^M$. Let $\mathbf{Q}(0)$ be the vector of initial queue lengths.

During each timeslot τ , the queue vector $\mathbf{Q}(\tau)$ is offered service described by a vector $\boldsymbol{\sigma}(\tau) = (\sigma_i(\tau))_{i=1}^M$ drawn from a given finite set $\mathcal{S} \subset \{0, 1\}^M$ of *feasible schedules*. Each queue $i \in \{1, 2, \dots, M\}$ has a dedicated exogenous arrival process $(A_i(\tau))_{\tau \geq 0}$, where $A_i(\tau)$ denotes the number of packets that arrive to queue i up to the beginning of timeslot τ , and $A_i(0) = 0$ for all i . We also let $a_i(\tau) = A_i(\tau+1) - A_i(\tau)$, which is the number of packets that arrive to queue i during timeslot τ . For simplicity, we assume that the $a_i(\cdot)$ are independent Bernoulli processes with parameter λ_i . We call $\boldsymbol{\lambda} = (\lambda_i)_{i=1}^M$ the *arrival rate vector*.

Given the service schedule $\boldsymbol{\sigma}(\tau) \in \mathcal{S}$ chosen at timeslot τ , the queues evolve according to the relation

$$Q_i(\tau+1) = [Q_i(\tau) - \sigma_i(\tau)]^+ + a_i(\tau).$$

In order to avoid trivialities, we assume, throughout the paper, the following.

ASSUMPTION 2.1. *For every queue i , there exists a $\boldsymbol{\sigma} \in \mathcal{S}$ such that $\sigma_i = 1$.*

An example: Input Queued (IQ) Switches.

The switched network model captures important instances of communication network scenarios (see [18] for various examples). Specifically, it faithfully models the packet-level operation of an input-queued (IQ) switch inside an Internet router. For an m -port IQ switch, it has m input and m output ports. It has a separate queue for each input-output pair (i, j) , denoted by Q_{ij} ,¹ for a total of $M = m^2$ queues. A schedule is required to match each input to exactly one output, and each output to exactly one input. Therefore, the set of schedules \mathcal{S} is

$$\left\{ \boldsymbol{\sigma} = (\sigma_{ij}) \in \{0, 1\}^{m \times m} : \sum_{k=1}^m \sigma_{ik} = \sum_{k=1}^m \sigma_{kj} = 1, \forall i, j \right\}.$$

We assume that the arrival process at each queue Q_{ij} is an independent Bernoulli process with mean λ_{ij} .

¹Here we deviate from our convention of indexing queues by a single subscript. This will ease exposition in the context of IQ switches, without causing confusion.

Capacity Region.

We define the capacity region $\boldsymbol{\Lambda}$ of a switched network as

$$\left\{ \boldsymbol{\lambda} \in \mathbb{R}_+^M : \boldsymbol{\lambda} \leq \sum_{\boldsymbol{\sigma} \in \mathcal{S}} \alpha_{\boldsymbol{\sigma}} \boldsymbol{\sigma}, \alpha_{\boldsymbol{\sigma}} \geq 0, \forall \boldsymbol{\sigma} \in \mathcal{S}, \sum_{\boldsymbol{\sigma} \in \mathcal{S}} \alpha_{\boldsymbol{\sigma}} < 1 \right\}.$$

It is called the capacity region because there exists a policy for which the Markov chain describing the network is positive recurrent iff $\boldsymbol{\lambda} \in \boldsymbol{\Lambda}$. We define the *load* induced by $\boldsymbol{\lambda} \in \boldsymbol{\Lambda}$, denoted by $\rho(\boldsymbol{\lambda})$, as

$$\rho(\boldsymbol{\lambda}) = \inf \left\{ \sum_{\boldsymbol{\sigma} \in \mathcal{S}} \alpha_{\boldsymbol{\sigma}} : \boldsymbol{\lambda} \leq \sum_{\boldsymbol{\sigma} \in \mathcal{S}} \alpha_{\boldsymbol{\sigma}} \boldsymbol{\sigma}, \alpha_{\boldsymbol{\sigma}} \geq 0, \forall \boldsymbol{\sigma} \in \mathcal{S} \right\}.$$

Note that $\rho(\boldsymbol{\lambda}) < 1$, for all $\boldsymbol{\lambda} \in \boldsymbol{\Lambda}$.

The Maximum-Weight- α Policy.

We now describe the so-called *Maximum-Weight- α* (MW- α) policy. For $\alpha > 0$, we use $\mathbf{Q}(\tau)^\alpha$ to denote the vector $(Q_i^\alpha(\tau))_{i=1}^M$. We define the *weight* of schedule $\boldsymbol{\sigma} \in \mathcal{S}$ to be $\boldsymbol{\sigma} \cdot \mathbf{Q}(\tau)^\alpha$. The MW- α policy chooses, at each timeslot τ , a schedule with the largest weight (breaking ties arbitrarily). Formally, during timeslot τ , the policy chooses a schedule $\boldsymbol{\sigma}(\tau)$ that satisfies

$$\boldsymbol{\sigma}(\tau) \cdot \mathbf{Q}(\tau)^\alpha = \max_{\boldsymbol{\sigma} \in \mathcal{S}} \boldsymbol{\sigma} \cdot \mathbf{Q}(\tau)^\alpha.$$

We define the maximum α -weight of the queue length vector \mathbf{Q} by $w_\alpha(\mathbf{Q}) = \max_{\boldsymbol{\sigma} \in \mathcal{S}} \boldsymbol{\sigma} \cdot \mathbf{Q}^\alpha$. When $\alpha = 1$, the policy is simply called the MW policy, and we use the notation $w(\mathbf{Q})$ instead of $w_1(\mathbf{Q})$. We take note of the fact that under the MW- α policy, the resulting Markov chain is known to be positive recurrent, for any $\boldsymbol{\lambda} \in \boldsymbol{\Lambda}$ (cf. [15]).

3. SUMMARY OF RESULTS

In this section, we summarize our main results for both the steady-state and the transient regime. The proofs are given in subsequent sections.

3.1 Stationary regime

The Markov chain $\mathbf{Q}(\cdot)$ that describes a switched network operating under the MW- α policy is known to be positive recurrent, as long as the system is underloaded, i.e., if $\boldsymbol{\lambda} \in \boldsymbol{\Lambda}$ or, equivalently, $\rho(\boldsymbol{\lambda}) < 1$. It is not hard to verify that this Markov chain is irreducible and aperiodic. Therefore, there exists a unique stationary distribution, which we will denote by $\boldsymbol{\pi}$. We use $\mathbb{E}_{\boldsymbol{\pi}}$ and $\mathbb{P}_{\boldsymbol{\pi}}$ to denote expectations and probabilities under $\boldsymbol{\pi}$.

Finiteness of Expected Queue-Size.

We establish that under the MW- α policy, the steady-state expected queue-size is finite, for any $\alpha \in (0, 1)$. (Recall that this result is already known when $\alpha \geq 1$.)

THEOREM 3.1. *Consider a switched network operating under the MW- α policy with $\alpha \in (0, 1)$, and assume that $\rho(\boldsymbol{\lambda}) < 1$. Then, the steady-state expected queue-size is finite, i.e.,*

$$\mathbb{E}_{\boldsymbol{\pi}} [\|\mathbf{Q}\|_1] < \infty.$$

Exponential Bound on Tail Probabilities.

For the MW- α policy, and for any $\alpha \in (0, \infty)$, we obtain an explicit exponential upper bound on the tail probabilities

of the queue-size, in steady-state. Our result involves two constants defined by

$$\bar{\nu} = \mathbb{E}[\|\mathbf{a}(1)\|_{\alpha+1}], \quad \gamma = \frac{1-\rho}{2M^{\frac{1}{\alpha+1}}},$$

where $\rho = \rho(\boldsymbol{\lambda})$.

THEOREM 3.2. *Consider a switched network operating under the MW- α policy, and assume that $\rho = \rho(\boldsymbol{\lambda}) < 1$. There exist positive constants B and B' such that for all $\ell \in \mathbb{Z}_+$:*

(a) *if $\alpha \geq 1$, then*

$$\mathbb{P}_\pi \left(\|\mathbf{Q}(\tau)\|_{\alpha+1} > B + 2M^{\frac{1}{\alpha+1}} \ell \right) \leq \left(\frac{\bar{\nu}}{\bar{\nu} + \gamma} \right)^{\ell+1};$$

(b) *if $\alpha \in (0, 1)$, then*

$$\mathbb{P}_\pi \left(\|\mathbf{Q}(\tau)\|_{\alpha+1} > B' + 10M^{\frac{1}{\alpha+1}} \ell \right) \leq \left(\frac{5\bar{\nu}}{5\bar{\nu} + \gamma} \right)^{\ell+1}.$$

Note that Theorem 3.1 could be obtained as a simple corollary of Theorem 3.2. On the other hand, our proof of Theorem 3.2 requires the finiteness of $\mathbb{E}_\pi[\|\mathbf{Q}\|_1]$, and so Theorem 3.1 needs to be established first.

In the Appendix of the full version of this paper [25], we comment on the tightness of our upper bounds by comparing them with explicit lower bounds that follow from the recent large deviations results in [24].

3.2 Transient regime

Here we provide a simple inequality on the maximal excursion of the queue-size over a finite time interval, under the MW- α policy, with $\alpha \geq 1$.

THEOREM 3.3. *Consider a switched network operating under the MW- α policy with $\alpha \geq 1$, and assume that $\rho(\boldsymbol{\lambda}) < 1$. Suppose that $\mathbf{Q}(0) = \mathbf{0}$. Let $Q_{\max}(\tau) = \max_{i \in \{1, \dots, M\}} Q_i(\tau)$, and $Q_{\max}^*(T) = \max_{\tau \in \{0, 1, \dots, T\}} Q_{\max}(\tau)$. Then, for any $b > 0$,*

$$\mathbb{P}(Q_{\max}^*(T) \geq b) \leq \frac{K(\alpha, M)T}{(1-\rho)^{\alpha-1} b^{\alpha+1}}, \quad (1)$$

for some positive constant $K(\alpha, M)$ depending only on α and M .

As an important application, we use Theorem 3.3 to prove a full state space collapse result,² for $\alpha \geq 1$, in Section 7.3. The precise statement can be found in Theorem 7.7.

