

From Local to Global Stability in Stochastic Processing Networks through Quadratic Lyapunov Functions

A. B. Dieker* J. Shin†

September 30, 2012

Abstract

We construct a generic, simple, and efficient scheduling policy for stochastic processing networks, and provide a general framework to establish its stability. Our policy is randomized and prioritized: with high probability it prioritizes jobs which have been least routed through the network. We show that the network is globally stable under this policy if there exists an appropriate quadratic ‘local’ Lyapunov function that provides a negative drift with respect to nominal loads at servers. Applying this generic framework, we obtain stability results for our policy in many important examples of stochastic processing networks: open multiclass queueing networks, parallel server networks, networks of input-queued switches, and a variety of wireless network models with interference constraints. Our main novelty is the construction of an appropriate ‘global’ Lyapunov function from quadratic ‘local’ Lyapunov functions, which we believe to be of broader interest.

1 Introduction

The past few decades have witnessed a surge in interest on the design and analysis of scheduling policies for stochastic networks, e.g., [13, 24, 18, 19, 12, 5]. One of the key insights from this body of work is that natural scheduling policies can lead to instability even when each server is nominally underloaded [15, 17, 20]. (There are several notions of stability for stochastic networks, but they intuitively entail that, in some sense, the number of customers in the system does not grow without bounds.) This insight stimulated a search for tools that can characterize the stability regions of scheduling policies, i.e., the exact conditions on the arrival and service rates under which a network is stabilized by a policy.

It is the objective of this paper to study a question of a different kind: is it possible to construct a generic, simple, and efficient scheduling policy for stochastic processing networks, which leads to a (globally) stable network if all servers are (locally) nominally underloaded in some sense? To our knowledge, we are the first to answer this question within the setting of stochastic processing networks, which constitute a large class of stochastic networks capable of modeling a variety of networked systems for communication, manufacturing, and service systems (e.g., [12]). To investigate this question, the key is to determine if and how jobs from different parts of the network should be treated differently when they share the same buffer.

*H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta GA USA 30332. Email: ton.dieker@isye.gatech.edu.

†School of Computer Science, Georgia Institute of Technology, Atlanta GA USA 30332. Email: jshin72@cc.gatech.edu.

Various existing scheduling policies are ‘throughput optimal’ in the sense that they achieve the largest possible stability region, but these policies suffer from significant drawbacks. They typically obtain the desired stability by framing the contention resolution between buffers as an appropriate global optimization problem. This optimization problem requires central coordination between network entities, and it is computationally hard to solve if the network is large. The resulting policies, such as the max-weight policy [18, 23] and the back-pressure policy [24, 5], are not scalable and cannot cope efficiently with large networks. As a result, these throughput optimal policies do not provide a satisfactory answer to the aforementioned question.

The computational challenges surrounding existing throughput optimal policies motivate the search for easily implementable scheduling policies with provable performance guarantees but not necessarily with the throughput optimality property. This has led to the analysis of simple greedy scheduling policies for a variety of special classes of stochastic processing networks: open multiclass queueing networks [3], input-queued switches [6], and wireless network models [16]. These policies play an important role in this paper, since our most critical assumption roughly requires any local network component to be nominally underloaded under any ‘maximal’ greedy scheduling policy.

Our main contribution is a randomized scheduling policy for stochastic processing networks which only requires coordination within local components (e.g., service stations), and which is computationally attractive since it is a kind of priority policy. With high probability, our policy prioritizes jobs which have been least routed, and we therefore call our policy the ε -Least Routed First Served (ε -LRFS) policy. Here $\varepsilon \geq 0$ is a small number which helps to make the meaning of “high probability” precise.

Our main technical tool is a novel framework to construct a ‘global’ Lyapunov function for a stochastic processing network through appropriate ‘local’ Lyapunov functions. If the local Lyapunov functions yield stability of the corresponding ‘local’ network components, then the global Lyapunov function allows us to conclude that the whole network is stable. A critical feature of our framework is that the Lyapunov functions we work with are quadratic. Through examples, we show that quadratic local Lyapunov functions can readily be found for wide classes of networks. We refer to [8, 14] for other uses of quadratic Lyapunov functions.

Our approach to construct an appropriate ‘global’ Lyapunov function using ‘local’ quadratic Lyapunov functions contrasts with the popular fluid model methodology for establishing stability of stochastic networks [3, 9, 20]. The fluid model framework essentially reduces the question of stochastic stability to a question of a related deterministic (fluid) system. In the case of reentrant lines, our policy reduces to the First Buffer First Served (FBFS) policy, which has been proven to be stable via ‘inductive’ fluid arguments [7]. A similar fluid induction argument can be expected to work for general multiclass networks (modulo some technical arguments), but a fluid induction argument cannot be expected to work in general. A disadvantage compared to fluid models is that we have to keep track of detailed system behavior such as remaining service times, but therein also lies the power of our approach. The generality of our framework presents challenges to the use of fluid methods, and instead we work directly with a global Lyapunov function. The connection with fluid techniques is discussed in more detail in Section 4.3.

Although we believe that our methodology could be of much wider use, we have chosen to work out two special classes of stochastic processing networks in order to describe the implications of our techniques in relatively simple yet powerful settings: parallel server networks (including multiclass queueing networks) and communication network models (including networks of input-queued switches and wireless networks).

Our work is related to a paper by Bramson [2], who shows that open multiclass queueing networks can be stabilized by the Earliest Due Date First Served (EDDFS) policy under the ‘local’ condition that every processing unit is nominally underloaded. However, it is not known whether the EDDFS policy achieves similar results beyond the open multiclass queueing network. In fact, it is *a priori* unclear how to formulate a ‘local’ condition for general stochastic processing networks; our notion of ‘local’ Lyapunov function plays this role in the present paper.

This paper is organized as follows. Section 2 introduces the class of stochastic processing networks we study in this paper. Section 3 formally introduces the ε -LRFS policy and presents our main results. Section 4 describes our main idea using a simple network, the Rybko-Stolyar network [20]. We specialize our result to parallel server and communication network models in Section 5. All proofs are given in Section 6.

2 A Class of Stochastic Processing Networks

A stochastic processing network (SPN) consists of a set $\mathcal{J} = \{1, \dots, I\}$ of buffers, a set $\mathcal{J} = \{1, \dots, J\}$ of activities and a set $\mathcal{K} = \{1, \dots, K\}$ of processors. Each buffer has infinite capacity and holds jobs that await service. The SPNs we study in this paper have the feature that activity $j \in \mathcal{J}$ can only process jobs from a single buffer $i_j \in \mathcal{J}$, and that j requires simultaneous possession of a set $\mathcal{K}_j \subseteq \mathcal{K}$ of processors. Let \mathcal{J}_i be the set of activities capable of processing buffer i , i.e.,

$$\mathcal{J}_i = \{j \in \mathcal{J} : i_j = i\}.$$

We say that two buffers i and ℓ are *activity-interchangeable* if $\{\mathcal{K}_j : j \in \mathcal{J}_i\} = \{\mathcal{K}_j : j \in \mathcal{J}_\ell\}$. We also say that buffers i and ℓ are *processor-independent* if $\bigcup_{j \in \mathcal{J}_i} \mathcal{K}_j$ and $\bigcup_{j \in \mathcal{J}_\ell} \mathcal{K}_j$ are disjoint.

Network state. We let $Q_i(t) \in \mathbb{Z}_+$ be the queue length of buffer i at time t , i.e., the number of jobs waiting in buffer i excluding those being processed. We write $V_i(t) \in \mathbb{R}_+$ for the sum of the remaining service requirements over all jobs in buffer i which are currently being processed at time t . We use $\sigma_j(t)$ to denote the activity level of activity j at time t , where we assume that $\boldsymbol{\sigma}(t) = [\sigma_j(t)] \in \{0, 1\}^J$, meaning that each activity is either fully employed or not employed at all. Thus, we say that the network is *non-processor-splitting*.

Routing. After departing from a buffer i , a job joins buffer $\ell \in \mathcal{J}$ with probability $P_{i\ell}$ and departs from the network with probability $1 - \sum_{\ell} P_{i\ell}$ (independently of everything else). We write P for the $I \times I$ matrix of routing probabilities.

Resource allocation. Each activity j decreases the remaining service requirement of the job it is processing at rate $\beta_j > 0$ if $\sigma_j(t) = 1$. It is not allowed for an activity to be interrupted before it finishes the service requirement of the job it is working on, i.e., the network is *non-preemptive*. We do not allow for multiple activities to work on the same job simultaneously. Furthermore, we assume that each processor k has unit capacity, i.e.,

$$\sum_{j \in \mathcal{J} : k \in \mathcal{K}_j} \sigma_j(t) \leq 1.$$

We note that this unit capacity assumption is not restrictive since in our non-processor-splitting network of $\boldsymbol{\sigma}(t) \in \{0, 1\}^J$, a single processor k with capacity $c_k \in \mathbb{N}$ can be replaced by c_k copies

with unit capacity, and identical activity structure inherited from the original processor. We further define the service rate vector $s(\mathbf{u}) = [s_i(\mathbf{u})] \in \mathbb{R}_+^I$ for a (scheduling) vector $\mathbf{u} = [u_j] \in \{0, 1\}^J$ through

$$s_i(\mathbf{u}) := \sum_{j \in \mathcal{J}_i} u_j \beta_j.$$

External arrivals. There are external arrivals to at least one buffer. For $0 \leq s \leq t$, let $A_i(s, t)$ be the number of external jobs arriving at buffer i during the time interval $[s, t)$. We assume that $E[A_i(0, t)^2] < \infty$ for all $t < \infty$ (i.e., bounded second moment) and

$$\limsup_{t \rightarrow \infty} E[A_i(t, t+1) \mid \{A_1(0, s), \dots, A_I(0, s) : 0 \leq s \leq t\}] < \alpha_i,$$

for constant $\alpha_i \in (0, \infty)$. We note that we allow for dependencies in the random processes $A_i(\cdot)$ and $A_\ell(\cdot)$ for two buffers i and ℓ . Even though it is possible for the α_i to exceed the external arrival rates, it is convenient to interpret α_i as the external arrival rate at buffer i . Similarly abusing terminology, we let $\boldsymbol{\lambda} = [\lambda_i]$ be the *effective* arrival rate vector, i.e.,

$$\boldsymbol{\lambda} := (I + P + P^2 + \dots) \boldsymbol{\alpha},$$

where $\boldsymbol{\alpha} = [\alpha_i]$. We also say that all routes are bounded (in length) if $\inf \{d \in \mathbb{N} : P^d = 0\} < \infty$.

Service requirements. Once a job in buffer i is selected for processing by activity $j \in \mathcal{J}_i$, it requires service for a random amount of time. We assume that all service times are independent, and that they are independent of the routing and external arrival processes. We also suppose that the service time distribution only depends on the buffer from which the job is processed. Writing Γ_i for a generic processing time at buffer i , we assume that

$$E[\Gamma_i] = m_i \quad \text{and} \quad E[\Gamma_i^2] < \infty,$$

for constants $m_i \in (0, \infty)$. Let $\boldsymbol{\rho} = [\rho_i]$ denote the nominal load, i.e., $\rho_i = \lambda_i m_i$. We write $W_i(t)$ for the (expected) immediate workload in buffer i , which we define to be

$$W_i(t) = m_i Q_i(t) + V_i(t). \tag{1}$$

This is the expected amount of work in the i -th buffer given $Q_i(t)$ and $V_i(t)$. We note that our notion of immediate workload is the conditional expectation of a more common definition, and that immediate workload is defined for each buffer (as opposed to resource). We say that the network is *synchronized* if for all $s, t \geq 0$, $i \in \mathcal{J}$ and $j \in \mathcal{J}$,

$$A_i(s, t) = A_i(s, \lfloor t \rfloor), \quad \text{and} \quad \Gamma_i = m_i = \beta_j = 1 \quad \text{with probability 1.}$$

That is, in a synchronized network, arrivals and service completions only occur at integer time epochs.

Maximal scheduling policies. We say that activity j is *maximal* in a (scheduling) vector $\mathbf{u} = [u_j] \in \{0, 1\}^J$ if there exists a processor $k \in \mathcal{K}_j$ such that

$$\sum_{\ell \in \mathcal{J}: k \in \mathcal{K}_\ell} u_\ell = 1,$$

i.e., either activity j uses each of the processors in \mathcal{K}_j under schedule \mathbf{u} or it cannot be employed without violating the unit capacity constraint for some processor in \mathcal{K}_j . Given a non-negative vector $\mathbf{w} = [w_i] \in \mathbb{R}_+^I$, activity j is called *maximal* in \mathbf{u} with respect to \mathbf{w} if $w_{i_j} = 0$ or j is maximal in \mathbf{u} . An activity is non-maximal if it is not maximal. If all activities are maximal in \mathbf{u} with respect to \mathbf{w} , we simply say that \mathbf{u} is maximal with respect to \mathbf{w} . We write $\mathcal{M}(\mathbf{w})$ for the set of maximal scheduling vectors with respect to \mathbf{w} , i.e.,

$$\mathcal{M}(\mathbf{w}) := \{\mathbf{u} \in \{0, 1\}^J : \mathbf{u} \text{ is maximal with respect to } \mathbf{w}\}.$$

Finally, we say that a scheduling policy is maximal if, under the policy, $\sigma(t) \in \mathcal{M}(\mathbf{Q}(t))$ for all $t \geq 0$.

3 Main Result

This section describes the scheduling policies which play a central role in this paper, and presents our main stability result.