4. MW- α POLICIES: A USEFUL DRIFT INEQUALITY

The key to many of our results is a *drift inequality* that holds for every $\alpha > 0$ and $\boldsymbol{\lambda} \in \boldsymbol{\Lambda}$. In this section, we shall state and prove this inequality. It will be used in Section 6 to prove Theorem 3.2. We remark that similar drift inequalities, but for a different Lyapunov function, have played an important role in establishing positive recurrence (cf. [22]) and multiplicative state space collapse (cf. [18]).

We will be making extensive use of a second-order mean value theorem [2], which we state below for easy reference.

²This is strong state space collapse and not full diffusion approximation.

PROPOSITION 4.1. *Let $g : \mathbb{R}^M \rightarrow \mathbb{R}$ be twice continuously differentiable over an open sphere S centered at a vector \mathbf{x} . Then, for any \mathbf{y} such that $\mathbf{x} + \mathbf{y} \in S$, there exists a $\theta \in [0, 1]$ such that*

$$g(\mathbf{x} + \mathbf{y}) = g(\mathbf{x}) + \mathbf{y}^T \nabla g(\mathbf{x}) + \frac{1}{2} \mathbf{y}^T H(\mathbf{x} + \theta \mathbf{y}) \mathbf{y}, \quad (2)$$

where $\nabla g(\mathbf{x})$ is the gradient of g at \mathbf{x} , and $H(\mathbf{x})$ is the Hessian of the function g at \mathbf{x} .

We now define the Lyapunov function that we will employ. For $\alpha \geq 1$, it will be simply the $(\alpha + 1)$ -norm $\|\mathbf{x}\|_{1+\alpha}$ of a vector \mathbf{x} . However, when $\alpha \in (0, 1)$, this function has unbounded second derivatives as we approach the boundary of \mathbb{R}_+^M . For this reason, our Lyapunov function will be a suitably smoothed version of $\|\cdot\|_{\alpha+1}$.

DEFINITION 4.2. *Define $f_\alpha : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ to be $f_\alpha(r) = r^\alpha$, when $\alpha \geq 1$, and*

$$f_\alpha(r) = \begin{cases} r^\alpha, & \text{if } r \geq 1, \\ (\alpha - 1)r^3 + (1 - \alpha)r^2 + r, & \text{if } r \leq 1, \end{cases}$$

when $\alpha \in (0, 1)$. Let $F_\alpha : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ be the antiderivative of f_α , so that $F_\alpha(r) = \int_0^r f_\alpha(s) ds$. The Lyapunov function $L_\alpha : \mathbb{R}_+^M \rightarrow \mathbb{R}_+$ is defined to be

$$L_\alpha(\mathbf{x}) = \left[(\alpha + 1) \sum_{i=1}^M F_\alpha(x_i) \right]^{\frac{1}{\alpha+1}}.$$

We will make heavy use of various properties of the functions f_α , F_α , and L_α , which we summarize in the following lemma. The proof is elementary and is omitted.

LEMMA 4.3. *Let $\alpha \in (0, 1)$. The function f_α has the following properties:*

- (i) *it is continuously differentiable with $f_\alpha(0) = 0$, $f_\alpha(1) = 1$, $f'_\alpha(0) = 1$, and $f'_\alpha(1) = \alpha$;*
- (ii) *it is increasing and, in particular, $f_\alpha(r) \geq 0$ for all $r \geq 0$;*
- (iii) *we have $r^\alpha - 1 \leq f_\alpha(r) \leq r^\alpha + 1$, for all $r \in [0, 1]$;*
- (iv) *$f'_\alpha(r) \leq 2$, for all $r \geq 0$.*

Furthermore, from (iii), we also have the following property of F_α :

- (iii') *$r^{\alpha+1} - 2 \leq (\alpha + 1)F_\alpha(r) \leq r^{\alpha+1} + 2$ for all $r \geq 0$.*

We are now ready to state the drift inequality.

THEOREM 4.4. *Consider a switched network operating under the MW- α policy, and assume that $\rho = \rho(\boldsymbol{\lambda}) < 1$. Then, there exists a constant $B > 0$, such that if $L_\alpha(\mathbf{Q}(\tau)) > B$, then*

$$\mathbb{E}[L_\alpha(\mathbf{Q}(\tau+1)) - L_\alpha(\mathbf{Q}(\tau)) \mid \mathbf{Q}(\tau)] \leq -\frac{1-\rho}{2} M^{\frac{1}{\alpha+1}-1}. \quad (3)$$

The proof of this drift inequality is quite tedious when $\alpha \neq 1$. To make the proof more accessible and to provide intuition, we first present the somewhat simpler proof for $\alpha = 1$. We then provide the proof for the case of general α , by considering separately the two cases where $\alpha > 1$ and $\alpha \in (0, 1)$.

We wish to draw attention here to the main difference from related drift inequalities in the literature. The usual

proof of stability involves the Lyapunov function $\|\mathbf{Q}\|_{\alpha+1}^{\alpha+1}$; for instance, for the standard MW policy, it involves a quadratic Lyapunov function. In contrast, we use $\|\mathbf{Q}\|_{\alpha+1}$ (or its smoothed version), which scales linearly along radial directions. In this sense, our approach is similar in spirit to [3], which employed piecewise linear Lyapunov functions to derive drift inequalities and then moment and tail bounds.

4.1 Proof of Theorem 4.4: $\alpha = 1$

In this section, we assume that $\alpha = 1$. As remarked earlier, we have $L_\alpha(\mathbf{x}) = \|\mathbf{x}\|_2$.

Suppose that $\|\mathbf{Q}(\tau)\|_2 > 0$. We claim that on every sample path, we have

$$\|\mathbf{Q}(\tau+1)\|_2 - \|\mathbf{Q}(\tau)\|_2 \leq \frac{\mathbf{Q}(\tau) \cdot \boldsymbol{\delta}(\tau) + \|\boldsymbol{\delta}(\tau)\|_2^2}{\|\mathbf{Q}(\tau)\|_2}, \quad (4)$$

where $\boldsymbol{\delta}(\tau) = \mathbf{Q}(\tau+1) - \mathbf{Q}(\tau)$. To see this, we proceed as follows. We have

$$\begin{aligned} & \left(\|\mathbf{Q}(\tau)\|_2 + \frac{\mathbf{Q}(\tau) \cdot \boldsymbol{\delta}(\tau) + \|\boldsymbol{\delta}(\tau)\|_2^2}{\|\mathbf{Q}(\tau)\|_2} \right)^2 \\ & \geq \|\mathbf{Q}(\tau)\|_2^2 + 2(\mathbf{Q}(\tau) \cdot \boldsymbol{\delta}(\tau) + \|\boldsymbol{\delta}(\tau)\|_2^2) \\ & \geq \|\mathbf{Q}(\tau)\|_2^2 + 2\mathbf{Q}(\tau) \cdot \boldsymbol{\delta}(\tau) + \|\boldsymbol{\delta}(\tau)\|_2^2 \\ & = \|\mathbf{Q}(\tau) + \boldsymbol{\delta}(\tau)\|_2^2 = \|\mathbf{Q}(\tau+1)\|_2^2. \end{aligned} \quad (5)$$

Note that

$$\begin{aligned} & \|\mathbf{Q}(\tau)\|_2^2 + \mathbf{Q}(\tau) \cdot \boldsymbol{\delta}(\tau) + \|\boldsymbol{\delta}(\tau)\|_2^2 \\ & = \left\| \mathbf{Q}(\tau) + \frac{\boldsymbol{\delta}(\tau)}{2} \right\|_2^2 + \frac{3}{4}\|\boldsymbol{\delta}(\tau)\|_2^2 \geq 0. \end{aligned}$$

We divide by $\|\mathbf{Q}(\tau)\|_2$, to obtain

$$\|\mathbf{Q}(\tau)\|_2 + \frac{\mathbf{Q}(\tau) \cdot \boldsymbol{\delta}(\tau) + \|\boldsymbol{\delta}(\tau)\|_2^2}{\|\mathbf{Q}(\tau)\|_2} \geq 0.$$

Therefore, we can take square roots of both sides of (5), without reversing the direction of the inequality, and the claimed inequality (4) follows.

Recall that $|\delta_i(\tau)| \leq 1$, because of the Bernoulli arrival assumption. It follows that $\|\boldsymbol{\delta}(\tau)\|_2 \leq M^{1/2}$. We now take the conditional expectation of both sides of (4). We have

$$\begin{aligned} & \mathbb{E} \left[\|\mathbf{Q}(\tau+1)\|_2 - \|\mathbf{Q}(\tau)\|_2 \mid \mathbf{Q}(\tau) \right] \\ & \leq \mathbb{E} \left[\frac{\mathbf{Q}(\tau) \cdot \mathbf{a}(\tau) - \mathbf{Q}(\tau) \cdot \boldsymbol{\sigma}(\tau) + M}{\|\mathbf{Q}(\tau)\|_2} \mid \mathbf{Q}(\tau) \right] \\ & = \frac{\sum_{i=1}^M Q_i(\tau) \mathbb{E}[a_i(\tau)] - \mathbf{Q}(\tau) \cdot \boldsymbol{\sigma}(\tau) + M}{\|\mathbf{Q}(\tau)\|_2} \\ & = \frac{\sum_{i=1}^M Q_i(\tau) \lambda_i - w(\mathbf{Q}(\tau)) + M}{\|\mathbf{Q}(\tau)\|_2} \\ & \leq \frac{M - (1-\rho)w(\mathbf{Q}(\tau))}{\|\mathbf{Q}(\tau)\|_2}. \end{aligned} \quad (6)$$

The last inequality above is justified as follows. From the definition of $\rho = \rho(\boldsymbol{\lambda})$, there exist constants $\alpha_\sigma \geq 0$ such that $\sum_{\sigma \in \mathcal{S}} \alpha_\sigma \leq \rho$, and

$$\boldsymbol{\lambda} \leq \sum_{\sigma \in \mathcal{S}} \alpha_\sigma \boldsymbol{\sigma}. \quad (7)$$

Therefore,

$$\begin{aligned} \sum_i Q_i(\tau) \lambda_i & = \mathbf{Q}(\tau) \cdot \boldsymbol{\lambda} \leq \sum_{\sigma \in \mathcal{S}} \alpha_\sigma \mathbf{Q}(\tau) \cdot \boldsymbol{\sigma} \\ & \leq \sum_{\sigma \in \mathcal{S}} \alpha_\sigma w(\mathbf{Q}(\tau)) \leq \rho w(\mathbf{Q}(\tau)). \end{aligned} \quad (8)$$

Let $Q_{\max}(\tau) = \max_{i=1}^M Q_i(\tau)$. Then,

$$\|\mathbf{Q}(\tau)\|_2 \leq (MQ_{\max}^2(\tau))^{\frac{1}{2}} = M^{\frac{1}{2}} Q_{\max}(\tau).$$

From Assumption 2.1, we have

$$w(\mathbf{Q}(\tau)) \geq Q_{\max}(\tau).$$

Therefore, the RHS of (6) can be upper bounded by

$$-(1-\rho)M^{-1/2} + \frac{M}{\|\mathbf{Q}(\tau)\|_2} \leq -\frac{1}{2}(1-\rho)M^{-1/2},$$

when $\|\mathbf{Q}(\tau)\|_2$ is sufficiently large.

4.2 Proof of Theorem 4.4: $\alpha > 1$

We wish to obtain an inequality similar to (6) for $L_\alpha(\mathbf{Q}(\cdot)) = \|\mathbf{Q}(\cdot)\|_{1+\alpha}$ under the MW- α policy, and we accomplish this using the second-order mean value theorem (cf. Proposition 4.1). Throughout this proof, we will drop the subscript $\alpha+1$ and use the notation $\|\cdot\|$ instead of $\|\cdot\|_{\alpha+1}$.

Consider the norm function

$$g(\mathbf{x}) = \|\mathbf{x}\| = (x_1^{\alpha+1} + \dots + x_M^{\alpha+1})^{\frac{1}{\alpha+1}}.$$

The first derivative is

$$\nabla g(\mathbf{x}) = \|\mathbf{x}\|^{-\alpha} (x_1^\alpha, \dots, x_M^\alpha) = \frac{\mathbf{x}^\alpha}{\|\mathbf{x}\|^\alpha}.$$

Let $H(\mathbf{x}) = [H_{ij}(\mathbf{x})]_{i,j=1}^M$ be the second derivative (Hessian) matrix of g . Then,

$$H_{ij}(\mathbf{x}) = \frac{\partial^2 g}{\partial x_i \partial x_j}(\mathbf{x}) = \delta_{ij} \frac{\alpha x_i^{\alpha-1}}{\|\mathbf{x}\|^\alpha} - \frac{\alpha x_i^\alpha x_j^\alpha}{\|\mathbf{x}\|^{2\alpha+1}},$$

where δ_{ij} is the Kronecker delta. By Proposition 4.1, for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}_+^M$, and with $\boldsymbol{\delta} = \mathbf{y} - \mathbf{x}$, there exists a $\theta \in [0, 1]$ for which

$$\begin{aligned} g(\mathbf{y}) & = g(\mathbf{x}) + \boldsymbol{\delta}^T \nabla g(\mathbf{x}) + \frac{1}{2} \boldsymbol{\delta}^T H(\mathbf{x} + \theta \boldsymbol{\delta}) \boldsymbol{\delta} \\ & = g(\mathbf{x}) + \|\mathbf{x}\|^{-\alpha} \left(\sum_i \delta_i x_i^\alpha \right) \\ & \quad + \frac{\alpha}{2} \|\mathbf{x} + \theta \boldsymbol{\delta}\|^{-\alpha} \left(\sum_i (x_i + \theta \delta_i)^{\alpha-1} \delta_i^2 \right) \\ & \quad - \frac{\alpha}{2} \|\mathbf{x} + \theta \boldsymbol{\delta}\|^{-1-2\alpha} \left(\sum_{i,j} (x_i + \theta \delta_i)^\alpha (x_j + \theta \delta_j)^\alpha \delta_i \delta_j \right) \\ & = g(\mathbf{x}) + \|\mathbf{x}\|^{-\alpha} \left(\sum_i \delta_i x_i^\alpha \right) \\ & \quad + \frac{\alpha}{2} \|\mathbf{x} + \theta \boldsymbol{\delta}\|^{-\alpha} \left(\sum_i (x_i + \theta \delta_i)^{\alpha-1} \delta_i^2 \right) \\ & \quad - \frac{\alpha}{2} \|\mathbf{x} + \theta \boldsymbol{\delta}\|^{-1-2\alpha} \left(\sum_i (x_i + \theta \delta_i)^\alpha \delta_i \right)^2. \end{aligned}$$

Using $\mathbf{x} = \mathbf{Q}(\tau)$, $\mathbf{y} = \mathbf{Q}(\tau+1)$ and $\boldsymbol{\delta}(\tau) = \mathbf{Q}(\tau+1) - \mathbf{Q}(\tau)$, we have

$$\begin{aligned} \|\mathbf{Q}(\tau+1)\| &= \|\mathbf{Q}(\tau)\| + \left[\frac{\sum_i \delta_i(\tau) Q_i^\alpha(\tau)}{\|\mathbf{Q}(\tau)\|^\alpha} \right] \\ &\quad + \frac{\alpha}{2} \left[\frac{\sum_i (Q_i(\tau) + \theta \delta_i(\tau))^{\alpha-1} \delta_i^2(\tau)}{\|\mathbf{Q}(\tau) + \theta \boldsymbol{\delta}(\tau)\|^\alpha} \right] \\ &\quad - \frac{\alpha}{2} \left[\frac{(\sum_i (Q_i(\tau) + \theta \delta_i(\tau))^\alpha \delta_i(\tau))^2}{\|\mathbf{Q}(\tau) + \theta \boldsymbol{\delta}(\tau)\|^{1+2\alpha}} \right]. \end{aligned} \quad (9)$$

Therefore, using the fact that $\delta_i(\tau) \in \{-1, 0, 1\}$, we have

$$\begin{aligned} &\|\mathbf{Q}(\tau+1)\| - \|\mathbf{Q}(\tau)\| \\ &\leq \left[\frac{\sum_i \delta_i(\tau) Q_i^\alpha(\tau)}{\|\mathbf{Q}(\tau)\|^\alpha} \right] + \frac{\alpha}{2} \left[\frac{\sum_i (Q_i(\tau) + \theta \delta_i(\tau))^{\alpha-1}}{\|\mathbf{Q}(\tau) + \theta \boldsymbol{\delta}(\tau)\|^\alpha} \right]. \end{aligned} \quad (10)$$

We take conditional expectations of both sides, given $\mathbf{Q}(\tau)$. To bound the first term on the RHS, we use the definition of the MW- α policy, the bound (7) on $\boldsymbol{\lambda}$, and the argument used to establish (8) in the proof of Theorem 4.4 for $\alpha = 1$ (with $w(\mathbf{Q}(\tau))$ replaced by $w_\alpha(\mathbf{Q}(\tau))$). We obtain

$$\mathbb{E} \left[\frac{\sum_i \delta_i(\tau) Q_i^\alpha(\tau)}{\|\mathbf{Q}(\tau)\|^\alpha} \mid \mathbf{Q}(\tau) \right] \leq -(1-\rho) \frac{w_\alpha(\mathbf{Q}(\tau))}{\|\mathbf{Q}(\tau)\|^\alpha}. \quad (11)$$

Note that

$$\begin{aligned} \|\mathbf{Q}(\tau)\|^\alpha &\leq (M Q_{\max}(\tau)^{\alpha+1})^{\frac{\alpha}{\alpha+1}} \\ &= M^{\frac{\alpha}{\alpha+1}} Q_{\max}^\alpha(\tau), \end{aligned} \quad (12)$$

and

$$w_\alpha(\mathbf{Q}(\tau)) \geq Q_{\max}^\alpha(\tau).$$

Therefore,

$$\mathbb{E} \left[\frac{\sum_i \delta_i(\tau) Q_i^\alpha(\tau)}{\|\mathbf{Q}(\tau)\|^\alpha} \mid \mathbf{Q}(\tau) \right] \leq -(1-\rho) M^{-\frac{\alpha}{1+\alpha}}. \quad (13)$$

Consider now the second term of the conditional expectation of the RHS of Inequality (10). Since $\alpha > 1$, and $\delta_i(\tau) \in \{-1, 0, 1\}$, the numerator of the expression inside the bracket satisfies

$$\sum_i (Q_i(\tau) + \theta \delta_i(\tau))^{\alpha-1} \leq M (Q_{\max}(\tau) + 1)^{\alpha-1},$$

and the denominator satisfies

$$\|\mathbf{Q}(\tau) + \theta \boldsymbol{\delta}(\tau)\|^\alpha \geq ([Q_{\max}(\tau) - 1]^+)^{\alpha},$$

where we use the notation $[c]^+ = 0 \vee c$. Thus,

$$\frac{\alpha}{2} \left[\frac{\sum_i (Q_i(\tau) + \theta \delta_i(\tau))^{\alpha-1}}{\|\mathbf{Q}(\tau) + \theta \boldsymbol{\delta}(\tau)\|^\alpha} \right] \leq \frac{\alpha}{2} \cdot \frac{M(Q_{\max} + 1)^{\alpha-1}}{([Q_{\max}(\tau) - 1]^+)^{\alpha}}.$$

Now if $\|\mathbf{Q}(\tau)\|$ is large enough, $Q_{\max}(\tau)$ is large enough, and $\frac{\alpha}{2} \cdot \frac{M(Q_{\max} + 1)^{\alpha-1}}{([Q_{\max}(\tau) - 1]^+)^{\alpha}}$ can be made arbitrarily small. Thus, the conditional expectation of the second term on the RHS of (10) can be made arbitrarily small for large enough $\|\mathbf{Q}(\tau)\|$. This fact, together with Inequality (13), implies that there exists $B > 0$ such that if $\|\mathbf{Q}(\tau)\| > B$, then

$$\mathbb{E} \left[\|\mathbf{Q}(\tau+1)\| - \|\mathbf{Q}(\tau)\| \mid \mathbf{Q}(\tau) \right] \leq -\frac{1-\rho}{2} M^{-\frac{\alpha}{1+\alpha}}.$$