Scheduling policies. Our policies require that each job maintains a ‘counter’ for the number of times it has been routed so far, where we follow the convention that counters start from 1. The counter of a job is increased even when a job is routed to a buffer it has previously visited, so the counter of a job could differ from the number of different stations it has visited. We also consider a partition $\{\mathcal{J}^{(h)} : h \in \mathcal{H} = \{1, \dots, H\}\}$ of buffers into components, such that

$$\mathcal{J} = \bigcup_{h \in \mathcal{H}} \mathcal{J}^{(h)},$$

where any two buffers from different components are processor-independent. One possible choice for the partition is $\mathcal{J} = \mathcal{J}^{(1)}$, but, as becomes apparent from the description of our policies below, a finer partition makes our policies more ‘distributed’. It is important to note that we allow for routing between components. We assume that each component $\mathcal{J}^{(h)}$ for $h \in \mathcal{H}$ maintains a ‘timer’ $\mathcal{T}^{(h)}(t) \in [0, 1]$ at time t which decreases at unit rate if $\mathcal{T}^{(h)}(t) > 0$.

To describe our maximal scheduling policies, we need the following notation. We write $\mathbf{Q}^{(h)}(t) = [Q_i(t)I_{\mathcal{J}^{(h)}}(i)]$ for the queue length information in component $\mathcal{J}^{(h)}$, where $I_S \in \{0, 1\}$ is the indicator function of the set S , i.e.,

$$I_S(i) = \begin{cases} 1 & \text{if } i \in S \\ 0 & \text{otherwise} \end{cases}.$$

The following policy plays a key role throughout this paper.

Definition 3.1 (Least Routed First Served (LRFS) policy). *For each $h \in \mathcal{H}$, whenever a new arrival and/or service completion occurs at time t in component $\mathcal{J}^{(h)}$, execute the following algorithm immediately after all arrivals and service completions have occurred:*

1. Find the set Σ of non-maximal activities in $\sigma(t)$ with respect to $\mathbf{Q}^{(h)}(t)$.
2. Find a job with the smallest counter among those in buffers $\{i_j : j \in \Sigma\}$, where ties are broken arbitrarily.

3. Choose an arbitrary activity $j \in \Sigma$ to process the job identified in step 2, i.e., set $\sigma_j(t) = 1$.
4. Repeat steps 1–3 until $\sigma(t)$ is maximal with respect to $\mathbf{Q}^{(h)}(t)$.

We next introduce the maximal scheduling policy which is of primary interest in this paper. It uses a small parameter $\epsilon > 0$ in order to deal with unbounded route lengths.

Definition 3.2 (ϵ -Least Routed First Served (ϵ -LRFS) policy). For each $h \in \mathcal{H}$, whenever a new arrival and/or service completion occurs at time t in component $\mathcal{J}^{(h)}$, execute the following algorithm immediately after all arrivals and service completions have occurred:

1. Find a buffer $i \in \mathcal{J}^{(h)}$ containing a job with the largest counter among those in buffers $\mathcal{J}^{(h)}$.
2. Find the set of non-maximal activities $\Sigma \subset \mathcal{J}_i$ in $\sigma(t)$.
3. If $\Sigma \neq \emptyset$ and $\mathcal{T}^{(h)}(t) = 0$,
 - 3-1. With probability ϵ , choose an arbitrary activity $j \in \Sigma$ to process the job identified in step 1, i.e., set $\sigma_j(t) = 1$. With probability $1 - \epsilon$, do nothing.
 - 3-2. Set $\mathcal{T}^{(h)}(t) = 1$.
4. Execute the LRFS policy for component $\mathcal{J}^{(h)}$.

We remark that the ϵ -LRFS policy is identical to the LRFS policy when either $\epsilon = 0$ or $\mathcal{T}^{(h)}(t) > 0$, i.e., execution of the first three steps is not necessary in these cases. Since we assume that each timer decreases at unit rate, step 3-1 can be executed at most once per component in any time interval of unit length.

The network process. Write $Q_{i,c}(t)$ for the queue length of jobs with counter c in buffer i at time t . Let $V_{i,c}^j(t)$ be the remaining service requirement of the job with counter c in buffer i at time t if it is processed by activity j , and set $V_{i,c}^j(t) = 0$ if j is not processing a job with counter c . The network state is described by

$$X(t) = \left[Q_{i,c}(t), V_{i,c}^j(t), \mathcal{T}^{(h)}(t) : i \in \mathcal{J}, j \in \mathcal{J}, c \in \mathbb{N}, h \in \mathcal{H} \right] \in \Omega_X := \left(\mathbb{Z}_+^I \times \mathbb{R}_+^{I \times J} \right)^\infty \times [0, 1]^H.$$

Note that $X(t)$ does not encode information on the external arrival processes. In particular, $\{X(t)\}$ is non-Markovian in general. We impose the convention that $\{X(t)\}$ has right-continuous sample paths. We define a norm on Ω_X through

$$|X(t)| = \sum_{i \in \mathcal{J}, c \in \mathbb{N}} Q_{i,c}(t) + \sum_{i \in \mathcal{J}, c \in \mathbb{N}, j \in \mathcal{J}} V_{i,c}^j(t).$$

We assume that $X(0) \in \Omega_X^*$ where $\Omega_X^* = \{X \in \Omega_X : |X| < \infty\}$. In synchronized networks, one has $V_{i,c}^j(t) \in \{0, 1\}$ for all $t \in \mathbb{Z}_+$ under the ϵ -LRFS policy. We refer to the process $\{X(t) : t \in \mathbb{R}_+\}$ operating under the ϵ -LRFS policy as the ϵ -LRFS process. If $\epsilon = 0$, then we simply refer to this process as the LRFS process.

Network stability. This paper uses the following notion of stability.

Definition 3.3. The ε -LRFS process $\{X(t)\}$ is called *queue-length-stable* if

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \int_0^t E[|X(s)|] ds < \infty, \quad (2)$$

for any given initial state $X(0) \in \Omega_X^*$.

We establish the queue-length-stability by constructing appropriate Lyapunov functions. Under some additional assumptions on the arrival processes and service time distributions, these Lyapunov functions can also be used to establish positive recurrence of the ε -LRFS process, cf. Condition (A3) in [4]. It is outside of the main scope of the current paper to work out the details. We also remark that our proof can find an explicit finite constant for the right-hand side of (2), but this requires tedious bookkeeping and we therefore do not carry out this analysis.

To show the desired stability, we need the following notion of a ‘local’ Lyapunov function.

Definition 3.4. We say that $\mathcal{L}_{\text{local}} : \mathbb{R}_+^I \rightarrow \mathbb{R}_+$ is a local Lyapunov function with slack parameter $\varepsilon \geq 0$ if there exist constants $\mathcal{C} \geq 0$, $\eta > 0$ such that for every pair $(\mathbf{w}, \mathbf{u}) \in \mathbb{R}_+^I \times \{0, 1\}^J$ satisfying $\mathbf{u} \in \mathcal{M}(\mathbf{w})$,

$$\mathcal{L}_{\text{local}}(\mathbf{w} + \boldsymbol{\rho} + \varepsilon \mathbf{m} - s(\mathbf{u})) \leq \mathcal{L}_{\text{local}}(\mathbf{w}) - \eta \|\mathbf{w}\|_1 + \mathcal{C}, \quad (3)$$

where $\mathbf{m} = [m_i] \in \mathbb{R}^I$.

The reason for this nomenclature is that in a ‘local’ network, i.e., a network without routing ($P = 0$), the above inequality for $\mathcal{L}_{\text{local}}$ provides the desired negative drift condition in the Foster-Lyapunov criteria [10], which implies network stability. Here w_i , ρ_i , and $s_i(\mathbf{u})$ can be interpreted as the immediate workload, the external workload arrival rate and the workload processing rate at buffer i , respectively. We refer to Sections 4 and 5 for examples of local Lyapunov functions.

Now we are ready to state the main theorem of this paper, which establishes ‘global’ stability using a ‘local’ quadratic Lyapunov function.

Theorem 3.5. Suppose that there exists a symmetric matrix $Z \in \mathbb{R}_+^{I \times I}$ such that $\mathcal{L}_{\text{local}}(x) = x^T Z x$ is a local Lyapunov function with slack $\varepsilon > 0$. Then the ε -LRFS process is queue-length-stable if one of the following conditions C1 and C2 is satisfied:

C1. The network is synchronized and $Z_{i\ell} \neq 0$ only if buffers i and ℓ are in the same component.

C2. $Z_{i\ell} \neq 0$ only if buffers i and ℓ are activity-interchangeable and every buffer i has an associated activity $j \in \mathcal{J}_i$ with $|\mathcal{K}_j| = 1$.

Furthermore, if all routes are bounded, then the LRFS process is queue-length-stable if one of the conditions C1 or C2’ is satisfied, where

C2’. $Z_{i\ell} \neq 0$ only if buffers i and ℓ are activity-interchangeable.

Theorem 3.5 implies that if any maximal policy is ‘quadratic’ stable in a stochastic processing network without routing under the external load $\boldsymbol{\rho}$, then ε -LRFS is stable in a stochastic processing network with routing under nominal load $< \boldsymbol{\rho}$ for some small $\varepsilon > 0$. Section 4 describes the main idea of the proof, and a full proof is presented in Section 6.

We remark that the requirement of activity-interchangeable buffers in Condition C2 can be relaxed slightly. Our proof of Theorem 3.5 also works when the following relaxed condition C3 replaces C2.

C3. $Z_{i\ell} \neq 0$ only if for every $j \in \mathcal{J}_i$, there exists $j' \in \mathcal{J}_\ell$ such that $\mathcal{K}_{j'} \subseteq \mathcal{K}_j$, and vice versa (i.e., for every $j \in \mathcal{J}_\ell$, there exists $j' \in \mathcal{J}_i$ such that $\mathcal{K}_{j'} \subseteq \mathcal{K}_j$). Also, every buffer i has an associated activity $j \in \mathcal{J}_i$ with $|\mathcal{K}_j| = 1$.

A corresponding condition $C3'$ can also replace $C2'$, where $C3'$ does not require the second part of $C3$.

4 Proof Ideas for Theorem 3.5

In this section, we describe the main idea in the proof of Theorem 3.5. We first present it in a very special network, the Rybko-Stolyar network [20], which allows us to summarize the main idea of the proof at a high level. We subsequently describe the challenges that have to be overcome to establish our result in the general case, and discuss the feasibility of an approach based on fluid models.

4.1 Rybko-Stolyar Network

This network consists of four activities associated to four different buffers and two processors. The first processor is required for activities 1 and 4, and the second for activities 2 and 3. Each activity decreases the remaining service requirement of the job it is currently processing at unit rate (i.e., $\beta_j = 1$). Customers (or jobs) arrive at the first and third buffers, and traverse the buffers deterministically in the order $1 \rightarrow 2$ or $3 \rightarrow 4$. The service time is deterministic and equal to m_i for buffer i , and the external arrival processes (at the first and third buffers) are independent Poisson processes with rate 1. The network is given in Figure 1, and a necessary condition for stability is

$$\rho_1 = m_1 + m_4 < 1 \quad \text{and} \quad \rho_2 = m_2 + m_3 < 1. \quad (4)$$




Figure 1: Illustration of the Rybko-Stolyar Network. In this diagram, there are four jobs in the system. Jobs are currently being processed and waiting for service are colored black and white, respectively.

The network state can be described by $X(t) = \{Q_i(t), V_i(t) : i = 1, 2, 3, 4\} \in \Omega_X := \mathbb{Z}_+^4 \times \mathbb{R}_+^4$, where $Q_i(t)$ is the queue length (i.e., the number of jobs waiting) of buffer i at time t and $V_i(t)$ is the remaining service time of the job currently being processed in buffer i at time t . Hence $V_i(t) = 0$ if no job is being processed in buffer i . We define a norm through $|X(t)| = \sum_i (V_i(t) + Q_i(t))$. In this network, it is known [20] that if processor 1 prioritizes buffer 4 and processor 2 prioritizes buffer 2, the network can be unstable even when the necessary condition (4) holds. On the other hand, LRFS prioritizes buffer 1 and 3, so it reduces to the First Buffer

First Served policy. This is known to be stable under the necessary stability condition (4). We next derive this stability result using our main idea.

Assuming LRFS and (4), we construct an appropriate Lyapunov function $\mathcal{L}_{\text{global}} : \Omega_X \rightarrow \mathbb{R}_+$ satisfying

$$E \left[\mathcal{L}_{\text{global}}(X(t+1)) - \mathcal{L}_{\text{global}}(X(t)) \mid X(t) \right] \leq -\gamma|X(t)| + \mathcal{C}, \quad \text{almost surely,} \quad (5)$$

where $\gamma, \mathcal{C} > 0$ are some constants. By first taking expectations with respect to the distribution of $X(t)$ and then integrating over t on both sides of the above inequality, we conclude that

$$\begin{aligned} \limsup_{t \rightarrow \infty} \frac{1}{t} \int_0^t E[|X(s)|] ds &\leq \limsup_{t \rightarrow \infty} \frac{1}{\gamma} \left(\mathcal{C} - \frac{\int_t^{t+1} E[\mathcal{L}_{\text{global}}(X(s))] ds - \int_0^1 E[\mathcal{L}_{\text{global}}(X(s))] ds}{t} \right) \\ &\leq \limsup_{t \rightarrow \infty} \frac{1}{\gamma} \left(\mathcal{C} + \frac{\int_0^1 E[\mathcal{L}_{\text{global}}(X(s))] ds}{t} \right) \\ &= \frac{\mathcal{C}}{\gamma} < \infty. \end{aligned} \quad (6)$$

We now proceed toward constructing the ‘global’ Lyapunov function $\mathcal{L}_{\text{global}}$ satisfying (5) based on a ‘local’ quadratic Lyapunov function $\mathcal{L}_{\text{local}}$. To this end, we first discuss how to construct the ‘local’ quadratic Lyapunov function. Consider a single-processor system with two buffers a and b , deterministic service times given by m_a and m_b , and independent Poisson arrival processes with rate 1 at each buffer (i.e., the total rate is 2). Hence, a necessary condition for stability is

$$m_a + m_b < 1.$$

Under a maximal (i.e., work-conserving) scheduling policy, the workload $W(t) = \sum_i W_i(t)$ at time t satisfies

$$W(t+1) = W(t) + \sum_{i=a,b} m_i A_i(t, t+1) - \int_t^{t+1} \mathbf{1}_{W(u)}^+ du,$$

where $W_i(t)$ is defined as the immediate workload at buffer i as in (1), $A_i(s, t)$ is the number of jobs arriving at buffer i during the time interval $[s, t)$ so that $E[A_i(s, t)] = t - s$ and we define

$$\mathbf{1}_x^+ = \begin{cases} 1 & \text{if } x > 0 \\ 0 & \text{otherwise} \end{cases}.$$

Hence, we have

$$E \left[W(t+1) - W(t) \mid W(t) \right] \leq m_a + m_b - \mathbf{1}_{W(t)-1}^+, \quad (7)$$

where we use that $\mathbf{1}_{W(u)}^+ \geq \mathbf{1}_{W(t)-1}^+$ for $u \in [t, t+1]$ since we assume deterministic service times.