4.3 Proof of Theorem 4.4: $\alpha \in (0, 1)$

The proof in this section is similar to that for the case $\alpha > 1$. We invoke Proposition 4.1 to write the drift term as a sum of terms, which we bound separately. Note that to use Proposition 4.1, we need L_α to be twice continuously differentiable. Indeed, by Lemma 4.3 (i), f_α is continuously differentiable, so its antiderivative F_α is twice continuously differentiable, and so is L_α . Thus, by the second order mean value theorem, we obtain an equation similar to Equation (9):

$$\begin{aligned} &L_\alpha(\mathbf{Q}(\tau+1)) - L_\alpha(\mathbf{Q}(\tau)) \\ &= \left[\frac{\sum_i \delta_i(\tau) f_\alpha(Q_i(\tau))}{L_\alpha^\alpha(\mathbf{Q}(\tau))} \right] + \frac{1}{2} \left[\frac{\sum_i f_\alpha'(Q_i(\tau) + \theta \delta_i(\tau)) \delta_i^2(\tau)}{L_\alpha^\alpha(\mathbf{Q}(\tau) + \theta \boldsymbol{\delta}(\tau))} \right] \\ &\quad - \frac{\alpha}{2} \left[\frac{(\sum_i \delta_i(\tau) f_\alpha(Q_i(\tau) + \theta \delta_i(\tau)))^2}{L_\alpha^{2\alpha+1}(\mathbf{Q}(\tau) + \theta \boldsymbol{\delta}(\tau))} \right]. \end{aligned} \quad (14)$$

Again, using the fact $\delta_i(\tau) \in \{-1, 0, 1\}$,

$$L_\alpha(\mathbf{Q}(\tau+1)) - L_\alpha(\mathbf{Q}(\tau)) \leq T_1 + T_2,$$

where

$$T_1 = \frac{\sum_i \delta_i(\tau) f_\alpha(Q_i(\tau))}{L_\alpha^\alpha(\mathbf{Q}(\tau))},$$

and

$$T_2 = \frac{1}{2} \left[\frac{\sum_i f_\alpha'(Q_i(\tau) + \theta \delta_i(\tau))}{L_\alpha^\alpha(\mathbf{Q}(\tau) + \theta \boldsymbol{\delta}(\tau))} \right].$$

Let us consider T_2 first. For $\alpha \in (0, 1)$, by Lemma 4.3 (iv), $f_\alpha'(r) \leq 2$ for all $r \geq 0$. Thus

$$T_2 \leq \frac{1}{2} \left[\frac{2M}{L_\alpha^\alpha(\mathbf{Q}(\tau) + \theta \boldsymbol{\delta}(\tau))} \right] = \frac{M}{L_\alpha^\alpha(\mathbf{Q}(\tau) + \theta \boldsymbol{\delta}(\tau))},$$

which becomes arbitrarily small when $L_\alpha(\mathbf{Q}(\tau))$ is large enough.

We now consider T_1 . Since $f_\alpha(r) \leq r^\alpha + 1$ for all $r \geq 0$ (cf. Lemma 4.3 (iii)), and $\delta_i(\tau) \in \{-1, 0, 1\}$,

$$T_1 \leq \frac{\sum_i \delta_i(\tau) Q_i^\alpha(\tau)}{L_\alpha^\alpha(\mathbf{Q}(\tau))} + \frac{M}{L_\alpha^\alpha(\mathbf{Q}(\tau))}.$$

When we take the conditional expectation, an argument similar to the one for the case $\alpha > 1$ yields

$$\mathbb{E} \left[\frac{\sum_i \delta_i(\tau) Q_i^\alpha(\tau)}{L_\alpha^\alpha(\mathbf{Q}(\tau))} \mid \mathbf{Q}(\tau) \right] \leq -(1-\rho) \frac{w_\alpha(\mathbf{Q}(\tau))}{L_\alpha^\alpha(\mathbf{Q}(\tau))}. \quad (15)$$

Again, as before, $w_\alpha(\mathbf{Q}(\tau)) \geq Q_{\max}^\alpha(\tau)$. For the denominator, by Lemma 4.3 (iii'), for any $r \geq 0$, we have $(\alpha + 1)F_\alpha(r) \leq r^{\alpha+1} + 2$. Thus

$$\begin{aligned} L_\alpha(\mathbf{Q}(\tau)) &\leq \left[\sum_i (Q_i(\tau) + 2)^{\alpha+1} \right]^{\frac{1}{\alpha+1}} \\ &\leq (M(Q_{\max}(\tau) + 2)^{\alpha+1})^{\frac{1}{\alpha+1}} \\ &= M^{\frac{1}{\alpha+1}} (Q_{\max}(\tau) + 2). \end{aligned}$$

Therefore,

$$\mathbb{E} \left[\frac{\sum_i \delta_i(\tau) Q_i^\alpha(\tau)}{L_\alpha^\alpha(\mathbf{Q}(\tau))} \mid \mathbf{Q}(\tau) \right] \leq -(1-\rho) M^{-\frac{\alpha}{\alpha+1}} \frac{Q_{\max}^\alpha(\tau)}{(Q_{\max} + 2)^\alpha}.$$

If $Q_{\max}(\tau)$ is large enough, we can further upper bound the RHS by, say, $-\frac{3}{4}(1-\rho)M^{-\frac{\alpha}{\alpha+1}}$.

Putting everything together, we have

$$\begin{aligned} & \mathbb{E} \left[L_\alpha(\mathbf{Q}(\tau+1)) - L_\alpha(\mathbf{Q}(\tau)) \mid \mathbf{Q}(\tau) \right] \\ & \leq -\frac{3}{4}(1-\rho)M^{-\frac{\alpha}{1+\alpha}} + \frac{M}{L_\alpha^\alpha(\mathbf{Q}(\tau))} + \mathbb{E}[T_2 \mid \mathbf{Q}(\tau)], \quad (16) \end{aligned}$$

if $Q_{\max}(\tau)$ is large enough. As before, if $L_\alpha(\mathbf{Q}(\tau))$ is large enough, then $Q_{\max}(\tau)$ is large enough, and T_2 and $\frac{M}{L_\alpha^\alpha(\mathbf{Q}(\tau))}$ can be made arbitrarily small. Thus, there exists $B > 0$ such that if $L_\alpha(\mathbf{Q}(\tau)) > B$, then

$$\mathbb{E} \left[L_\alpha(\mathbf{Q}(\tau+1)) - L_\alpha(\mathbf{Q}(\tau)) \mid \mathbf{Q}(\tau) \right] \leq -\frac{1}{2}(1-\rho)M^{-\frac{\alpha}{1+\alpha}}.$$

5. PROOF OF THEOREM 3.1

In this section, we fix some $\alpha \in (0, 1)$ and prove that the MW- α policy induces finite steady-state expected queue lengths. The key to our proof is the use of the Lyapunov function $\Phi(\mathbf{x}) = L_\alpha^2(\mathbf{x})$. This is to be contrasted with the use of the standard Lyapunov function, $\sum_i x_i^{1+\alpha}$, in the literature, or the ‘‘norm’’-Lyapunov function $L_\alpha(\mathbf{x})$ that we used in establishing the drift inequality of Theorem 4.4.

Throughout the proof, we drop the subscript α from L_α , F_α , and f_α , as they are clear from the context. We also use $\|\mathbf{x}\|$ to denote the $(\alpha+1)$ -norm of the vector \mathbf{x} , again dropping the subscript.

As usual, we consider the conditional expected drift at time τ ,

$$D(\mathbf{Q}(\tau)) = \mathbb{E} \left[\Phi(\mathbf{Q}(\tau+1)) - \Phi(\mathbf{Q}(\tau)) \mid \mathbf{Q}(\tau) \right].$$

Recall the notation $Q_{\max}(\tau) = \max\{Q_1(\tau), \dots, Q_M(\tau)\}$. Since for $Q_{\max} < 2$, $D(\mathbf{Q}(\tau))$ is bounded by a constant, we assume throughout the proof that $Q_{\max}(\tau) \geq 2$. As in the proof of Theorem 4.4 for the case $\alpha \in (0, 1)$, we shall use the second order mean value theorem to obtain a bound on $D(\mathbf{Q}(\tau))$. Using the definition $\Phi(\mathbf{x}) = L^2(\mathbf{x})$, we have

$$[\nabla\Phi(\mathbf{x})]_i = 2L(\mathbf{x})\frac{\partial L(\mathbf{x})}{\partial x_i} = 2f(x_i)L^{1-\alpha}(\mathbf{x}), \quad (17)$$

and

$$\begin{aligned} \frac{\partial^2\Phi}{\partial x_i\partial x_j}(\mathbf{x}) &= 2\frac{\partial L(\mathbf{x})}{\partial x_i} \cdot \frac{\partial L(\mathbf{x})}{\partial x_j} + 2L(\mathbf{x})\frac{\partial^2 L(\mathbf{x})}{\partial x_i\partial x_j} \\ &= 2\frac{f(x_i)f(x_j)}{L^{2\alpha}(\mathbf{x})} + 2L(\mathbf{x})\left(\delta_{ij}\frac{f'(x_i)}{L^\alpha(\mathbf{x})} - \frac{\alpha f(x_i)f(x_j)}{L^{2\alpha+1}(\mathbf{x})}\right) \\ &= 2(1-\alpha)\frac{f(x_i)f(x_j)}{L^{2\alpha}(\mathbf{x})} + 2\delta_{ij}f'(x_i)L^{1-\alpha}(\mathbf{x}). \end{aligned} \quad (18)$$

Using the second order mean value theorem and the notation $\mathbf{Q}(\tau+1) = \mathbf{Q}(\tau) + \boldsymbol{\delta}(\tau)$, we have, for some $\theta \in [0, 1]$,

$$\begin{aligned} & \Phi(\mathbf{Q}(\tau+1)) - \Phi(\mathbf{Q}(\tau)) \\ & \leq 2L^{1-\alpha}(\mathbf{Q}(\tau)) \left(\sum_i f(Q_i(\tau))\delta_i(\tau) \right) \\ & \quad + L^{1-\alpha}(\mathbf{Q}(\tau) + \theta\boldsymbol{\delta}(\tau)) \left(\sum_i f'(Q_i(\tau) + \theta\delta_i(\tau)) \right) \\ & \quad + (1-\alpha) \frac{(\sum_i f(Q_i(\tau) + \theta\delta_i(\tau))\delta_i(\tau))^2}{L^{2\alpha}(\mathbf{Q}(\tau) + \theta\boldsymbol{\delta}(\tau))}. \end{aligned} \quad (19)$$

Let us denote the three terms on the RHS of (19) as \bar{T}_1 , \bar{T}_2 and \bar{T}_3 respectively, so that

$$\begin{aligned} \bar{T}_1 &= 2L^{1-\alpha}(\mathbf{Q}(\tau)) \left(\sum_i f(Q_i(\tau))\delta_i(\tau) \right), \\ \bar{T}_2 &= L^{1-\alpha}(\mathbf{Q}(\tau) + \theta\boldsymbol{\delta}(\tau)) \left(\sum_i f'(Q_i(\tau) + \theta\delta_i(\tau)) \right), \\ \text{and } \bar{T}_3 &= (1-\alpha) \frac{(\sum_i f(Q_i(\tau) + \theta\delta_i(\tau))\delta_i(\tau))^2}{L^{2\alpha}(\mathbf{Q}(\tau) + \theta\boldsymbol{\delta}(\tau))}. \end{aligned}$$

We consider these terms one at a time.

a) By Lemma 4.3 (iii), $f(r) \leq r^\alpha + 1$. Using the fact that $\delta_i(\tau) \in \{-1, 0, 1\}$, we obtain

$$\bar{T}_1 \leq 2L^{1-\alpha}(\mathbf{Q}(\tau)) \left(M + \sum_i Q_i^\alpha(\tau)\delta_i(\tau) \right).$$

When we take a conditional expectation, an argument similar to the one in earlier sections yields

$$\mathbb{E} \left[\sum_i Q_i^\alpha(\tau)\delta_i(\tau) \mid \mathbf{Q}(\tau) \right] \leq -(1-\rho)w_\alpha(\mathbf{Q}(\tau)).$$

Thus,

$$\begin{aligned} \mathbb{E} \left[\bar{T}_1 \mid \mathbf{Q}(\tau) \right] &\leq -2(1-\rho)w_\alpha(\mathbf{Q}(\tau))L^{1-\alpha}(\mathbf{Q}(\tau)) \\ &\quad + 2ML^{1-\alpha}(\mathbf{Q}(\tau)). \end{aligned}$$

In general, for $r, s \geq 0$ and $\beta \in [0, 1]$,

$$(r+s)^\beta \leq r^\beta + s^\beta. \quad (20)$$

Now, by Lemma 4.3 (iii'), $r^{\alpha+1} - 2 \leq (\alpha+1)F(r) \leq r^{\alpha+1} + 2$, so

$$\sum_i x_i^{\alpha+1} - 2M \leq (\alpha+1) \sum_i F(x_i) \leq \sum_i x_i^{\alpha+1} + 2M.$$

We use inequality (20), with $r = x_i^{\alpha+1}$, $s = 2M$, and $\beta = (1-\alpha)/(1+\alpha) \in (0, 1)$, to obtain

$$\begin{aligned} L^{1-\alpha}(\mathbf{x}) &= \left((\alpha+1) \sum_i F(x_i) \right)^{\frac{1-\alpha}{1+\alpha}} \\ &\leq \left(2M + \sum_i x_i^{\alpha+1} \right)^{\frac{1-\alpha}{1+\alpha}} \\ &\leq (2M)^{\frac{1-\alpha}{1+\alpha}} + \left(\sum_i x_i^{\alpha+1} \right)^{\frac{1-\alpha}{1+\alpha}} \\ &= (2M)^{\frac{1-\alpha}{1+\alpha}} + \|\mathbf{x}\|^{1-\alpha}. \end{aligned}$$

A similar argument, based on inequality (20), with $r = (\alpha+1)F(x_i)$ and $s = 2M$, yields

$$\|\mathbf{x}\|^{1-\alpha} - (2M)^{\frac{1-\alpha}{1+\alpha}} \leq L^{1-\alpha}(\mathbf{x}).$$

We also know that

$$w_\alpha(\mathbf{Q}(\tau)) \geq Q_{\max}^\alpha(\tau) \geq M^{-\frac{\alpha}{\alpha+1}} \|\mathbf{Q}(\tau)\|^\alpha.$$

Putting all these facts together, we obtain

$$\begin{aligned}
& \mathbb{E} \left[\bar{T}_1 \mid \mathbf{Q}(\tau) \right] \\
& \leq -2(1-\rho)w_\alpha(\mathbf{Q}(\tau))L^{1-\alpha}(\mathbf{Q}(\tau)) + 2ML^{1-\alpha}(\mathbf{Q}(\tau)) \\
& \leq -2(1-\rho)M^{-\frac{\alpha}{\alpha+1}} \|\mathbf{Q}(\tau)\|^\alpha \left(\|\mathbf{Q}(\tau)\|^{1-\alpha} - (2M)^{\frac{1-\alpha}{1+\alpha}} \right) \\
& \quad + 2M \left((2M)^{\frac{1-\alpha}{1+\alpha}} + \|\mathbf{Q}(\tau)\|^{1-\alpha} \right) \\
& = -2(1-\rho)M^{-\frac{\alpha}{\alpha+1}} \|\mathbf{Q}(\tau)\| + 2M\|\mathbf{Q}(\tau)\|^{1-\alpha} \\
& \quad + 2^{\frac{2}{1+\alpha}}(1-\rho)M^{\frac{1-2\alpha}{1+\alpha}} \|\mathbf{Q}(\tau)\|^\alpha + (2M)^{\frac{2}{1+\alpha}}. \tag{21}
\end{aligned}$$

b) We now consider the term \bar{T}_2 . Since $\alpha \in (0, 1)$, we have $f'(r) \leq 2$ for all $r \geq 0$. Since we also have $\theta \in [0, 1]$ and $\delta_i(\tau) \in \{-1, 0, 1\}$, and using the fact that $L^{1-\alpha}(\mathbf{x}) \leq (2M)^{\frac{1-\alpha}{1+\alpha}} + \|\mathbf{x}\|^{1-\alpha}$, we have

$$\begin{aligned}
\bar{T}_2 & \leq 2ML(\mathbf{Q}(\tau) + \theta\delta(\tau))^{1-\alpha} \\
& \leq 2M \left((2M)^{\frac{1-\alpha}{1+\alpha}} + \|\mathbf{Q}(\tau) + \theta\delta(\tau)\|^{1-\alpha} \right) \\
& = (2M)^{\frac{2}{1+\alpha}} + 2M\|\mathbf{Q}(\tau) + \theta\delta(\tau)\|^{1-\alpha}.
\end{aligned}$$

Now $\|\mathbf{Q}(\tau) + \theta\delta(\tau)\| \leq \|\mathbf{Q}(\tau) + \mathbf{1}\| \leq \|\mathbf{Q}(\tau)\| + \|\mathbf{1}\| = \|\mathbf{Q}(\tau)\| + M^{\frac{1}{\alpha+1}}$. Since $\alpha \in (0, 1)$, we have $0 < 1 - \alpha < 1$, and so

$$\begin{aligned}
\|\mathbf{Q}(\tau) + \theta\delta(\tau)\|^{1-\alpha} & \leq \left(\|\mathbf{Q}(\tau)\| + M^{\frac{1}{\alpha+1}} \right)^{1-\alpha} \\
& \leq \|\mathbf{Q}(\tau)\|^{1-\alpha} + M^{\frac{1-\alpha}{\alpha+1}}.
\end{aligned}$$

Putting everything together, we have

$$\begin{aligned}
\bar{T}_2 & \leq (2M)^{\frac{2}{1+\alpha}} + 2M \left(\|\mathbf{Q}(\tau)\|^{1-\alpha} + M^{\frac{1-\alpha}{\alpha+1}} \right) \\
& = (2 + 2^{\frac{2}{1+\alpha}})\alpha M^{\frac{2}{1+\alpha}} + 2M\|\mathbf{Q}(\tau)\|^{1-\alpha}. \tag{22}
\end{aligned}$$

c) We finally consider \bar{T}_3 . For notational convenience, we write $\mathbf{x} = \mathbf{Q}(\tau) + \theta\delta(\tau)$, and let $x_{\max} = \max\{x_1, \dots, x_M\}$. Note that since $\delta_i(\tau) \in \{-1, 0, 1\}$, $\theta \in [0, 1]$, and we assumed that $Q_{\max} \geq 2$, we always have $x_{\max} \geq 1$. We consider the numerator and the denominator separately. First use the facts that $f(r) \geq 0$ for all $r \geq 0$ (cf. Lemma 4.3 (ii)), and $\delta_i(\tau) \in \{-1, 0, 1\}$, to obtain

$$\left(\sum_i f(x_i) \delta_i(\tau) \right)^2 \leq \left(\sum_i f(x_i) \right)^2.$$

Since f is increasing in r (cf. Lemma 4.3 (ii)),

$$\left(\sum_i f(x_i) \right)^2 \leq (Mf(x_{\max}))^2 = M^2 f^2(x_{\max}).$$

Thus,

$$\left(\sum_i f(x_i) \delta_i(\tau) \right)^2 \leq M^2 f^2(x_{\max}).$$

Next, since $F(r) = \int_0^r f(s) ds$ and $f \geq 0$, we have $F \geq 0$ as well. Thus,

$$L^{2\alpha}(\mathbf{x}) = \left((\alpha + 1) \sum_i F(x_i) \right)^{\frac{2\alpha}{\alpha+1}} \geq ((\alpha + 1)F(x_{\max}))^{\frac{2\alpha}{\alpha+1}},$$

and so

$$\bar{T}_3 \leq (1 - \alpha) \frac{M^2 f^2(x_{\max})}{((\alpha + 1)F(x_{\max}))^{\frac{2\alpha}{\alpha+1}}}.$$

We will show that \bar{T}_3 is bounded above by a positive constant, whenever $x_{\max} \geq 1$. Indeed, by Lemma 4.3 (iii) and (iii'), as $x_{\max} \rightarrow \infty$,

$$\frac{f^2(x_{\max})}{x_{\max}^{2\alpha}} \rightarrow 1 \quad \text{and} \quad \frac{((\alpha + 1)F(x_{\max}))^{\frac{2\alpha}{\alpha+1}}}{x_{\max}^{2\alpha}} \rightarrow 1,$$

so

$$(1 - \alpha) \frac{M^2 f^2(x_{\max})}{((\alpha + 1)F(x_{\max}))^{\frac{2\alpha}{\alpha+1}}} \rightarrow (1 - \alpha)M^2$$

as $x_{\max} \rightarrow \infty$. Using the continuity of f and F for $x_{\max} \geq 1$, it follows that there exists a constant $\tilde{K} > 0$ such that

$$\bar{T}_3 \leq (1 - \alpha) \frac{M^2 f^2(x_{\max})}{((\alpha + 1)F(x_{\max}))^{\frac{2\alpha}{\alpha+1}}} \leq \tilde{K}, \tag{23}$$

whenever $x_{\max} \geq 1$.