Using this, it follows that for some finite constant \mathcal{C} ,

$$\begin{aligned}
E \left[W(t+1)^2 - W(t)^2 \mid W(t) \right] &= E \left[2W(t)(W(t+1) - W(t)) \mid W(t) \right] \\
&\quad + E \left[(W(t+1) - W(t))^2 \mid W(t) \right] \\
&\leq 2W(t)E \left[W(t+1) - W(t) \mid W(t) \right] + \mathcal{C} \\
&\leq 2W(t) \left(m_a + m_b - \mathbf{1}_{W(t)-1}^+ \right) + \mathcal{C} \\
&< -2(1 - m_a - m_b)W(t) + \mathcal{C} + 2, \tag{8}
\end{aligned}$$

where we use $E[A_i(t, t+1)^2] < \infty$ for the first inequality. This shows that $W(t)^2 = (W_1(t) + W_2(t))^2$ is a suitable Lyapunov function for our ‘local’ single-processor system under the necessary stability condition $m_a + m_b < 1$.

This observation on the single-processor system motivates the following quadratic local Lyapunov function $\mathcal{L}_{\text{local}} : \mathbb{R}_+^I \rightarrow \mathbb{R}_+$ for the Rybko-Stolyar network:

$$\mathcal{L}_{\text{local}}(x) = x^T Z x = x^T \begin{pmatrix} 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 \\ 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 \end{pmatrix} x = (x_1 + x_4)^2 + (x_2 + x_3)^2, \quad \text{for } x = (x_1, x_2, x_3, x_4). \tag{9}$$

One can easily check that it satisfies (3) with (small) slack $\varepsilon > 0$ under the necessary stability requirements $m_1 + m_4 < 1$ and $m_2 + m_3 < 1$. We propose the following global Lyapunov function $\mathcal{L}_{\text{global}}$:

$$\mathcal{L}_{\text{global}}(X(t)) = \mathcal{L}_{\text{local}}(W_1(t), V_2(t), W_3(t), V_4(t)) + \xi \mathcal{L}_{\text{local}}(W_1(t), \widehat{W}_2(t), W_3(t), \widehat{W}_4(t)),$$

where the new parameters ξ and $\widehat{W}_i(t)$ shall be defined explicitly. We remind the reader that our goal is to prove (5).

First, a similar calculation as for the single-processor case in (7) yields that under the LRFS policy,

$$\begin{aligned}
E \left[W_1(t+1) + V_4(t+1) - W_1(t) - V_4(t) \mid X(t) \right] &\leq -\mathbf{1}_{W_1(t)+V_4(t)-1}^+ + m_1, \\
E \left[W_3(t+1) + V_2(t+1) - W_3(t) - V_2(t) \mid X(t) \right] &\leq -\mathbf{1}_{W_3(t)+V_2(t)-1}^+ + m_3.
\end{aligned}$$

Hence, as for (8), one can conclude that for some constant \mathcal{C} ,

$$E \left[(W_1(t+1) + V_4(t+1))^2 - (W_1(t) + V_4(t))^2 \mid X(t) \right] \leq -2(1 - m_1)(W_1(t) + V_4(t)) + \mathcal{C}, \tag{10}$$

$$E \left[(W_3(t+1) + V_2(t+1))^2 - (W_3(t) + V_2(t))^2 \mid X(t) \right] \leq -2(1 - m_3)(W_3(t) + V_2(t)) + \mathcal{C}, \tag{11}$$

where the precise value of \mathcal{C} can be different from line to line.

We note that the sum $(W_1(t) + V_4(t))^2 + (W_3(t) + V_2(t))^2$ is not a suitable choice for $\mathcal{L}_{\text{global}}$ since it does not include $Q_2(t)$ and $Q_4(t)$ (or $W_2(t)$ and $W_4(t)$). To address this issue, we further use

$$\widehat{W}_2(t) := W_2(t) + m_2 Q_1(t) \quad \text{and} \quad \widehat{W}_4(t) := W_4(t) + m_4 Q_3(t).$$

We refer to \widehat{W}_2 and \widehat{W}_4 as the total workload in buffer 2 and 4, respectively. Using this notation, one finds that under the LRFS policy,

$$\begin{aligned} E \left[W_1(t+1) + \widehat{W}_4(t+1) - W_1(t) - \widehat{W}_4(t) \mid X(t) \right] &\leq -\mathbf{1}_{W_1(t)+W_4(t)-1}^+ + m_1 + m_4, \\ E \left[W_3(t+1) + \widehat{W}_2(t+1) - W_3(t) - \widehat{W}_2(t) \mid X(t) \right] &\leq -\mathbf{1}_{W_2(t)+W_3(t)-1}^+ + m_2 + m_3. \end{aligned}$$

The above equalities can be used to obtain ‘negative drift terms’ for $W_2(t)$ and $W_4(t)$, which are missing in (10) and (11). Namely, for some constant \mathcal{C} , we obtain

$$\begin{aligned} E \left[\left(W_1(t+1) + \widehat{W}_4(t+1) \right)^2 - \left(W_1(t) + \widehat{W}_4(t) \right)^2 \mid X(t) \right] \\ \leq 2 \left(W_1(t) + \widehat{W}_4(t) \right) E \left[W_1(t+1) + \widehat{W}_4(t+1) - W_1(t) - \widehat{W}_4(t) \mid X(t) \right] + \mathcal{C} \\ \leq 2 \left(W_1(t) + \widehat{W}_4(t) \right) \left(-\mathbf{1}_{W_1(t)+W_4(t)-1}^+ + m_1 + m_4 \right) + \mathcal{C} \\ \leq -2(1 - m_1 - m_4) (W_1(t) + W_4(t)) + 2(m_1 + m_4)m_4 Q_3(t) + \mathcal{C} + 2 \\ \leq -2(1 - \rho_{\max}) (W_1(t) + W_4(t)) + 2m^* \rho_{\max} W_3(t) + \mathcal{C} + 2, \end{aligned} \quad (12)$$

where $\rho_{\max} = \max\{\rho_1, \rho_2\}$ and $m^* = \max_{i,j} \left\{ \frac{m_i}{m_j} \right\}$. Similarly,

$$\begin{aligned} E \left[\left(W_3(t+1) + \widehat{W}_2(t+1) \right)^2 - \left(W_3(t) + \widehat{W}_2(t) \right)^2 \mid X(t) \right] \\ \leq -2(1 - \rho_{\max}) (W_3(t) + W_2(t)) + 2m^* \rho_{\max} W_1(t) + \mathcal{C} + 2. \end{aligned} \quad (13)$$

Observe that there are positive terms $W_3(t)$ and $W_1(t)$ in (12) and (13), respectively. The key idea behind our proof is that the positive terms can be canceled out by appropriately summing (10), (11), (12) and (13). Indeed, we define the desired Lyapunov function $\mathcal{L}_{\text{global}}$ as

$$\mathcal{L}_{\text{global}}(X) = (W_1 + V_4)^2 + (W_3 + V_2)^2 + \xi \left(W_1 + \widehat{W}_4 \right)^2 + \xi \left(W_3 + \widehat{W}_2 \right)^2,$$

where we choose $\xi = \frac{1-\rho_{\max}}{2m^* \rho_{\max}}$. Combining (10), (11), (12) and (13), we conclude that

$$E \left[\mathcal{L}_{\text{global}}(X(t+1)) - \mathcal{L}_{\text{global}}(X(t)) \mid X(t) \right] \leq -2\xi(1 - \rho_{\max})|X(t)| + \mathcal{C}.$$

This completes the proof of (5), and hence the desired stability (6).

4.2 Beyond the Rybko-Stolyar Network

The preceding subsection presents the main idea behind our construction of a ‘global’ Lyapunov function $\mathcal{L}_{\text{global}}$ using a ‘local’ Lyapunov function $\mathcal{L}_{\text{local}}$ (i.e., from the single-processor system) in

the specific example of the Rybko-Stolyar network. The construction of $\mathcal{L}_{\text{global}}$ relies on summing $\mathcal{L}_{\text{local}}$ terms inductively by exploring certain maximality properties of the LRFS policy at each iteration. In general networks there are several difficulties which do not arise in the Rybko-Stolyar network, and this section discusses the ideas and arguments needed to overcome them.

A first challenge we have overcome arises in networks with unbounded route lengths (i.e., $\inf\{d \in \mathbb{N} : P^d = 0\} = \infty$). In that case, the above inductive procedure does not terminate. For this reason, we propose a variant of the LRFS policy, the ε -LRFS policy, which occasionally processes a job with the largest counter. Intuitively speaking, this additional mechanism in ε -LRFS can control the jobs with large counters, whereas LRFS cannot.

A second challenge we have surmounted is that the construction of $\mathcal{L}_{\text{global}}$ in the Rybko-Stolyar network starts from a simple local Lyapunov function in a single-server system, but it is not clear whether similar arguments go through for general local Lyapunov functions and stochastic processing networks. We require Condition *C2* to resolve this issue. It is readily seen that the local Lyapunov function (9) used in the Rybko-Stolyar network satisfies this condition. The condition can be relaxed under some additional conditions on the arrival processes and service time distributions. For example, in synchronized networks, Condition *C1* can be used instead of *C2*.

A further challenge in the general case relates to the definition of $\widehat{W}_i(t)$. In the Rybko-Stolyar network, it is the sum of workloads along a path of buffers, with i as the last buffer. This definition only applies to networks with deterministic routing. In the general case we use several notions of total workload. To allow for stochastic routing, we construct a new process $\{Y(t)\}$ from $\{X(t)\}$ with deterministic routing. This process is essentially identical to $\{X(t)\}$, but we enlarge the state space to incorporate routing information. We construct a Lyapunov function $\mathcal{L}_{\text{global}}$ for $\{Y(t)\}$, which we use to establish the stability of $\{Y(t)\}$ and hence the stability of $\{X(t)\}$.

In summary, we construct the Lyapunov function $\mathcal{L}_{\text{global}}$ for general networks as the sum of three parts:

$$\mathcal{L}_{\text{global}}(Y(t)) = \sum_{c=1}^D \left(\frac{v}{2\mathcal{C}}\right)^c \mathcal{L}_{\text{local}}\left(\widehat{W}_{\leq c}(t)\right) + \frac{2\mathcal{C}}{\beta_{\min}} \|\mathbf{V}(t)\|_2^2 + \frac{\xi}{2\mathcal{C}} \mathcal{G}(Y(t))^2.$$

The specific notation used here is not important; we refer readers to Section 6 for the definitions used. To prove stability, we need to argue that this function satisfies a so-called negative-drift condition. The first term, i.e., the finite sum, comes from the inductive construction under LRFS, appropriately truncated. For the Rybko-Stolyar network, this is the only part we need. The first part produces the desired negative drift for jobs with low counters, but it gives a positive drift in terms of remaining service requirements $\mathbf{V}(t)$ as a by-product (albeit not in the Rybko-Stolyar network under the assumptions of the preceding subsection). The second term in our Lyapunov function ($\|\mathbf{V}(t)\|_2^2$) has a negative drift and compensates the positive drift incurred by the first term. The third term in our Lyapunov function ($\mathcal{G}(Y(t))^2$) controls the high-counter jobs under the mechanism which is present in the ε -LRFS policy but not in the LRFS policy (step 3 in Definition 3.2). This additional mechanism allows us to establish a negative drift for the last term. By appropriately weighing each of the three terms, we derive the desired negative drift condition for the Lyapunov function $\mathcal{L}_{\text{global}}$.

4.3 Connection with Fluid Models

As mentioned in the previous subsection, our approach relies on an inductive argument based on job counters. For the Rybko-Stolyar network and more generally for multiclass networks, fluid models can be used to give relatively simple proofs of our results. Thus, a more detailed discussion on the connection with fluid models together with its pros and cons is warranted.

The fluid approach consists of two main steps. In the first step, by scaling time and space, one proves convergence of the queueing process to the solution of a system of deterministic equations known as the *fluid model*. In the second step, one proves that this fluid model is stable, i.e., that it eventually reaches the origin. Stability of the fluid model can be established through the construction of a Lyapunov function for the fluid model, or in some cases one can obtain fluid stability through direct methods such as induction. Once fluid stability has been established, one can apply general theorems to deduce that the stochastic model is also stable (in a certain sense), see for instance [3].

It might be possible to establish existence of a fluid model and to prove that the stochastic model converges to the fluid model in the setting of the present paper, and that our ‘global’ Lyapunov function might work to prove fluid stability. Comparing our approach with this proof strategy, a disadvantage of the fluid model is that one needs to establish convergence to the fluid model, while a disadvantage of our approach is that we have to keep track of detailed state information such as residual service times.