Putting together the bounds (21), (22), and (23) for \bar{T}_1 , \bar{T}_2 , and \bar{T}_3 , respectively, we conclude that, for $x_{\max} \geq 1$,

$$\begin{aligned}
D(\mathbf{Q}(\tau)) & \leq -2(1-\rho)M^{-\frac{\alpha}{\alpha+1}} \|\mathbf{Q}(\tau)\| + 2M\|\mathbf{Q}(\tau)\|^{1-\alpha} \\
& \quad + 4(1-\rho)M^{\frac{1-2\alpha}{1+\alpha}} \|\mathbf{Q}(\tau)\|^\alpha + (2M)^{\frac{2}{1+\alpha}} \\
& \quad + (2 + 2^{\frac{2}{1+\alpha}})\alpha M^{\frac{2}{1+\alpha}} + 2M\|\mathbf{Q}(\tau)\|^{1-\alpha} + \tilde{K} \\
& = -\bar{A}\|\mathbf{Q}(\tau)\| + C_1\|\mathbf{Q}(\tau)\|^{1-\alpha} + C_2\|\mathbf{Q}(\tau)\|^\alpha + K, \tag{24}
\end{aligned}$$

for some positive constants \bar{A} , C_1 , C_2 and K . Since $\alpha \in (0, 1)$, the $\|\mathbf{Q}(\tau)\|$ term dominates. In particular, there exist positive constants A and D such that as long as $\max_i Q_i(\tau) \geq D$, we have

$$D(\mathbf{Q}(\tau)) \leq -A\|\mathbf{Q}(\tau)\| + K. \tag{25}$$

On the other hand, on the bounded set where $\max_i Q_i(\tau) \leq D$, the drift $D(\mathbf{Q}(\tau))$ is also bounded by a constant. By suitably redefining the constant K , we conclude that Eq. (25) holds for all possible values of $\mathbf{Q}(\tau)$.

The drift condition (25) is the standard Foster-Lyapunov criterion for the Lyapunov function Φ and implies the positive recurrence of the Markov chain $\mathbf{Q}(\cdot)$ under the MW- α policy, for $\alpha \in (0, 1)$. The irreducibility and aperiodicity of the underlying Markov chain implies the existence of a unique stationary distribution π as well as ergodicity. Let \mathbf{Q}_∞ be a random variable distributed according to π . Then $\mathbf{Q}(\tau)$ converges to \mathbf{Q}_∞ in distribution. Using Skorohod's representation theorem, we can embed the random vectors $\mathbf{Q}(\tau)$ in a suitable probability space so that they converge to \mathbf{Q}_∞ almost surely. With this embedding, $\|\mathbf{Q}(\tau)\| \rightarrow \|\mathbf{Q}_\infty\|$, and $(\sum_{\tau=0}^{T-1} \|\mathbf{Q}(\tau)\|)/T \rightarrow \|\mathbf{Q}_\infty\|$, almost surely. Using Fatou's Lemma, we have

$$\begin{aligned}
\mathbb{E}[\|\mathbf{Q}_\infty\|] & = \mathbb{E} \left[\liminf_{T \rightarrow \infty} \frac{1}{T} \sum_{\tau=0}^{T-1} \|\mathbf{Q}(\tau)\| \right] \\
& \leq \liminf_T \mathbb{E} \left[\frac{1}{T} \sum_{\tau=0}^{T-1} \|\mathbf{Q}(\tau)\| \right].
\end{aligned}$$

On the other hand, the drift inequality (25) is well known to imply that the RHS above is finite; see, e.g., Lemma 4.1 of [7]. This proves that $\mathbb{E}[\|\mathbf{Q}_\infty\|] < \infty$. By the equivalence of norms, the result for $\|\mathbf{Q}\|_1$ follows as well.

6. EXPONENTIAL BOUND UNDER MW- α

In this section we derive an exponential upper bound on the tail probability of the stationary queue-size distribution, under the MW- α policy.

6.1 Proof of Theorem 3.2: $\alpha \geq 1$

The proof of Theorem 3.2 relies on the following proposition, and the drift inequality established in Theorem 4.4.

PROPOSITION 6.1. *Consider a switched network operating under the MW- α policy with $\alpha \geq 1$, and arrival rate vector λ with $\rho = \rho(\lambda) < 1$. Let π be the unique stationary distribution of the Markov chain $\mathbf{Q}(\cdot)$. Suppose that for all τ ,*

$$\left| \|\mathbf{Q}(\tau + 1)\|_{\alpha+1} - \|\mathbf{Q}(\tau)\|_{\alpha+1} \right| \leq \nu_{\max}.$$

Furthermore, suppose that for some constants $B > 0$ and $\gamma > 0$, and whenever $\|\mathbf{Q}(\tau)\|_{1+\alpha} > B$, we have

$$\mathbb{E}[\|\mathbf{Q}(\tau + 1)\|_{\alpha+1} - \|\mathbf{Q}(\tau)\|_{\alpha+1} \mid \mathbf{Q}(\tau)] \leq -\gamma.$$

Then for any $\ell \in \mathbb{Z}_+$,

$$\mathbb{P}_{\pi}(\|\mathbf{Q}(\tau)\|_{\alpha+1} > B + 2\nu_{\max}\ell) \leq \left(\frac{\bar{\nu}}{\bar{\nu} + \gamma}\right)^{\ell+1}.$$

Proposition 6.1 follows immediately from the following Lemma, which is a minor adaptation of Lemma 1 of [3]. An interested reader may refer to the proof of Theorem 1(a) in [3] to see how Lemma 6.2 leads to the bound claimed in Proposition 6.1.

LEMMA 6.2. *Under the same assumptions in Proposition 6.1, and for any $c > B - \nu_{\max}$,*

$$\begin{aligned} & \mathbb{P}_{\pi}(\|\mathbf{Q}(\tau)\|_{\alpha+1} > c + \nu_{\max}) \\ & \leq \left(\frac{\bar{\nu}}{\bar{\nu} + \gamma}\right) \mathbb{P}_{\pi}(\|\mathbf{Q}(\tau)\|_{\alpha+1} > c - \nu_{\max}). \end{aligned} \quad (26)$$

PROOF. Since this Lemma is a minor adaptation of Lemma 1 in [3], we only indicate the changes to the proof of Lemma 1 in [3] that lead to our claimed result. First let us point out that the proof in [3] makes use of the finiteness of the expected value of the Lyapunov function under the stationary distribution π . In our case, the Lyapunov function in question is $\|\cdot\|_{\alpha+1}$, and the finiteness follows from Theorem 3.1 by noticing that all norms are equivalent.

As in the proof of Lemma 1 in [3], define $\hat{\Phi}(\mathbf{x}) = \max\{c, \|\mathbf{x}\|_{\alpha+1}\}$. Note that the maximal change in $\hat{\Phi}(\mathbf{x})$ in one time step is at most ν_{\max} . As in [3], we consider all \mathbf{x} satisfying $c - \nu_{\max} < \|\mathbf{x}\|_{\alpha+1} \leq c + \nu_{\max}$. Then,

$$\begin{aligned} & \mathbb{E}[\hat{\Phi}(\mathbf{Q}(\tau + 1)) \mid \mathbf{Q}(\tau) = \mathbf{x}] - \hat{\Phi}(\mathbf{x}) \\ & \leq \sum_{\mathbf{x}': \|\mathbf{x}'\| > \|\mathbf{x}\|} p(\mathbf{x}, \mathbf{x}') (\|\mathbf{x}'\| - \|\mathbf{x}\|) \\ & \leq \mathbb{E}[\|\mathbf{a}(\tau)\|] = \bar{\nu}. \end{aligned}$$

The proof of Lemma 1 in [3] essentially used ν_{\max} as an upper bound on $\bar{\nu}$. For our result, we keep $\bar{\nu}$ and then proceed as in the proof in [3]. \square

Completing the Proof of Theorem 3.2 ($\alpha \geq 1$).

Now the proof of Theorem 3.2 follows immediately from Proposition 6.1 by noticing that Theorem 4.4 provides the desired drift inequality, and the maximal change in $\|\mathbf{Q}(\tau)\|_{1+\alpha}$ in one time step is at most $\nu_{\max} = M^{\frac{1}{1+\alpha}}$, because each queue can receive at most one arrival and have at most one departure per time step.

6.2 Proof of Theorem 3.2: $\alpha \in (0, 1)$

The proof for the case $\alpha \in (0, 1)$ is entirely parallel to that in the previous section and we do not reproduce it here.

7. TRANSIENT ANALYSIS

In this section, we present a transient analysis of the MW- α policy with $\alpha \geq 1$. First we present a general maximal lemma, which is then specialized to the switched network. In particular, we prove a drift inequality for the Lyapunov function $\tilde{L}(\mathbf{x}) = \frac{1}{\alpha+1} \sum_i x_i^{\alpha+1}$. We combine the drift inequality with the maximal lemma to obtain a maximal inequality for the switched network. We then apply the maximal inequality to prove full state space collapse for $\alpha \geq 1$.

7.1 The Key Lemma

Our analysis relies on the following lemma:

LEMMA 7.1. *Let $(\mathcal{F}_n)_{n \in \mathbb{Z}_+}$ be a filtration on a probability space. Let $(X_n)_{n \in \mathbb{Z}_+}$ be a nonnegative \mathcal{F}_n -adapted stochastic process that satisfies*

$$\mathbb{E}[X_{n+1} \mid \mathcal{F}_n] \leq X_n + B_n \quad (27)$$

where B_n 's are nonnegative random variables (not necessarily \mathcal{F}_n -adapted) with finite means. Let $X_n^* = \max\{X_0, \dots, X_n\}$ and suppose that $X_0 = 0$. Then, for any $a > 0$ and any $T \in \mathbb{Z}_+$,

$$\mathbb{P}(X_T^* \geq a) \leq \frac{\sum_{n=0}^{T-1} \mathbb{E}[B_n]}{a}.$$

This lemma is a simple consequence of the following standard maximal inequality for nonnegative supermartingales (see for example, Exercise 4, Section 12.4, of [9]):

THEOREM 7.2. *Let $(\mathcal{F}_n)_{n \in \mathbb{Z}_+}$ be a filtration on a probability space. Let $(Y_n)_{n \in \mathbb{Z}_+}$ be a nonnegative \mathcal{F}_n -adapted supermartingale, i.e., for all n ,*

$$\mathbb{E}[Y_{n+1} \mid \mathcal{F}_n] \leq Y_n.$$

Let $Y_T^* = \max\{Y_0, \dots, Y_T\}$. Then,

$$\mathbb{P}(Y_T^* \geq a) \leq \frac{\mathbb{E}[Y_0]}{a}.$$

PROOF OF LEMMA 7.1. First note that if we take the conditional expectation on both sides of (27), given \mathcal{F}_n , we have

$$\begin{aligned} \mathbb{E}[X_{n+1} \mid \mathcal{F}_n] & \leq \mathbb{E}[X_n \mid \mathcal{F}_n] + \mathbb{E}[B_n \mid \mathcal{F}_n] \\ & = X_n + \mathbb{E}[B_n \mid \mathcal{F}_n]. \end{aligned}$$

Fix $T \in \mathbb{Z}_+$. For any $n \leq T$, define

$$Y_n = X_n + \mathbb{E}\left[\sum_{k=n}^{T-1} B_k \mid \mathcal{F}_n\right].$$

Then

$$\begin{aligned}\mathbb{E}[Y_{n+1} \mid \mathcal{F}_n] &= \mathbb{E}[X_{n+1} \mid \mathcal{F}_n] \\ &\quad + \mathbb{E} \left[\mathbb{E} \left[\sum_{k=n+1}^{T-1} B_k \mid \mathcal{F}_{n+1} \right] \mid \mathcal{F}_n \right] \\ &\leq X_n + \mathbb{E}[B_n \mid \mathcal{F}_n] + \mathbb{E} \left[\sum_{k=n+1}^{T-1} B_k \mid \mathcal{F}_n \right] \\ &= Y_n.\end{aligned}$$

Thus, Y_n is an \mathcal{F}_n -adapted supermartingale; furthermore, by definition, Y_n is non-negative for all n . Therefore, by Theorem 7.2,

$$\mathbb{P}(Y_T^* \geq a) \leq \frac{\mathbb{E}[Y_0]}{a} = \frac{\mathbb{E} \left[\sum_{k=0}^{T-1} B_k \right]}{a}.$$

But $Y_n \geq X_n$ for all n , since the B_k are nonnegative. Thus

$$\mathbb{P}(X_T^* \geq a) \leq \mathbb{P}(Y_T^* \geq a) \leq \frac{\mathbb{E} \left[\sum_{k=0}^{T-1} B_k \right]}{a}.$$

□

We have the following corollary of Lemma 7.1 in which we take all the B_n equal to the same constant:

COROLLARY 7.3. *Let \mathcal{F}_n , X_n and X_n^* be as in Lemma 7.1. Suppose that*

$$\mathbb{E}[X_{n+1} \mid \mathcal{F}_n] \leq X_n + B,$$

for all $n \geq 0$, where B is a nonnegative constant. Then, for any $a > 0$ and any $T \in \mathbb{Z}_+$,

$$\mathbb{P}(X_T^* \geq a) \leq \frac{BT}{a}.$$

7.2 The Maximal Inequality for Switched Networks

We employ the Lyapunov function

$$\tilde{L}(\mathbf{x}) = \frac{1}{\alpha + 1} \sum_{i=1}^M x_i^{\alpha+1}, \quad (28)$$

to study the MW- α policy. This is the Lyapunov function that was used in [15] to establish positive recurrence of the chain $\mathbf{Q}(\cdot)$ under the MW- α policy. Below we fine-tune the proof in [15] to obtain a more precise bound.

LEMMA 7.4. *Let $\alpha \geq 1$. For a switched network model operating under the MW- α policy with $\rho = \rho(\boldsymbol{\lambda}) < 1$, we have:*

$$\mathbb{E}[\tilde{L}(\mathbf{Q}(\tau+1)) - \tilde{L}(\mathbf{Q}(\tau)) \mid \mathbf{Q}(\tau)] \leq \frac{\bar{K}(\alpha, M)}{(1-\rho)^{\alpha-1}}, \quad (29)$$

where $\bar{K}(\alpha, M)$ is a constant depending only on α and M .

PROOF. We employ the same strategy as in previous sections. By the second-order mean value theorem, there exists $\theta \in [0, 1]$ such that

$$\begin{aligned}\tilde{L}(\mathbf{Q}(\tau+1)) - \tilde{L}(\mathbf{Q}(\tau)) &= \frac{1}{\alpha+1} \sum_{i=1}^M ((Q_i(\tau) + \delta_i(\tau))^{\alpha+1} - Q_i^{\alpha+1}(\tau)) \\ &= \sum_{i=1}^M Q_i^\alpha(\tau) \delta_i(\tau) + \sum_{i=1}^M \alpha(Q_i(\tau) + \theta \delta_i(\tau))^{\alpha-1} \delta_i^2(\tau).\end{aligned}$$

Let us bound the second term on the RHS. We have

$$\begin{aligned}&\sum_{i=1}^M \alpha(Q_i(\tau) + \theta \delta_i(\tau))^{\alpha-1} \delta_i^2(\tau) \\ &\leq \sum_{i=1}^M \alpha(Q_i(\tau) + \theta)^{\alpha-1} \leq \sum_{i=1}^M \alpha(Q_i(\tau) + 1)^{\alpha-1} \\ &\leq \alpha \sum_{i=1}^M (2^{\alpha-1} Q_i^{\alpha-1}(\tau) + 1) = \alpha 2^{\alpha-1} \sum_{i=1}^M Q_i^{\alpha-1}(\tau) + \alpha M \\ &\leq \alpha 2^{\alpha-1} M Q_{\max}^{\alpha-1}(\tau) + \alpha M.\end{aligned}$$

The third inequality follows because when $Q_i(\tau) \geq 1$, $(Q_i(\tau) + 1)^{\alpha-1} \leq (2Q_i(\tau))^{\alpha-1} = 2^{\alpha-1} Q_i^{\alpha-1}(\tau)$, and when $Q_i(\tau) = 0$, $(Q_i(\tau) + 1)^{\alpha-1} = 1$.

Let us now take conditional expectations. From Section 4, we know that

$$\begin{aligned}\mathbb{E} \left[\sum_{i=1}^M Q_i^\alpha(\tau) \delta_i(\tau) \mid \mathbf{Q}(\tau) \right] &\leq -(1-\rho) w_\alpha(\mathbf{Q}(\tau)) \\ &\leq -(1-\rho) Q_{\max}^\alpha(\tau).\end{aligned}$$

Thus, if we combine the inequalities above, we have

$$\begin{aligned}\mathbb{E}[\tilde{L}(\mathbf{Q}(\tau+1)) - \tilde{L}(\mathbf{Q}(\tau)) \mid \mathbf{Q}(\tau)] &\leq -(1-\rho) Q_{\max}^\alpha(\tau) + \alpha 2^{\alpha-1} M Q_{\max}^{\alpha-1}(\tau) + \alpha M.\end{aligned} \quad (30)$$

It is a simple exercise in calculus to see that the RHS of (30) is maximized at $Q_{\max}(\tau) = (\alpha-1)2^{\alpha-1}M/(1-\rho)$, giving the maximum value

$$\frac{(\alpha-1)^{\alpha-1} 2^{\alpha(\alpha-1)} M^\alpha}{(1-\rho)^{\alpha-1}} + \alpha M = O((1-\rho)^{1-\alpha}).$$

□

Proof of Theorem 3.3.

Let $b > 0$. Then

$$\begin{aligned}\mathbb{P}(Q_{\max}^*(T) \geq b) &= \mathbb{P} \left(\frac{1}{\alpha+1} (Q_{\max}^*(T))^{\alpha+1} \geq \frac{1}{\alpha+1} b^{\alpha+1} \right) \\ &\leq \mathbb{P} \left(\max_{\tau \in \{0, \dots, T\}} \tilde{L}(\mathbf{Q}(\tau)) \geq \frac{1}{\alpha+1} b^{\alpha+1} \right).\end{aligned}$$

Now, by Lemma 7.4 and Corollary 7.3,

$$\begin{aligned}\mathbb{P} \left(\max_{\tau \in \{0, \dots, T\}} \tilde{L}(\mathbf{Q}(\tau)) \geq \frac{1}{\alpha+1} b^{\alpha+1} \right) &\leq \frac{(\alpha+1)\bar{K}(\alpha, M)T}{(1-\rho)^{\alpha-1} b^{\alpha+1}} \\ &= \frac{K(\alpha, M)T}{(1-\rho)^{\alpha-1} b^{\alpha+1}},\end{aligned}$$

where $K(\alpha, M) = (\alpha+1)\bar{K}(\alpha, M)$.

7.3 Full State Space Collapse for $\alpha \geq 1$

Throughout this section, we assume that we are given $\alpha \geq 1$, and correspondingly, the Lyapunov function $\tilde{L}(\mathbf{x}) = \frac{1}{\alpha+1} \sum_{i=1}^M x_i^{\alpha+1}$. To state the full state space collapse result for $\alpha \geq 1$, we need some preliminary definitions and the statement of the multiplicative state space collapse result.

Let Σ be the convex hull of \mathcal{S} (the set of feasible schedules), and let $\bar{\boldsymbol{\Lambda}}$ be defined by

$$\bar{\boldsymbol{\Lambda}} = \left\{ \boldsymbol{\lambda} \in \mathbb{R}_+^M : \boldsymbol{\lambda} \leq \boldsymbol{\sigma}' \text{ componentwise, for some } \boldsymbol{\sigma}' \in \Sigma \right\}.$$

Note that this is the closure of the capacity region Λ defined earlier. Recall the definition of the load $\rho(\lambda)$ of an arrival rate vector λ . It is clear that $\lambda \in \bar{\Lambda}$ iff $\rho(\lambda) \leq 1$. Define $\partial\Lambda$ the set of *critical* arrival rate vectors:

$$\partial\Lambda = \bar{\Lambda} - \Lambda = \left\{ \lambda \in \bar{\Lambda} : \rho(\lambda) = 1 \right\}.$$

Now consider the linear optimization problem, named $\text{DUAL}(\lambda)$ in [18]:

$$\begin{aligned} & \text{maximize} && \xi \cdot \lambda \\ & \text{subject to} && \max_{\sigma \in \mathcal{S}} \xi \cdot \sigma \leq 1, \\ & && \xi \in \mathbb{R}_+^M. \end{aligned}$$

For $\lambda \in \partial\Lambda$, the optimal value of the objective in $\text{DUAL}(\lambda)$ is 1 (cf.[18]). The set of optimal solutions to $\text{DUAL}(\lambda)$ is a bounded polyhedron, and we let $\mathcal{S}^* = \mathcal{S}^*(\lambda)$ be the set of its extreme points.