Another possible approach to establish fluid stability is to use an inductive argument, which may seem particularly attractive given our construction of job counters and the suitability of an induction argument in existing work on fluid models [7]. However, this approach has inherent challenges. The base step in an inductive approach could use the ‘local’ Lyapunov function $\mathcal{L}_{\text{local}}$ to argue that the fluid level of jobs with counter 1 vanishes after some finite time T_1 . It would then use $\mathcal{L}_{\text{local}}$ to argue that the fluid level of jobs with counter 2 vanishes after some finite time $T_2 > T_1$, and so forth. To carry out this argument, one has to show that $\mathcal{L}_{\text{local}}$ satisfies a certain negative-drift condition under the assumption that high-priority counter 1 jobs vanish on a fluid scale. The latter only yields a guarantee on the ‘average’ or ‘long-run’ behavior of the jobs with counter 1, whereas one needs ‘short-term’ network state information to establish the negative-drift condition for jobs with counter 2. Indeed, under our scheduling policy, jobs with counter 1 (even when vanishing on a fluid scale) can significantly influence the dynamics of jobs with counter 2 depending on the complexity of the network. Therefore, the base of the induction approach is too weak to be used in the induction step for general networks since one needs more detailed information than the time-average given by the fluid approach.

In special cases such as multiclass networks, one may not need quadratic Lyapunov functions and it may be possible to establish the stability of our counter-based policy using fluid induction without quadratic Lyapunov functions. However, in general (e.g., for networks of switches), we need quadratic Lyapunov functions since they are the only available tool to establish stability for single-hop networks.

5 Examples

In this section, we provide applications of Theorem 3.5 to various special stochastic processing networks. We consider parallel server networks (including multiclass queueing networks) in Section 5.1 and communication networks (including wireless networks and networks of input-queued switches) in Section 5.2. They are examples of non-synchronized and synchronized

networks, respectively. In all of these important examples, suitable local Lyapunov functions are easy to find.

5.1 Open Multiclass Queueing Networks and Parallel Server Networks

In this section, we consider special stochastic processing networks known as *parallel server networks*. These networks are characterized by the following assumption.

A1. Each activity is processed by exactly one processor and processes exactly one buffer, i.e.,

$$|\mathcal{K}_j| = 1, \quad \text{for all } j \in \mathcal{J}.$$

Figure 2 illustrates the relations between buffers, activities and processors in parallel server networks. Our notion of ‘parallel server network’ generalizes the well-studied parallel server systems [11] by adding stochastic routing dynamics between buffers. It also includes open multiclass queueing networks [3] as a special case, which additionally require

$$|\mathcal{J}_i| = 1, \quad \text{for all } i \in \mathcal{J}. \tag{14}$$

In open multiclass queueing networks, buffers and activities are in one-to-one correspondence and they are referred to as *classes*. The Rybko-Stolyar in Section 4 is an instance of open multiclass queueing networks.




Figure 2: Two examples of parallel server networks. They do not satisfy Assumption A2. The leftmost diagram illustrates a parallel server network with five buffers $\{a, b, c, d, e\}$, three processors $\{1, 2, 3\}$ and eight activities $\{(a, 1), (b, 1), (b, 2), (b, 3), (c, 2), (d, 3), (e, 2), (e, 3)\}$. Once a job in buffer a, b, c, d completes its service requirement, it joins buffer b, c, d, e (i.e., $P_{ab} = P_{bc} = P_{cd} = P_{de} = 1$), respectively. Once a job in buffer e completes its service requirement, it leaves the network. The rightmost diagram illustrates a parallel server network consisting of three ‘local’ parallel server systems where jobs are routed between local systems.

A parallel server network naturally defines a bipartite graph $(\mathcal{J}, \mathcal{K}, \mathcal{J})$ such that each activity in \mathcal{J} defines an edge between buffers \mathcal{J} and processors \mathcal{K} . Requirement (14) of open multiclass queueing networks imposes the additional restriction that each vertex in \mathcal{J} has degree one. We further consider the following strengthening of Assumption A1.

A2. $(\mathcal{J}, \mathcal{K}, \mathcal{J})$ is a union of disjoint complete bipartite graphs $\{(\mathcal{J}^{(h)}, \mathcal{K}^{(h)}, \mathcal{J}^{(h)}) : h \in \mathcal{H}\}$, i.e.,

$$\cup_{h \in \mathcal{H}} \mathcal{J}^{(h)} = \mathcal{J}, \quad \cup_{h \in \mathcal{H}} \mathcal{K}^{(h)} = \mathcal{K}, \quad \text{and} \quad \cup_{h \in \mathcal{H}} \mathcal{J}^{(h)} = \mathcal{J}.$$

This assumption implies that two buffers in the same component are activity-interchangeable, even though they may differ with respect to routing, external arrivals or service requirements. One can easily check that open multiclass queueing networks always satisfy this assumption, while the parallel server networks in Figure 2 do not. Assumption A2 is useful because it enables us to establish a necessary condition for stability and it allows us to find a suitable local Lyapunov function satisfying Condition C2 of Theorem 3.5. However, Theorem 3.5 is applicable to general networks as long as one can find a ‘good’ local Lyapunov function satisfying Condition C2. Figure 3 gives examples of parallel server networks satisfying Assumption A2.




Figure 3: Two examples of parallel server networks satisfying Assumption A2. The rightmost diagram illustrates the Rybko-Stolyar network described in Section 4.

Necessary condition for stability. We now aim to obtain a necessary condition for stability of a parallel server network. Under Assumption A2, a necessary condition to stabilize the network is that for every n ,

$$\rho^{(h)} := \sum_{i \in \mathcal{J}^{(h)}} \rho_i < \beta^{(h)} := \sum_{j \in \mathcal{J}^{(h)}} \beta_j. \quad (15)$$

It is clear that the above condition is required for stability since $\rho^{(h)}$ and $\beta^{(h)}$ describe the total nominal load and the maximum processing rate, respectively, at the local component $(\mathcal{J}^{(h)}, \mathcal{K}^{(h)}, \mathcal{J}^{(h)})$.

Local Lyapunov function. As in Section 4, the single-processor example is the main building block. We define the local Lyapunov function as

$$\mathcal{L}_{\text{local}}(x) = \sum_{h \in \mathcal{H}} \left(\sum_{i \in \mathcal{J}^{(h)}} x_i \right)^2 \quad \text{for } x = (x_1, \dots, x_I). \quad (16)$$

We now show that this function satisfies (3) with some slack $\varepsilon \geq 0$ as long as the necessary condition (15) for stability is satisfied. For given vectors $\mathbf{w} = [w_i] \in \mathbb{R}_+^I$ and $\mathbf{u} = [u_j] \in \mathcal{M}(\mathbf{w})$,

maximality implies that, on writing $w^{(h)} = \sum_{i \in \mathcal{J}^{(h)}} w_i$,

$$s^{(h)} = s^{(h)}(\mathbf{u}) := \sum_{j \in \mathcal{J}^{(h)}} u_j \beta_j \geq \sum_{j \in \mathcal{J}^{(h)}} \mathbf{1}_{w^{(h)}}^+ \beta_j = \mathbf{1}_{w^{(h)}}^+ \beta^{(h)},$$

where we use Assumption A2 and we recall that $\mathbf{1}_x^+ = 1$ if $x > 0$, and $\mathbf{1}_x^+ = 0$ otherwise. Thus, we have

$$\begin{aligned} \mathcal{L}_{\text{local}}(\mathbf{w} + \boldsymbol{\rho} + \varepsilon \mathbf{m} - s(\mathbf{u})) - \mathcal{L}_{\text{local}}(\mathbf{w}) &= \sum_{h \in \mathcal{H}} \left[\left(w^{(h)} + \rho^{(h)} + \varepsilon m^{(h)} - s^{(h)} \right)^2 - \left(w^{(h)} \right)^2 \right] \\ &\leq \mathcal{C} + 2 \sum_{h \in \mathcal{H}} w^{(h)} \left(\rho^{(h)} + \varepsilon m^{(h)} - s^{(h)} \right) \\ &\leq \mathcal{C} + 2 \sum_{h \in \mathcal{H}} w^{(h)} \left(\rho^{(h)} + \varepsilon m^{(h)} - \mathbf{1}_{w^{(h)}}^+ \beta^{(h)} \right) \\ &= \mathcal{C} + 2 \sum_{h \in \mathcal{H}} w^{(h)} \left(\rho^{(h)} + \varepsilon m^{(h)} - \beta^{(h)} \right), \end{aligned}$$

where \mathcal{C} is some constant and we define $m^{(h)} := \sum_{i \in \mathcal{J}^{(h)}} m_i$. Therefore, $\mathcal{L}_{\text{local}}$ is a local Lyapunov function with slack ε for

$$0 \leq \varepsilon < \min_{h \in \mathcal{H}} \frac{\beta^{(h)} - \rho^{(h)}}{m^{(h)}},$$

where the right-hand side is positive if (15) holds.

Stability of LRFS policies. We now formulate the main results of this paper for open multi-class networks and parallel server networks. Under Assumption A2, the local Lyapunov function (16) satisfies Condition C2 of Theorem 3.5. Therefore, we obtain the following proposition as a corollary. We remind the reader that open multiclass queueing networks are special instances of parallel server networks, and that Assumption A2 automatically holds for these networks.

Proposition 5.1. *If a stochastic processing network satisfies Assumption A2 with $\rho^{(h)} < \beta^{(h)}$ for all n , then*

- *The ε -LRFS process is queue-length-stable for any $\varepsilon \in \left(0, \min_{h \in \mathcal{H}} \frac{\beta^{(h)} - \rho^{(h)}}{m^{(h)}} \right)$.*
- *The LRFS process is queue-length-stable if all routes are bounded in length.*

We note that the ε -LRFS policy admits a simpler description in a stochastic processing network satisfying Assumption A2, since a job can be processed by any processor in the partition, i.e., Σ in Definition 3.2 is non-empty whenever a processor is idle and capable of processing a job. Indeed, the ε -LRFS policy reduces to the following work-conserving randomized priority policy: whenever a processor k is idle at time t and there are jobs capable of being processed by k ,

- Process a job with the smallest counter with probability $1 - \varepsilon \left(1 - \mathbf{1}_{\mathcal{J}^{(h)}(t)}^+ \right)$, otherwise process a job with the largest counter.
- Set $\mathcal{J}^{(h)}(t) = 1$ if $\mathcal{J}^{(h)}(t) = 0$,

where $\mathcal{J}^{(h)}$ is the (local) component of buffers associated with processor k . Proposition 5.1 implies that the ε -LRFS policy can achieve ‘almost’ the full capacity region (15) by choosing a small $\varepsilon > 0$.

Our proof of Proposition 5.1 provides a different proof for some results that have been established using fluid model techniques. For example, in reentrant lines, the ε -LRFS policy for $\varepsilon = 0$ is identical to the well-known First Buffer First Served (FBFS) policy. Our proposition implies that the FBFS policy is throughput optimal in all reentrant lines, which has been proved originally in [7].

5.2 Communication Networks

We now consider examples of synchronized stochastic processing networks described in Section 2, i.e., $m_i = \beta_j = 1$, for all $i \in \mathcal{I}, j \in \mathcal{J}$. In particular, we consider the following additional assumption on synchronized stochastic processing networks.

B1. Each buffer has exactly one associated activity, i.e.,

$$\mathcal{J}_i = \{j_i\} \quad \text{for all } i \in \mathcal{I}.$$

Hence, we write $\mathcal{K}_i := \mathcal{K}_{j_i}$.

We again remark that Assumption *B1* facilitates a suitable local Lyapunov function for Theorem 3.5. However, even if Assumption *B1* does not hold, Theorem 3.5 is applicable to synchronized stochastic processing networks as long as one can find a ‘good’ local Lyapunov function. Synchronized stochastic processing networks satisfying Assumption *B1* include various communication network models of unit-sized packets: networks of input-queued switches [18, 6], wireless network models with primary interference constraints [21] and independent-set interference constraints [22]. We refer the corresponding references for detailed descriptions of the network models. As a concrete example, we write out the details of the wireless network model with primary interference constraints.

Wireless networks with primary interference constraints. Consider a network of n nodes represented by $V = \{1, \dots, n\}$ and a set of directed paths $\{P_1, P_2, \dots\}$. Unit-size packets arrive at the ingress node of each path as per an exogenous arrival process. Assume that the network is synchronized, i.e., each packet departs from a node at time $t \in \mathbb{Z}_+$ and arrives at the next node on its route at time $t + 1$. The primary interference constraint means that each node can either send or receive (it cannot do both) one packet at the time. A scheduling policy (or algorithm) decides which packets transmit at each (discrete) time instance. Figure 4 illustrates a wireless network of four nodes with primary interference constraints.

Necessary condition for stability. As in the parallel server networks, one can obtain the following necessary condition to stabilize a stochastic processing network satisfying Assumption *B1*: for all $k \in \mathcal{K}$,

$$\rho^{(k)} := \sum_{i \in \mathcal{I}: k \in \mathcal{K}_i} \rho_i < 1. \quad (17)$$

This is because Assumption *B1* implies that each buffer has at most one associated activity, and the processing rate is 1.




Figure 4: Example of a wireless network with primary interference constraints with four nodes $\{1, 2, 3, 4\}$ and five paths $\{1 \rightarrow 2, 2 \rightarrow 1, 1 \rightarrow 3, 3 \rightarrow 1, 2 \rightarrow 4 \rightarrow 3\}$. There is a buffer on each (directed) edge on each path, i.e., six buffers in total. Unit-size packets arrive at the ingress buffer (i.e., the first node) of each path. In the leftmost diagram, two links $2 \rightarrow 4$ and $3 \rightarrow 1$ are transmitting (unit-sized) packets. Once a packet is transmitted, it leaves the network if it arrives at the destination node (i.e., the last node on its path). The rightmost diagram illustrates the corresponding stochastic processing network, where the relation between buffers, activities and processors induces a hypergraph (each buffer requires two processors)

Local Lyapunov function. We consider the following local Lyapunov function:

$$\mathcal{L}_{\text{local}}(x) = \sum_{i \in \mathcal{J}} \sum_{k \in \mathcal{K}_i} \sum_{\ell \in \mathcal{J}: k \in \mathcal{K}_\ell} x_i x_\ell, \quad \text{for } x = (x_1, \dots, x_I).$$

We now proceed toward proving that condition (3) holds. For given vectors $\mathbf{w} = [w_i] \in \mathbb{R}_+^I$ and $\mathbf{u} = [u_j] \in \mathcal{M}(\mathbf{w})$, maximality implies that

$$\sum_{k \in \mathcal{K}_i} \sum_{\ell \in \mathcal{J}: k \in \mathcal{K}_\ell} u_\ell \geq \mathbf{1}_{w_i}^+.$$

Thus, we have that, on writing $\nu_i = \rho_i + \varepsilon m_i$,

$$\begin{aligned} \mathcal{L}_{\text{local}}(\mathbf{w} + \boldsymbol{\rho} + \varepsilon \mathbf{m} - s(\mathbf{u})) - \mathcal{L}_{\text{local}}(\mathbf{w}) &\leq \mathcal{C} + 2 \sum_{i \in \mathcal{J}} w_i \left(\sum_{k \in \mathcal{K}_i} \sum_{\ell \in \mathcal{J}: k \in \mathcal{K}_\ell} \nu_\ell - \sum_{k \in \mathcal{K}_i} \sum_{\ell \in \mathcal{J}: k \in \mathcal{K}_\ell} u_\ell \right) \\ &\leq \mathcal{C} + 2 \sum_{i \in \mathcal{J}} w_i \left(\sum_{k \in \mathcal{K}_i} \sum_{\ell \in \mathcal{J}: k \in \mathcal{K}_\ell} \nu_\ell - \mathbf{1}_{w_i}^+ \right) \\ &= \mathcal{C} + 2 \sum_{i \in \mathcal{J}} w_i \left(\sum_{k \in \mathcal{K}_i} \sum_{\ell \in \mathcal{J}: k \in \mathcal{K}_\ell} \nu_\ell - 1 \right), \end{aligned}$$

where \mathcal{C} is some constant. Therefore, (3) holds if $\nu^{(k)} := \sum_{i \in \mathcal{J}: k \in \mathcal{K}_i} \nu_i < \frac{1}{\max_{i \in \mathcal{J}} |\mathcal{K}_i|}$ for all $k \in \mathcal{K}$. This is equivalent to

$$0 \leq \varepsilon < \min_{k \in \mathcal{K}} \frac{\frac{1}{\max_{i \in \mathcal{J}} |\mathcal{K}_i|} - \rho^{(k)}}{m^{(k)}},$$

where we define $m^{(k)} := \sum_{i \in \mathcal{J}: k \in \mathcal{K}_i} m_i$. The above interval is non-empty as long as $\rho^{(k)} < \frac{1}{\max_{i \in \mathcal{J}} |\mathcal{K}_i|}$ for all $k \in \mathcal{K}$.

Stability of LRFS policies. We now state the main result for synchronized stochastic networks satisfying Assumption *B1*. Since the network is synchronized, we obtain the following proposition as a corollary of Theorem 3.5.

Proposition 5.2. *If a stochastic processing network is synchronized and satisfies Assumption B1 with $\rho^{(k)} < \frac{1}{\max_{i \in \mathcal{J}} |\mathcal{K}_i|}$ for all $k \in \mathcal{K}$, then*

- *The ε -LRFS process is queue-length-stable for any $\varepsilon \in \left(0, \min_{k \in \mathcal{K}} \frac{\frac{1}{\max_{i \in \mathcal{J}} |\mathcal{K}_i|} - \rho^{(k)}}{m^{(k)}}\right)$.*
- *The LRFS process is queue-length-stable if all routes are bounded in length.*

Proposition 5.2 implies that the ε -LRFS policy can achieve a $\frac{1}{\max_{i \in \mathcal{J}} |\mathcal{K}_i|}$ fraction of the capacity region (17). For networks of input-queued switches [18, 6] and wireless networks with primary interference constraints [21], it is easy to see that $\max_{i \in \mathcal{J}} |\mathcal{K}_i| = 2$, and hence the ε -LRFS policy achieves 50% of the capacity region. For wireless networks with general independent set constraints [22], $\max_{i \in \mathcal{J}} |\mathcal{K}_i|$ is the maximum number of interfering neighbors, i.e., the maximum degree of the underlying interference graph.

In large-scale networks, distributed scheduling schemes with low complexity have gained much attention recently, even though they usually perform worse than centralized ones with high complexity. We remark that the 50% throughput result of a greedy scheduling algorithm has previously been established for input-queued switches (i.e., no routing between buffers) by Dai et al. [6]. Proposition 5.2 generalizes this to ‘networks’ of input-queued switches operated under stochastic routing between buffers (or local switches). In wireless networks with primary interference constraints, Wu et al. [25] establish the 50% throughput result of the LRFS policy assuming deterministic routing (i.e., fixed routes between nodes), while Proposition 5.2 allows stochastic routing.

6 Proof of Theorem 3.5

To prove the desired stability of the ε -LRFS process $\{X(t)\}$, we construct a new process $\{Y(t)\}$, which is almost identical to $\{X(t)\}$, but it has a larger state space. The main idea of the proof is to construct a Lyapunov function for the ‘larger’ process $\{Y(t)\}$, which implies the stability of $\{Y(t)\}$ and therefore the stability of $\{X(t)\}$.

The description of $\{Y(t)\}$ is as follows. Consider the stochastic processing network setup in Section 2, and let $\mathcal{P} = \{P_1, P_2, \dots\}$ be the collection of all possible paths of buffers (allowing repetitions) of length at most D (i.e., $|\mathcal{P}| \leq (I+1)^D$), where $D \in \mathbb{N}$ is some finite constant to be determined later. We assume that when a job enters the network, it pre-determines the first D buffers on its route. After being processed from these D buffers, jobs perform the usual stochastic routing as described in Section 2. Let $i_{nm} \in \mathcal{J}$, $Q_{nm}(t) \in \mathbb{Z}_+$ and $V_{nm}^j(t) \in \mathbb{R}_+$ denote

the m -th buffer on path P_n , the number of jobs waiting in buffer i_{nm} and the remaining service requirement of the job in buffer i_{nm} being processed by activity j at time t , respectively. If activity j is not processing a job in buffer i_{nm} , then $V_{nm}^j(t) = 0$. Furthermore, as before, let $Q_{i,c}(t)$ be the queue length (i.e., the number of jobs waiting for service, excluding those being processed) with counter c in buffer i at time t . $V_{i,c}^j(t)$ is defined to be the remaining service requirement of the job with counter c in buffer i being processed by activity j at time t . We then define

$$Y(t) = \left[Q_{nm}(t), Q_{i,c}(t), V_{nm}^j(t), V_{i,c}^j(t), \mathcal{J}^{(h)}(t) \right] \in \Omega_Y := \mathbb{Z}_+^\infty \times \mathbb{R}_+^\infty \times [0, 1]^H,$$

where n, m, c, h are positive integers such that $n \leq (I+1)^D$, $m \leq D$, $c > D$ and $h \leq H$.

As for $\{X(t)\}$, we impose the convention that $\{Y(t)\}$ has right-continuous sample paths. We define the norm $|Y(t)|$ through

$$|Y(t)| = \sum_{n \leq (I+1)^D, m \leq D} Q_{nm}(t) + \sum_{i \in \mathcal{J}, c > D} Q_{i,c}(t) + \sum_{n \leq (I+1)^D, m \leq D, j \in \mathcal{J}} V_{nm}^j(t) + \sum_{i \in \mathcal{J}, c > D, j \in \mathcal{J}} V_{i,c}^j(t).$$

One can define a natural projection P such that the distribution of $P(Y(t))$ is identical to that of $X(t)$ and $|P(Y(t))| = |X(t)|$ given that $Y(0)$ is drawn appropriately from the preimage $P^{-1}(X(0))$ of $X(0)$. Intuitively speaking, $\{Y(t)\}$ tosses random coins to determine routes in advance, and since the scheduling decisions of ε -LRFS are independent of these coins, the natural projection of $\{Y(t)\}$ ignoring these pre-determined coin flips provides exactly the dynamics of $\{X(t)\}$. Hence, it suffices to prove that the (bigger) process $\{Y(t) : t \in \mathbb{R}_+\}$ is queue-length-stable, i.e.,

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \int_0^t E[|Y(s)|] ds < \infty, \quad (18)$$

for any given initial state $Y(0) \in \Omega_Y^* := \{Y \in \Omega_Y : |Y| < \infty\}$. In essence, this follows from the following proposition, which is proved in Section 6.1.

Proposition 6.1. *If the conditions of Theorem 3.5 hold, then there exist constants $N, T, D, \zeta, \mathcal{C} \in (0, \infty)$ and a Lyapunov function $\mathcal{L}_{\text{global}} : \Omega_Y \rightarrow \mathbb{R}_+ \cup \{\infty\}$ such that for all $t \geq N$,*

$$E[\mathcal{L}_{\text{global}}(Y(t+T)) \mid \mathcal{F}(t)] \leq \mathcal{L}_{\text{global}}(Y(t)) - \zeta|Y(t)| + \mathcal{C}, \quad \text{almost surely,}$$

and $\sup_{Y \in \Omega_Y^*} \mathcal{L}_{\text{global}}(Y)/|Y|^2 < \infty$, where we define the filtration $\{\mathcal{F}(t) : t \geq 0\}$ by

$$\mathcal{F}(t) := \sigma\{Y(s) : 0 \leq s \leq t\}.$$

Now we describe how Proposition 6.1 implies (18), and hence the conclusion of Theorem 3.5. First one can observe that $E[|Y(t)|^2] < \infty$ since $Y(0) \in \Omega_Y^*$ and we assume bounded second moments on arrivals and service times. Since $\sup_{Y \in \Omega_Y^*} \mathcal{L}_{\text{global}}(Y)/|Y|^2 < \infty$ from Proposition 6.1, it follows that

$$E[\mathcal{L}_{\text{global}}(Y(t+T))], E[\mathcal{L}_{\text{global}}(Y(t))], E[|Y(t)|] < \infty. \quad (19)$$

Combining Proposition 6.1 and (19) yields that for all $t \geq N$,

$$E[\mathcal{L}_{\text{global}}(Y(t+T))] \leq E[\mathcal{L}_{\text{global}}(Y(t))] - \zeta E[|Y(t)|] + \mathcal{C}.$$

Therefore, we have that for $t \geq N$,

$$\begin{aligned} \int_{T+N}^{T+t} E[\mathcal{L}_{\text{global}}(Y(s))] ds &= \int_N^t E[\mathcal{L}_{\text{global}}(Y(s+T))] ds \\ &\leq \int_N^t E[\mathcal{L}_{\text{global}}(Y(s))] ds - \zeta \int_N^t E[|Y(s)|] ds + \mathcal{C}(t-N), \end{aligned}$$

which implies that for $t \geq T+N$,

$$\frac{\zeta}{t-N} \int_N^t E[|Y(s)|] ds \leq \mathcal{C} + \frac{1}{t-N} \int_N^{T+N} E[\mathcal{L}_{\text{global}}(Y(s))] ds.$$

The right-hand side of the above inequality converges to \mathcal{C} as $t \rightarrow \infty$. This leads to the desired conclusion (18).

6.1 Proof of Proposition 6.1

The choice of N in Proposition 6.1 comes from our assumption on the external arrival processes in Section 2, namely, that there exists some $N < \infty$ such that for all $i \in \mathcal{J}$ and $t \geq N$,

$$E\left[A_i(t, t+1) \mid \{A_1(0, s), \dots, A_I(0, s) : 0 \leq s \leq t\}\right] \leq \alpha_i.$$

For notational convenience, we assume $N = 0$ in the proof of Proposition 6.1. Namely, we assume that for all $t \geq 0$,

$$E\left[A_i(t, t+1) \mid \mathcal{F}(t)\right] \leq \alpha_i < \infty. \quad (20)$$

All the proof arguments are applicable to the general case $N > 0$.

We first define some further notation in Section 6.1.1. The skeleton of the proof of Proposition 6.1 is described in Section 6.1.2, and it uses three key lemmas. The proofs of these lemmas follow.

6.1.1 Notation

The quantities defined below are simple functions of the network state $Y(t)$. We use bold symbols to denote vectors of quantities, e.g., $\mathbf{Q}_c(t) = [Q_{i,c}(t)]$ and $\mathbf{W}_{\leq c}(t) = [W_{i,\leq c}(t)]$.

Queues. For $c \in \mathbb{N}$, let $Q_{i,c}(t)$, $Q_{i,<c}(t)$ and $Q_{i,\leq c}(t)$ be the number of waiting jobs with counter c , $< c$ and $\leq c$ in buffer i at time t , respectively. That is, we set

$$\begin{aligned} Q_{i,c}(t) &= \sum_{n=1}^{(I+1)^D} Q_{nc}(t) I_{\{i\}}(i_{nc}) \quad \text{for } c \leq D, \\ Q_{i,<c}(t) &= \sum_{c'=1}^{c-1} Q_{i,c'}(t) \quad Q_{i,\leq c}(t) = \sum_{c'=1}^c Q_{i,c'}(t), \end{aligned}$$

where we recall that $I_{\{i\}}(i_{nc}) = 1$ if $i_{nc} = i$ and $I_{\{i\}}(i_{nc}) = 0$ otherwise. Furthermore, for $c \leq D$, we let $\widehat{Q}_{i,\leq c}(t)$ be the number of waiting jobs with counter $\leq c$ ‘in or destined for’ buffer i at time t . Namely,

$$\widehat{Q}_{i,\leq c}(t) = Q_{i,\leq c}(t) + \sum_{n=1}^{(I+1)^D} \sum_{m=1}^c \sum_{r=m+1}^c Q_{nm}(t) I_{\{i\}}(i_{nr}),$$

where the first term ‘ $Q_{i,\leq c}(t)$ ’ and the second term ‘ $\sum_{n=1}^{(I+1)^D} \dots$ ’ on the right-hand side count the numbers of waiting jobs currently in buffer i and destined for buffer i , respectively.