Fix $\lambda \in \partial\Lambda$. We then consider the optimization problem $\text{ALGD}(w)$:

$$\begin{aligned} & \text{minimize} && \tilde{L}(x) \\ & \text{subject to} && \xi \cdot x \geq w_\xi \text{ for all } \xi \in \mathcal{S}^*(\lambda), \\ & && x \in \mathbb{R}_+^M. \end{aligned}$$

We know from [18] that $\text{ALGD}(w)$ has a unique solution. We now define the *lifting map*:

DEFINITION 7.5. Fix some $\lambda \in \partial\Lambda$. The lifting map $\Delta^\lambda : \mathbb{R}_+^{|\mathcal{S}^*(\lambda)|} \rightarrow \mathbb{R}_+^M$ maps w to the unique solution to $\text{ALGD}(w)$. We also define the workload map $W^\lambda : \mathbb{R}_+^M \rightarrow \mathbb{R}_+^{|\mathcal{S}^*(\lambda)|}$ by $W^\lambda(\mathbf{q}) = (\xi \cdot \mathbf{q})_{\xi \in \mathcal{S}^*(\lambda)}$.

Fix $\lambda \in \partial\Lambda$. Consider a sequence of switched networks indexed by $r \in \mathbb{N}$, operating under the MW- α policy (recall that $\alpha \geq 1$ here), all with the same number M of queues and feasible schedules. Suppose that $\lambda^r \in \Lambda$ for all r , and that $\lambda^r = \lambda - \Gamma/r$, for some $\Gamma \in \mathbb{R}_+^M$. For simplicity, suppose that all networks start with empty queues. Consider the following central limit scaling,

$$\hat{\mathbf{q}}^r(t) = \mathbf{Q}^r(r^2t)/r, \quad (31)$$

where $\mathbf{Q}^r(\tau)$ is the queue size vector of the r th network at time τ , and where we extend the domain of $\mathbf{Q}^r(\cdot)$ to \mathbb{R}_+ by linear interpolation in each interval $(\tau - 1, \tau)$.

We are finally ready to state the multiplicative state space collapse result (Theorem 8.2 in [18]):

THEOREM 7.6. Fix $T > 0$, and let

$$\|\mathbf{x}(\cdot)\| = \sup_{i \in \{1, \dots, M\}, 0 \leq t \leq T} |x_i(t)|.$$

Under the above assumptions, for any $\varepsilon > 0$,

$$\lim_{r \rightarrow \infty} \mathbb{P} \left(\frac{\|\hat{\mathbf{q}}^r(\cdot) - \Delta^\lambda(W^\lambda(\hat{\mathbf{q}}^r(\cdot)))\|}{\|\hat{\mathbf{q}}^r(\cdot)\| \vee 1} < \varepsilon \right) = 1.$$

We now state and prove the full state space collapse result.

THEOREM 7.7. Under the same assumptions in Theorem 7.6, and for any $\varepsilon > 0$,

$$\lim_{r \rightarrow \infty} \mathbb{P} \left(\|\hat{\mathbf{q}}^r(\cdot) - \Delta^\lambda(W^\lambda(\hat{\mathbf{q}}^r(\cdot)))\| < \varepsilon \right) = 1.$$

PROOF. First note that since $\lambda^r = \lambda - \Gamma/r$, the corresponding loads satisfy $\rho_r \leq 1 - C/r$, for some positive constant $C > 0$. By Theorem 3.3, for any $b > 0$,

$$\begin{aligned} \mathbb{P} \left(\max_{\tau \in \{0, 1, \dots, r^2T\}} Q_{\max}^r(\tau) \geq b \right) & \leq \frac{K(\alpha, M)r^2T}{(1 - \rho)^{\alpha-1}b^{\alpha+1}} \\ & \leq \frac{K(\alpha, M)r^{1+\alpha}T}{C^{\alpha-1}b^{\alpha+1}}. \end{aligned}$$

Then with $a = b/r$ and under the scaling in (31),

$$\mathbb{P}(\|\hat{\mathbf{q}}^r(\cdot)\| \geq a) \leq \frac{K(\alpha, M)}{C^{\alpha-1}} \frac{T}{a^{\alpha+1}},$$

for any $a > 0$.

For notational convenience, we write

$$D(r) = \|\hat{\mathbf{q}}^r(\cdot) - \Delta^\lambda(W^\lambda(\hat{\mathbf{q}}^r(\cdot)))\|.$$

Then, for any $a > 1$,

$$\begin{aligned} \mathbb{P}(D(r) \geq \varepsilon) & \leq \mathbb{P} \left(\frac{D(r)}{\|\hat{\mathbf{q}}^r(\cdot)\|} > \frac{\varepsilon}{a} \text{ or } \|\hat{\mathbf{q}}^r(\cdot)\| \geq a \right) \\ & \leq \mathbb{P} \left(\frac{D(r)}{\|\hat{\mathbf{q}}^r(\cdot)\|} > \frac{\varepsilon}{a} \right) + \mathbb{P}(\|\hat{\mathbf{q}}^r(\cdot)\| \geq a). \end{aligned}$$

Note that by Theorem 7.6, the first term on the RHS goes to 0 as $r \rightarrow \infty$, for any $a > 0$. The second term on the RHS can be made arbitrarily small by taking a sufficiently large. Thus, $\mathbb{P}(D(r) \geq \varepsilon) \rightarrow 0$ as $r \rightarrow \infty$. This concludes the proof. \square

8. DISCUSSION

The results in this paper can be viewed from two different perspectives. On the one hand, they provide much new information on the qualitative behavior (e.g., finiteness of expected backlog, bounds on steady-state tail probabilities and finite-horizon maximum excursion probabilities, etc.) of MW- α policies for switched network models. On the other hand, at a technical level, our results highlight the importance of choosing a suitable Lyapunov function: even if a network is shown to be stable by using a particular Lyapunov function, different choices may lead to more powerful bounds.

The methods and results in this paper extend in two directions. First, all of the results, suitably restated, remain valid for multihop networks under backpressure- α policies. Second, the same is true for flow-level models of the type considered in [11]. These extensions will be reported elsewhere.

Acknowledgements

This work was supported by NSF project CCF 0728554.

9. REFERENCES

- [1] M. Andrews, K. Kumaran, K. Ramanan, A. Stolyar, R. Vijayakumar, and P. Whiting. Scheduling in a queuing system with asynchronously varying service rates. *Probability in the Engineering and Informational Sciences*, 18(02):191–217, 2004.
- [2] D. Bertsekas. *Nonlinear Programming*. Athena Scientific, 1999.

- [3] D. Bertsimas, D. Gamarnik, and J. N. Tsitsiklis. Performance of multiclass Markovian queueing networks via piecewise linear Lyapunov functions. *The Annals of Applied Probability*, 11(4):1384–1428, 2001.
- [4] J. Dai and W. Lin. Asymptotic optimality of maximum pressure policies in stochastic processing networks. *Annals of Applied Probability*, 18(6):2239–2299, 2008.
- [5] J. Dai and S. Meyn. Stability and convergence of moments for multiclass queueing networks via fluid limit models. *IEEE Transactions on Automatic Control*, 40(11):1889–1904, 1995.
- [6] J. Dai and B. Prabhakar. The throughput of switches with and without speed-up. In *Proceedings of IEEE Infocom*, pages 556–564, 2000.
- [7] L. Georgiadis, M. Neely, and L. Tassiulas. *Resource Allocation and Cross-Layer Control in Wireless Networks*. Foundations and Trends in Networking, Now Publishers, 2006.
- [8] P. Giaccone, B. Prabhakar, and D. Shah. Randomized scheduling algorithms for high-aggregate bandwidth switches. *IEEE Journal on Selected Areas in Communications High-performance electronic switches/routers for high-speed internet*, 21(4):546–559, 2003.
- [9] G. Grimmett and D. Stirzaker. *Probability and Random Processes*. Oxford University Press, 2001.
- [10] J. M. Harrison. Brownian models of open processing networks: Canonical representation of workload. *Annals of Applied Probability*, 10:75–103, 2000.
- [11] F. Kelly and R. Williams. Fluid model for a network operating under a fair bandwidth-sharing policy. *The Annals of Applied Probability*, 14(3):1055–1083, 2004.
- [12] I. Keslassy and N. McKeown. Analysis of scheduling algorithms that provide 100% throughput in input-queued switches. In *Proceedings of Allerton Conference on Communication, Control and Computing*, 2001.
- [13] S. Kittipiyakul, T. Javidi, and V. Subramanian. Many-Sources Large Deviations for Max-Weight Scheduling. In *Communication, Control, and Computing, 2008 46th Annual Allerton Conference on*, pages 1495–1502, 2008.
- [14] N. McKeown. iSLIP: a scheduling algorithm for input-queued switches. *IEEE Transaction on Networking*, 7(2):188–201, 1999.
- [15] N. McKeown, V. Anantharam, and J. Walrand. Achieving 100% throughput in an input-queued switch. In *Proceedings of IEEE Infocom*, pages 296–302, 1996.
- [16] S. Rajagopalan, D. Shah, and J. Shin. A network adiabatic theorem: an efficient randomized protocol for contention resolution. In *ACM Sigmetrics/Performance*, 2009.
- [17] D. Shah and D. J. Wischik. Optimal scheduling algorithm for input queued switch. In *Proceeding of IEEE INFOCOM*, 2006.
- [18] D. Shah and D. J. Wischik. The teleology of scheduling algorithms for switched networks under light load, critical load, and overload. <http://web.mit.edu/~devavrat/www/shahwischik.pdf>, 2007-09.
- [19] A. Stolyar. Maxweight scheduling in a generalized switch: State space collapse and workload minimization in heavy traffic. *The Annals of Applied Probability*, 14(1):1–53, 2004.
- [20] A. Stolyar. Large deviations of queues sharing a randomly time-varying server. *Queueing Systems*, 59(1):1–35, 2008.
- [21] L. Tassiulas. Linear complexity algorithms for maximum throughput in radio networks and input queued switches. In *IEEE INFOCOM*, volume 2, pages 533–539, 1998.
- [22] L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Transactions on Automatic Control*, 37:1936–1948, 1992.
- [23] V. J. Venkataramanan and X. Lin. Structural properties of LDP for queue-length based wireless scheduling algorithms. In *45th Annual Allerton Conference on Communication, Control, and Computing*, September 2007.
- [24] V. J. Venkataramanan and X. Lin. On the large-deviations optimality of scheduling policies minimizing the drift of a lyapunov function. In *47th Annual Allerton Conference on Communication, Control, and Computing*, September 2009.
- [25] D. Shah, J. N. Tsitsiklis, and Y. Zhong. Qualitative properties of α -weighted scheduling policies. http://www.mit.edu/~zhyu4118/sig_paper.pdf.