Workloads. Let $V^j(t)$ denote the remaining service requirement of the job being processed by activity j at time t , and $V_i(t)$ the total remaining service requirement of the jobs being processed in buffer i at time t (multiple jobs can be processed from the same buffer by different processors). Similarly, $V_{i,c}(t), V_{i,<c}(t), V_{i,\leq c}(t)$ stands for the total remaining service requirement of jobs with counter $c, < c, \leq c$ being processed in buffer i at time t , respectively. We furthermore define the following quantities:

if the network is not synchronized,

$$\begin{aligned} W_{i,<c}(t) &= m_i Q_{i,<c}(t) + V_i(t) \\ W_{i,\leq c}(t) &= m_i Q_{i,\leq c}(t) + V_i(t) \\ \widehat{W}_{i,\leq c}(t) &= m_i \widehat{Q}_{i,\leq c}(t) + V_i(t). \end{aligned}$$

if the network is synchronized,

$$\begin{aligned} W_{i,<c}(t) &= m_i Q_{i,<c}(t) + V_{i,<c}(t) \\ W_{i,\leq c}(t) &= m_i Q_{i,\leq c}(t) + V_{i,\leq c}(t) \\ \widehat{W}_{i,\leq c}(t) &= m_i \widehat{Q}_{i,\leq c}(t) + V_{i,\leq c}(t). \end{aligned}$$

We use different definitions for these variables depending on whether the network is synchronized or not, since the strategy of the proof differs in each case, e.g., see Section 6.1.4. We also note that in a synchronized network, $V_{i,c}(t), V_{i,<c}(t), V_{i,\leq c}(t) \in \{0, 1\}$ for $t \in \mathbb{Z}_+$. The quantities $W_{i,<c}(t)$ and $W_{i,\leq c}(t)$ are (expected) immediate workloads, since they only involve work that is in buffer i at time t . The quantities $\widehat{W}_{i,\leq c}(t)$ are (expected) total workloads, since they incorporate work currently in the system which will be routed to buffer i , regardless where the work resides in the network at time t .

Three types of jobs & weights. We distinguish three types of jobs:

- Type 1. Jobs with counter $> D$
- Type 2. Jobs on a path of length $= D$ in \mathcal{P} .
- Type 3. Jobs on a path of length $< D$ in \mathcal{P} .

Hence, the counters of jobs of Type 2 and 3 cannot exceed D . We further note that each job (regardless whether it is currently being processed or not) can compute its expected total remaining service requirement in the future (under the network process $\{Y(t)\}$), and we call this quantity the weight of the job. Namely, if a job is not currently being processed, its weight is

$$\sum_{i \in \mathcal{J}} m_i E[\text{the number of times it visits buffer } i \text{ before it leaves the network}],$$

which includes the current buffer of the job. Thus, the weight of a waiting job, not currently being processed, can be calculated as follows.

Type 1. The weight of a waiting job of Type 1 in buffer $i \in \mathcal{J}$ is

$$\mathbf{e}_i^T (I + P + P^2 + \dots) \mathbf{m} = \mathbf{e}_i^T (I - P)^{-1} \mathbf{m},$$

where $\mathbf{m} = [m_i] \in \mathbb{R}_+^I$ and $\mathbf{e}_i \in \{0, 1\}^I$ is the unit vector with zeros except for its i -th coordinate which equals to one (both are column vectors).

Type 2. The weight of a waiting job of Type 2 in buffer i_{nm} on a path P_n of length D (i.e., $|P_n| = D$) is

$$\mathbf{e}_{i_{nD}}^T (P + P^2 + P^3 + \dots) \mathbf{m} + \sum_{k=m}^D m_{i_{nk}} = \mathbf{e}_{i_{nD}}^T P (I - P)^{-1} \mathbf{m} + \sum_{k=m}^D m_{i_{nk}},$$

where we note that i_{nD} is the last buffer on P_n .

Type 3. The weight of a waiting job of Type 3 in buffer i_{nm} on a path P_n of length $< D$ is

$$\sum_{k=m}^{|P_n|} m_{i_{nk}}.$$

On the other hand, for jobs currently being processed, each weight is calculated as

the current remaining service requirement

$$+ \sum_{i \in \mathcal{J}} m_i E[\text{the number of times it visits buffer } i \text{ before it leaves the network}],$$

where the latter number does not include the current buffer of the job. Now let $M_1(t)$, $M_2(t)$ and $M_3(t)$ be the total weights of jobs of Type 1, 2 and 3 at time t , respectively. These total weights are fully determined by the network state information $Y(t)$:

$$\begin{aligned} M_1(t) &= \sum_{i \in \mathcal{J}, c > D} Q_{i,c}(t) \mathbf{e}_i^T (I - P)^{-1} \mathbf{m} + \sum_{i \in \mathcal{J}, c > D, j \in \mathcal{J}} \mathbf{1}_{V_{i,c}^j}^+ \left[V_{i,c}^j(t) + \mathbf{e}_i^T P (I - P)^{-1} \mathbf{m} \right], \\ M_2(t) &= \sum_{n \leq (I+1)^D, m \leq D, |P_n|=D} Q_{nm}(t) \left[\mathbf{e}_{i_{nD}}^T P (I - P)^{-1} \mathbf{m} + \sum_{k=m}^D m_{i_{nk}} \right] \\ &\quad + \sum_{n \leq (I+1)^D, m \leq D, j \in \mathcal{J}, |P_n|=D} \mathbf{1}_{V_{nm}^j}^+ \left[V_{nm}^j(t) + \mathbf{e}_{i_{nD}}^T P (I - P)^{-1} \mathbf{m} + \sum_{k=m+1}^D m_{i_{nk}} \right], \\ M_3(t) &= \sum_{n \leq (I+1)^D, m \leq D, |P_n| < D} \left[Q_{nm}(t) \sum_{k=m}^{|P_n|} m_{i_{nk}} \right] \\ &\quad + \sum_{n \leq (I+1)^D, m \leq D, j \in \mathcal{J}, |P_n| < D} \mathbf{1}_{V_{nm}^j}^+ \left[V_{nm}^j(t) + \sum_{k=m+1}^{|P_n|} m_{i_{nk}} \right], \end{aligned}$$

where we recall that $\mathbf{1}_x^+ = 1$ if $x > 0$, and $\mathbf{1}_x^+ = 0$ otherwise.

6.1.2 Three Auxiliary Lemmas and Proof of Proposition 6.1

We state and prove the following three key lemmas, which we prove in Section 6.1.3, 6.1.4 and 6.1.5, respectively. To simplify notation, we will use \mathcal{C} to denote a finite constant which only depends on the matrix Z given in Theorem 3.5 or on the predefined network parameters from Section 2. Its precise value can be different from line to line.

Lemma 6.2. *There exists a constant $\mathcal{C} < \infty$ such that for all $t \geq 0$,*

$$E [\|\mathbf{V}(t+1)\|_2^2 \mid \mathcal{F}(t)] \leq \|\mathbf{V}(t)\|_2^2 - \beta_{\min} \|\mathbf{V}(t)\|_1 + \mathcal{C},$$

where $\beta_{\min} = \min_{j \in \mathcal{J}} \beta_j$.

Lemma 6.3. *Suppose that there exists a symmetric matrix $Z \in \mathbb{R}_+^{I \times I}$ such that $\mathcal{L}_{\text{local}}(x) = x^T Z x$ is a local Lyapunov function with slack $\varepsilon \geq 0$, and that either Condition C1 or C2' from Theorem 3.5 holds. Then given $D \in \mathbb{N}$, there exist constants $v, \mathcal{C} = \mathcal{C}(D) \in (0, \infty)$ such that for all $c \leq D$, $t \geq 0$,*

$$\begin{aligned} E \left[\mathcal{L}_{\text{local}} \left(\widehat{\mathbf{W}}_{\leq c}(t+1) \right) \mid \mathcal{F}(t) \right] \\ \leq \mathcal{L}_{\text{local}} \left(\widehat{\mathbf{W}}_{\leq c}(t) \right) - v \|\mathbf{Q}_{\leq c}(t)\|_1 + \mathcal{C} (\|\mathbf{V}(t)\|_1 + \|\mathbf{Q}_{< c}(t)\|_1 + 1). \end{aligned}$$

Lemma 6.4. *Consider $\varepsilon > 0$. If the network is synchronized or if Condition C2 from Theorem 3.5 holds, then there exist constants $D, T \in \mathbb{N}$ and $\mathcal{C}, \gamma_1, \gamma_2 \in (0, \infty)$ such that for all $t \geq 0$,*

$$E [\mathcal{G}(Y(t+T))^2 \mid \mathcal{F}(t)] \leq \mathcal{G}(Y(t))^2 - \gamma_2 \|\mathbf{Q}_{> D}(t)\|_1 + \mathcal{C} (\|\mathbf{Q}_{\leq D}(t)\|_1 + \|\mathbf{V}(t)\|_1 + 1),$$

where $\mathcal{G}(Y(t)) := M_1(t) + M_2(t) + \gamma_1 \|\mathbf{V}(t)\|_1$.

Lemma 6.4 is not needed for the proof of Proposition 6.1 if all routes are bounded. Hence, for networks with bounded routes, $\varepsilon = 0$ is allowed and only Condition C2' is needed for the desired stability. The right-hand sides of the inequalities in Lemmas 6.2 – 6.4 provide ‘negative drifts’ on $\|\mathbf{V}(t)\|_1$, $\|\mathbf{Q}_{\leq D}(t)\|_1$ and $\|\mathbf{Q}_{> D}(t)\|_1$, respectively. By appropriately weighing the functions in these three lemmas, we shall construct an appropriate (global) Lyapunov function $\mathcal{L}_{\text{global}}$. To this end, from Lemma 6.2, we obtain

$$\begin{aligned} E [\|\mathbf{V}(t+T)\|_2^2 \mid \mathcal{F}(t)] &\leq E [\|\mathbf{V}(t+T-1)\|_2^2 \mid \mathcal{F}(t)] - \beta_{\min} E [\|\mathbf{V}(t+T-1)\|_1 \mid \mathcal{F}(t)] + \mathcal{C}, \\ &\leq E [\|\mathbf{V}(t+T-1)\|_2^2 \mid \mathcal{F}(t)] + \mathcal{C} \\ &\leq E [\|\mathbf{V}(t+T-2)\|_2^2 \mid \mathcal{F}(t)] + 2\mathcal{C} \\ &\quad \dots \\ &\leq E [\|\mathbf{V}(t+1)\|_2^2 \mid \mathcal{F}(t)] + (T-1)\mathcal{C} \\ &\leq \|\mathbf{V}(t)\|_2^2 - \beta_{\min} \|\mathbf{V}(t)\|_1 + T\mathcal{C}, \end{aligned} \tag{21}$$

where T is the constant from Lemma 6.4. We write this inequality as

$$E \left[\frac{1}{\beta_{\min}} \|\mathbf{V}(t+T)\|_2^2 \mid \mathcal{F}(t) \right] \leq \frac{1}{\beta_{\min}} \|\mathbf{V}(t)\|_2^2 - \|\mathbf{V}(t)\|_1 + \mathcal{C}, \tag{22}$$

where the constant \mathcal{C} is redefined appropriately. We now argue that, similarly, Lemma 6.3 implies that

$$\begin{aligned} E \left[\mathcal{L}_{\text{local}} \left(\widehat{\mathbf{W}}_{\leq c}(t+T) \right) \mid \mathcal{F}(t) \right] \\ \leq \mathcal{L}_{\text{local}} \left(\widehat{\mathbf{W}}_{\leq c}(t) \right) - v \|\mathbf{Q}_{\leq c}(t)\|_1 + \mathcal{C} (\|\mathbf{V}(t)\|_1 + \|\mathbf{Q}_{< c}(t)\|_1 + 1), \end{aligned} \quad (23)$$

where \mathcal{C} is some (large) constant which may differ from the one in Lemma 6.3. To see why this holds, we use the same argument that led to (21) with the additional observation that, for some constant \mathcal{C}' ,

$$\sup_{s \in \{0,1,\dots,T\}} (E [\|\mathbf{V}(t+s)\|_1 \mid \mathcal{F}(t)] - \|\mathbf{V}(t)\|_1) \leq \mathcal{C}', \quad (24)$$

$$\sup_{s \in \{0,1,\dots,T\}} (E [\|\mathbf{Q}(t+s)\|_1 \mid \mathcal{F}(t)] - \|\mathbf{Q}(t)\|_1) \leq \mathcal{C}'. \quad (25)$$

Here (24) can be derived along the lines of (21), and (25) follows from (20) and the observation that the change in queue length is majorized by the number of external job arrivals.

We now show that Proposition 6.1 follows from these three lemmas, where the last lemma is not needed if all routes are bounded. We consider the following Lyapunov function $\mathcal{L}_{\text{global}}$:

- If all routes are bounded,

$$\mathcal{L}_{\text{global}}(Y(t)) := \sum_{c=1}^D \left(\frac{v}{2\mathcal{C}} \right)^c \mathcal{L}_{\text{local}} \left(\widehat{\mathbf{W}}_{\leq c}(t) \right) + \frac{2\mathcal{C}}{\beta_{\min}} \|\mathbf{V}(t)\|_2^2,$$

where we choose $D < \infty$ such that $P^D = 0$, and v comes from Lemma 6.3.

- Otherwise,

$$\mathcal{L}_{\text{global}}(Y(t)) := \sum_{c=1}^D \left(\frac{v}{2\mathcal{C}} \right)^c \mathcal{L}_{\text{local}} \left(\widehat{\mathbf{W}}_{\leq c}(t) \right) + \frac{2\mathcal{C}}{\beta_{\min}} \|\mathbf{V}(t)\|_2^2 + \frac{\xi}{2\mathcal{C}} \mathcal{G}(Y(t))^2,$$

where v and D is from Lemma 6.3 and 6.4, respectively, $\xi := \xi(v, \mathcal{C}, D) = v \left(\frac{v}{2\mathcal{C}} \right)^D$ and \mathcal{C} is a large enough constant chosen so that it can be used for Lemma 6.4 as well as for (22) and (23).

We focus on proving Proposition 6.1 for the case of unbounded routes, but all arguments go through for the other case. Without loss of generality, we assume that

$$\xi, v, \gamma_2 < 1 < \mathcal{C}.$$

The property $\sup_{Y \in \Omega_Y^*} \mathcal{L}_{\text{global}}(Y)/|Y|^2 < \infty$ in Proposition 6.1 is readily seen to hold. To derive the negative drift property, we observe that Lemma 6.4 in conjunction with (22) and (23) imply

that

$$\begin{aligned}
& E [\mathcal{L}_{\text{global}}(Y(t+T)) - \mathcal{L}_{\text{global}}(Y(t)) \mid \mathcal{F}(t)] \\
& \leq \mathcal{C}^* - 2\mathcal{C}\|\mathbf{V}(t)\|_1 + \sum_{c=1}^D \left(\frac{v}{2\mathcal{C}}\right)^c \left[-v\|\mathbf{Q}_{\leq c}(t)\|_1 + \mathcal{C}\|\mathbf{Q}_{< c}(t)\|_1 + \mathcal{C}\|\mathbf{V}(t)\|_1 \right] \\
& \quad - \frac{\gamma_2\xi}{2\mathcal{C}}\|\mathbf{Q}_{> D}(t)\|_1 + \frac{\xi}{2} (\|\mathbf{Q}_{\leq D}(t)\|_1 + \|\mathbf{V}(t)\|_1 + 1) \\
& \leq \mathcal{C}^* - \mathcal{C}\|\mathbf{V}(t)\|_1 + \sum_{c=1}^D \left(\frac{v}{2\mathcal{C}}\right)^c \left[-v\|\mathbf{Q}_{\leq c}(t)\|_1 + \mathcal{C}\|\mathbf{Q}_{< c}(t)\|_1 \right] \\
& \quad - \frac{\gamma_2\xi}{2\mathcal{C}}\|\mathbf{Q}_{> D}(t)\|_1 + \frac{\xi}{2} (\|\mathbf{Q}_{\leq D}(t)\|_1 + \|\mathbf{V}(t)\|_1 + 1),
\end{aligned}$$

where \mathcal{C}^* is some (large enough) constant and we use $v \leq \mathcal{C}$ for the last inequality. The sum in this expression can be bounded as follows:

$$\begin{aligned}
& \sum_{c=1}^D \left(\frac{v}{2\mathcal{C}}\right)^c \left[-v\|\mathbf{Q}_{\leq c}(t)\|_1 + \mathcal{C}\|\mathbf{Q}_{< c}(t)\|_1 \right] \\
& = -v \sum_{c=1}^D \left(\frac{v}{2\mathcal{C}}\right)^c \|\mathbf{Q}_{\leq c}(t)\|_1 + \frac{v}{2} \sum_{c=1}^D \left(\frac{v}{2\mathcal{C}}\right)^{c-1} \|\mathbf{Q}_{< c}(t)\|_1 \\
& = -v \left(\frac{v}{2\mathcal{C}}\right)^D \|\mathbf{Q}_{\leq D}(t)\|_1 - \left(v - \frac{v}{2}\right) \sum_{c=1}^{D-1} \left(\frac{v}{2\mathcal{C}}\right)^c \|\mathbf{Q}_{\leq c}(t)\|_1 \\
& \leq -v \left(\frac{v}{2\mathcal{C}}\right)^D \|\mathbf{Q}_{\leq D}(t)\|_1 = -\xi\|\mathbf{Q}_{\leq D}(t)\|_1.
\end{aligned}$$

After combining the preceding two displays, we obtain the desired negative drift property:

$$\begin{aligned}
& E [\mathcal{L}_{\text{global}}(Y(t+T)) - \mathcal{L}_{\text{global}}(Y(t)) \mid \mathcal{F}(t)] \\
& \leq \mathcal{C}^* - \mathcal{C}\|\mathbf{V}(t)\|_1 - \xi\|\mathbf{Q}_{\leq D}(t)\|_1 - \frac{\gamma_2\xi}{2\mathcal{C}}\|\mathbf{Q}_{> D}(t)\|_1 + \frac{\xi}{2} (\|\mathbf{Q}_{\leq D}(t)\|_1 + \|\mathbf{V}(t)\|_1 + 1) \\
& \leq \mathcal{C}^* + \frac{\xi}{2} - \frac{\gamma_2\xi}{2\mathcal{C}} (\|\mathbf{V}(t)\|_1 + \|\mathbf{Q}_{\leq D}(t)\|_1 + \|\mathbf{Q}_{> D}(t)\|_1) \\
& = \mathcal{C}^* + \frac{\xi}{2} - \frac{\gamma_2\xi}{2\mathcal{C}} (\|\mathbf{V}(t)\|_1 + \|\mathbf{Q}(t)\|_1),
\end{aligned}$$

where we use $\xi, v, \gamma_2 < 1$ and $\mathcal{C} > 1$. This completes the proof of Proposition 6.1.

6.1.3 Proof of Lemma 6.2

Recall that $\|\mathbf{V}(t)\|_2^2$ is the sum of squares of the remaining service requirements of jobs being processed by some activity at time t , i.e.,

$$\|\mathbf{V}(t)\|_2^2 = \sum_{j \in \mathcal{J}} V^j(t)^2 \quad \text{and} \quad \|\mathbf{V}(t)\|_1 = \sum_{j \in \mathcal{J}} V^j(t).$$

On the event $\{V^j(t) \leq \beta_j\}$, activity j has to restart before time $t + 1$ and hence,

$$\begin{aligned}
& E [V^j(t+1)^2 \mid \mathcal{F}(t)] \\
& \leq E [V^j(t+1)^2 \mid V^j(t+1) > 0, \mathcal{F}(t)] \\
& \leq E [\Gamma_{i_j, t+1}^2 \mid V^j(t+1) > 0, \mathcal{F}(t)] \\
& = E [E [\Gamma_{i_j, t+1}^2 \mid V^j(t+1) > 0, \mathcal{F}(t), J^{(t+1)}] \mid V^j(t+1) > 0, \mathcal{F}(t)],
\end{aligned}$$

where we let $\Gamma_{i_j, t+1}$ denote the service time generated by the job being processed by activity j at time $t + 1$ and we write

$$J^{(t+1)} := t + 1 - \frac{\Gamma_{i_j, t+1} - V^j(t+1)}{\beta_j} \in [t, t+1]$$

for the time when this job starts its service. Note that, again on the event $\{V^j(t) \leq \beta_j\}$,

$$\begin{aligned}
& E [\Gamma_{i_j, t+1}^2 \mid V^j(t+1) > 0, \mathcal{F}(t), J^{(t+1)}] \\
& = \int_0^\infty \Pr [\Gamma_{i_j, t+1}^2 > x \mid V^j(t+1) > 0, \mathcal{F}(t), J^{(t+1)}] dx \\
& = \int_0^\infty \Pr [\Gamma_{i_j}^2 > x \mid \Gamma_{i_j} > \beta_j (t+1 - J^{(t+1)})] dx \\
& = \int_0^\infty \frac{\Pr [\Gamma_{i_j}^2 > x, \Gamma_{i_j} > \beta_j (t+1 - J^{(t+1)})]}{\Pr [\Gamma_{i_j} > \beta_j (t+1 - J^{(t+1)})]} dx \\
& \leq \int_0^\infty \frac{\Pr [\Gamma_{i_j}^2 > x]}{\Pr [\Gamma_{i_j} > \beta_j]} dx \\
& = \frac{E[\Gamma_{i_j}^2]}{\Pr[\Gamma_{i_j} > \beta_j]},
\end{aligned}$$

where we recall that Γ_{i_j} stands for a generic service time for buffer i_j . We have thus established that on $\{V^j(t) \leq \beta_j\}$,

$$E [V^j(t+1)^2 \mid \mathcal{F}(t)] \leq \begin{cases} \frac{E[\Gamma_{i_j}^2]}{\Pr[\Gamma_{i_j} > \beta_j]} & \text{if } \Pr[\Gamma_{i_j} > \beta_j] > 0 \\ (\beta_j)^2 & \text{otherwise} \end{cases}. \quad (26)$$

On the other hand, on the event $\{V^j(t) > \beta_j\}$,

$$V^j(t+1)^2 = (V^j(t) - \beta_j)^2 \leq V^j(t)^2 - \beta_j V^j(t) + \beta_j^2. \quad (27)$$

Hence, combining (26) and (27), we find that for some constant $\mathcal{C} < \infty$,

$$E[V^j(t+1)^2 \mid \mathcal{F}(t)] \leq V^j(t)^2 - \beta_j V^j(t) + \mathcal{C},$$

which leads to the desired conclusion of Lemma 6.2.

6.1.4 Proof of Lemma 6.3

For notational convenience, we stick to the case $t = 0$ in the conclusion of Lemma 6.3. Namely, we show that

$$E \left[\mathcal{L}_{\text{local}} \left(\widehat{\mathbf{W}}_{\leq c}(1) \right) \mid \mathcal{F}(0) \right] \leq \mathcal{L}_{\text{local}} \left(\widehat{\mathbf{W}}_{\leq c}(0) \right) - v \|\mathbf{Q}_{\leq c}(0)\|_1 + \mathcal{C} (\|\mathbf{V}(0)\|_1 + \|\mathbf{Q}_{< c}(0)\|_1 + 1).$$

However, all arguments go through for general $t > 0$.

Non-synchronized network. We first consider the case when the network may not be synchronized. The first step in the proof is the observation that the schedule $\boldsymbol{\sigma}(t)$ under the ε -LRFS policy is always maximal with respect to the vector $\mathbf{w} = [w_i] = [m_i Q_{i, \leq c}(t)]$ for any $c \geq 1$. Consequently, since the local quadratic Lyapunov function $\mathcal{L}_{\text{local}}(x) = x^T Z x$ satisfies (3), we obtain that for $t \in [0, 1]$,

$$2 \sum_{(i, \ell) \in \mathcal{J} \times \mathcal{J}} Z_{i\ell} w_\ell \left(\rho_i + \varepsilon m_i - \sum_{j \in \mathcal{J}_i} \beta_j \sigma_j(t) \right) \leq -\eta \|\mathbf{w}\|_1 + \mathcal{C},$$

where we further use the observation that

$$\begin{aligned} & \left(w_i + \rho_i + \varepsilon m_i - \sum_{j \in \mathcal{J}_i} \beta_j \sigma_j(t) \right) \left(w_\ell + \rho_\ell + \varepsilon m_\ell - \sum_{j \in \mathcal{J}_\ell} \beta_j \sigma_j(t) \right) - w_i w_\ell \\ & \leq w_\ell \left(\rho_i + \varepsilon m_i - \sum_{j \in \mathcal{J}_i} \beta_j \sigma_j(t) \right) + w_i \left(\rho_\ell + \varepsilon m_\ell - \sum_{j \in \mathcal{J}_\ell} \beta_j \sigma_j(t) \right) + \mathcal{C}. \end{aligned} \quad (28)$$

We remind the reader that we write \mathcal{C} for a finite constant which may differ from line to line. Since we set $w_i = m_i Q_{i, \leq c}(t)$, it follows that

$$2 \sum_{(i, \ell) \in \mathcal{J} \times \mathcal{J}} Z_{i\ell} m_\ell Q_{\ell, \leq c}(t) \left(\rho_i + \varepsilon m_i - \sum_{j \in \mathcal{J}_i} \beta_j \sigma_j(t) \right) \leq -\eta \sum_{i \in I} m_i Q_{i, \leq c}(t) + \mathcal{C}.$$

After taking conditional expectations given $\mathcal{F}(0)$ on both sides in the above inequality, we obtain

$$2 \sum_{(i, \ell) \in \mathcal{J} \times \mathcal{J}} Z_{i\ell} m_\ell \cdot E \left[Q_{\ell, \leq c}(t) \left(\rho_i + \varepsilon m_i - \sum_{j \in \mathcal{J}_i} \beta_j \sigma_j(t) \right) \mid \mathcal{F}(0) \right] \leq -\eta \sum_{i \in I} m_i Q_{i, \leq c}(0) + \mathcal{C},$$

where we use that $E[Q_{\ell, \leq c}(t) \mid \mathcal{F}(0)] \geq Q_{\ell, \leq c}(0) - \mathcal{C}'$ for $t \in [0, 1]$ for some constant \mathcal{C}' . This can be verified by suppressing any arrivals and letting all activities work, and then using the standard fact from renewal theory that any renewal function is finite. Similarly, we obtain from $E \left[Q_{\ell, \leq c}(t) \mid \sum_{j \in \mathcal{J}_i} \beta_j \sigma_j(t), \mathcal{F}(0) \right] \geq Q_{\ell, \leq c}(0) - \mathcal{C}'$ for $t \in [0, 1]$ that

$$2 \sum_{(i, \ell) \in \mathcal{J} \times \mathcal{J}} Z_{i\ell} m_\ell Q_{\ell, \leq c}(0) \cdot E \left[\rho_i + \varepsilon m_i - \sum_{j \in \mathcal{J}_i} \beta_j \sigma_j(t) \mid \mathcal{F}(0) \right] \leq -\eta \sum_{i \in I} m_i Q_{i, \leq c}(0) + \mathcal{C}, \quad (29)$$

where the constant \mathcal{C} again has to be redefined appropriately. We leave this inequality for later use.

The second step in the proof is to bound $E \left[\widehat{W}_{i,\leq c}(1) - \widehat{W}_{i,\leq c}(0) \mid \mathcal{F}(0) \right]$ for fixed i . Let $\widehat{A}_{i,\leq c}$ be the number of job arrivals contributing to an increase in $\widehat{Q}_{i,\leq c}(\cdot)$ during the time interval $[0, 1)$. Then, one can check that for $\widehat{\mathbf{A}}_{\leq c} = \left[\widehat{A}_{i,\leq c} \right]$,

$$E \left[\widehat{\mathbf{A}}_{\leq c} \mid \mathcal{F}(0) \right] \leq (I + P + \dots + P^c) \boldsymbol{\alpha} \leq \boldsymbol{\lambda},$$

since we assume (20). Define the following quantities.

- $D_{i,\leq c}$ and $D_{i,>c}$ are the numbers of jobs in buffer i which start their service during the time interval $[0, 1)$ and with counter $\leq c$ and $> c$, respectively.
- $R_{i,\leq c}$ and $R_{i,>c}$ are the total amounts of service times generated by jobs contributing to $D_{i,\leq c}$ and $D_{i,>c}$, respectively, again during the time interval $[0, 1)$. We stress that the contribution of each job to these quantities may exceed the service time it receives during the interval $[0, 1)$.
- $R_{i,>c}^{(L)}$ are the total amounts of service times generated by jobs contributing to $D_{i,>c}$ due to the LRFS policy (i.e., step 4 in Definition 3.2). In addition, we set $R_{i,>c}^{(M)} = R_{i,>c} - R_{i,>c}^{(L)}$.

Then, we have that

$$\Delta_{i,\leq c} := \widehat{W}_{i,\leq c}(1) - \widehat{W}_{i,\leq c}(0) = m_i \left(\widehat{A}_{i,\leq c} - D_{i,\leq c} \right) + R_{i,\leq c} + R_{i,>c} - \sum_{j \in \mathcal{J}_i} \int_0^1 \beta_j \sigma_j(t) dt. \quad (30)$$

Taking conditional expectations given $\mathcal{F}(0)$ on both sides, we obtain

$$\begin{aligned} & E \left[\Delta_{i,\leq c} \mid \mathcal{F}(0) \right] \\ & \leq m_i \lambda_i + E \left[-m_i D_{i,\leq c} + R_{i,\leq c} \mid \mathcal{F}(0) \right] + E \left[R_{i,>c} \mid \mathcal{F}(0) \right] - E \left[\sum_{j \in \mathcal{J}_i} \int_0^1 \beta_j \sigma_j(t) dt \mid \mathcal{F}(0) \right] \\ & = \rho_i + E \left[R_{i,>c} \mid \mathcal{F}(0) \right] - E \left[\sum_{j \in \mathcal{J}_i} \int_0^1 \beta_j \sigma_j(t) dt \mid \mathcal{F}(0) \right]. \end{aligned} \quad (31)$$

Now we bound $E \left[R_{i,>c} \mid \mathcal{F}(0) \right]$ in the above inequality, or equivalently $E \left[R_{i,>c}^{(L)} \mid \mathcal{F}(0) \right] + E \left[R_{i,>c}^{(M)} \mid \mathcal{F}(0) \right]$ since $R_{i,>c} = R_{i,>c}^{(M)} + R_{i,>c}^{(L)}$. First, one can check that

$$E \left[R_{i,>c}^{(M)} \mid \mathcal{F}(0) \right] \leq \varepsilon m_i.$$

This is because the expected number of jobs in buffer i which start their service during the time interval $[0, 1)$ due to step 3-1 of the ε -LRFS policy in Definition 3.2 is at most ε since each timer $\mathcal{T}^{(h)}(t)$ is zero at most once during this time interval. On the other hand, to bound $E \left[R_{i,>c}^{(L)} \mid \mathcal{F}(0) \right]$, consider activity-interchangeable buffers ℓ and i (which includes $\ell = i$). Let \mathcal{E}_ℓ be the event that every job in the queue $Q_{\ell,\leq c}(0)$ (i.e., jobs with counter $\leq c$ waiting in buffer

ℓ at time 0) starts service during the time interval $[0, 1]$. One can observe that $R_{i,>c}^{(L)} = 0$ on the complementary event $\bar{\mathcal{E}}_\ell$. We let

$$N := \sum_{w=1}^{Q_{\ell,\leq c}(0) - |\mathcal{J}_\ell|} \Gamma_{\ell,w},$$

where $\{\Gamma_{\ell,w}\}$ are identical random variables with mean m_ℓ and variance $\varsigma_\ell^2 < \infty$. We first consider the case when

$$E[N \mid Q_{\ell,\leq c}(0)] = m_\ell(Q_{\ell,\leq c}(0) - |\mathcal{J}_\ell|) > \beta_{\max}|\mathcal{J}_\ell|. \quad (32)$$

Since \mathcal{E}_ℓ occurs only if at least $Q_{\ell,\leq c}(0) - |\mathcal{J}_\ell|$ jobs complete their service requirements, we obtain the following on the event that (32) holds:

$$\begin{aligned} \Pr[\mathcal{E}_\ell \mid \mathcal{F}(0)] &\leq \Pr[N \leq \beta_{\max}|\mathcal{J}_\ell| \mid \mathcal{F}(0)] \\ &= \Pr[N - E[N \mid Q_{\ell,\leq c}(0)] \leq \beta_{\max}|\mathcal{J}_\ell| - E[N \mid Q_{\ell,\leq c}(0)] \mid \mathcal{F}(0)] \\ &\leq \Pr\left[(N - E[N \mid Q_{\ell,\leq c}(0)])^2 \geq (E[N \mid Q_{\ell,\leq c}(0)] - \beta_{\max}|\mathcal{J}_\ell|)^2 \mid \mathcal{F}(0)\right] \\ &\leq \frac{\varsigma_\ell^2(Q_{\ell,\leq c}(0) - |\mathcal{J}_\ell|)}{(E[N \mid Q_{\ell,\leq c}(0)] - \beta_{\max}|\mathcal{J}_\ell|)^2} \\ &= \frac{\varsigma_\ell^2(Q_{\ell,\leq c}(0) - |\mathcal{J}_\ell|)}{(m_\ell(Q_{\ell,\leq c}(0) - |\mathcal{J}_\ell|) - \beta_{\max}|\mathcal{J}_\ell|)^2} \\ &\leq \frac{\mathcal{C}}{Q_{\ell,\leq c}(0) + 1}, \end{aligned} \quad (33)$$

where $\beta_{\max} = \max_{j \in \mathcal{J}} \beta_j$, \mathcal{C} is some (finite) constant depending on ς_ℓ , m_ℓ , $|\mathcal{J}_\ell|$ and we use Markov's inequality in conjunction with (32). On the event that (32) does not hold, i.e., when $Q_{\ell,\leq c}(0)$ is bounded above by $|\mathcal{J}_\ell| \left(\beta_{\max} + \frac{1}{m_\ell}\right)$, one can redefine the constant \mathcal{C} so that (33) holds. Hence, (33) always holds.

Using (33), it follows that

$$\begin{aligned} E[R_{i,>c} \mid \mathcal{F}(0)] &= E[R_{i,>c}^{(M)} \mid \mathcal{F}(0)] + E[R_{i,>c}^{(L)} \mid \mathcal{F}(0)] \\ &\leq \varepsilon m_i + \Pr[\mathcal{E}_\ell \mid \mathcal{F}(0)] \cdot E[R_{i,>c}^{(L)} \mid \mathcal{E}_\ell, \mathcal{F}(0)] \\ &\leq \varepsilon m_i + \frac{\mathcal{C}}{Q_{\ell,\leq c}(0) + 1} \cdot E[R_{i,>c}^{(L)} \mid \mathcal{E}_\ell, \mathcal{F}(0)] \\ &\leq \varepsilon m_i + \frac{\mathcal{C}}{Q_{\ell,\leq c}(0) + 1}, \end{aligned}$$

where the last inequality requires that the constant \mathcal{C} has to be redefined appropriately, since $E[R_{i,>c}^{(L)} \mid \mathcal{E}_\ell, \mathcal{F}(0)] \leq \mathcal{C}'$ as can be seen using arguments similar to those leading up to (26). Together with (31), this leads to

$$E[\Delta_{i,\leq c} \mid \mathcal{F}(0)] \leq \rho_i + \varepsilon m_i + \frac{\mathcal{C}}{Q_{\ell,\leq c}(0) + 1} - E\left[\sum_{j \in \mathcal{J}_i} \int_0^1 \beta_j \sigma_j(t) dt \mid \mathcal{F}(0)\right], \quad (34)$$

for any activity-interchangeable buffers ℓ and i .

The third step in the proof for the non-synchronized case is to prove the conclusion of Lemma 6.3. By a similar argument as in (28), the claim follows with $v = \eta \min_{i \in \mathcal{J}} m_i$ after we show that

$$2 \sum_{(i,\ell) \in \mathcal{J} \times \mathcal{J}} Z_{i\ell} \widehat{W}_{\ell, \leq c}(0) E[\Delta_{i, \leq c} \mid \mathcal{F}(0)] \leq -\eta \sum_{i \in \mathcal{J}} m_i Q_{i, \leq c}(0) + \mathcal{C} \sum_{i \in \mathcal{J}} [V_i(0) + Q_{i, < c}(0) + 1], \quad (35)$$

where we use that

$$E[\Delta_{i, \leq c} \Delta_{\ell, \leq c} \mid \mathcal{F}(0)] \leq \mathcal{C} \quad (36)$$

for all i, ℓ and some constant \mathcal{C} . To see that (36) holds, it suffices to show that $E[\Delta_{i, \leq c}^2 \mid \mathcal{F}(0)] \leq \mathcal{C}$ by the Cauchy-Schwarz inequality. This can be shown using (30) and

$$E[\widehat{A}_{i, \leq c}^2 \mid \mathcal{F}(0)], E[D_{i, \leq c}^2 \mid \mathcal{F}(0)], E[R_i^2 \mid \mathcal{F}(0)] \leq \mathcal{C}$$

where these bounds can be derived using arguments similar to those leading up to (26). Since $Z_{i\ell}, E[\Delta_{i, \leq c} \mid \mathcal{F}(0)] \leq \mathcal{C}$ and

$$\widehat{W}_{\ell, \leq c}(0) \leq W_{\ell, \leq c}(0) + \sum_{i \in \mathcal{J}} m_i Q_{i, < c}(0) = m_\ell Q_{\ell, \leq c}(0) + V_\ell(0) + \sum_{i \in \mathcal{J}} m_i Q_{i, < c}(0),$$

the inequality in (35) reduces to

$$2 \sum_{(i,\ell) \in \mathcal{J} \times \mathcal{J}} Z_{i\ell} m_\ell Q_{\ell, \leq c}(0) E[\Delta_{i, \leq c} \mid \mathcal{F}(0)] \leq -\eta \sum_{i \in \mathcal{J}} m_i Q_{i, \leq c}(0) + \mathcal{C}. \quad (37)$$

We prove this using (29) and (34) in conjunction with Condition $C2'$ as follows:

$$\begin{aligned} & 2 \sum_{(i,\ell) \in \mathcal{J} \times \mathcal{J}} Z_{i\ell} m_\ell Q_{\ell, \leq c}(0) E[\Delta_{i, \leq c} \mid \mathcal{F}(0)] \\ & \leq 2 \sum_{(i,\ell) \in \mathcal{J} \times \mathcal{J}} Z_{i\ell} m_\ell Q_{\ell, \leq c}(0) \left(\rho_i + \varepsilon m_i + \frac{\mathcal{C}}{Q_{\ell, \leq c}(0) + 1} - E \left[\sum_{j \in \mathcal{D}_i} \int_0^1 \beta_j \sigma_j(t) dt \mid \mathcal{F}(0) \right] \right) \\ & \leq 2 \sum_{(i,\ell) \in \mathcal{J} \times \mathcal{J}} Z_{i\ell} m_\ell Q_{\ell, \leq c}(0) \left(\rho_i + \varepsilon m_i - E \left[\sum_{j \in \mathcal{D}_i} \int_0^1 \beta_j \sigma_j(t) dt \mid \mathcal{F}(0) \right] \right) + \mathcal{C} \\ & \leq -\eta \sum_{i \in \mathcal{J}} m_i Q_{i, \leq c}(0) + \mathcal{C}, \end{aligned}$$

where we again remind the reader that the constant \mathcal{C} may differ from line to line. This completes the proof of Lemma 6.3 for non-synchronized networks.

Synchronized network. Now we consider the case when the network is synchronized, i.e., Condition $C1$. We establish the same three steps as in the non-synchronized case. In the non-synchronized case, we used the fact that the schedule $\sigma(t)$ under the ε -LRFS policy is maximal with respect to $[m_i Q_{i, \leq c}(t)]$, for which we required Condition $C2'$. In synchronized networks, as a first step in the proof, we use a different (stronger) maximality property, which allows us to relax Condition $C2'$ to Condition $C1$. To this end, we introduce some necessary notation. We let

$\sigma_{j,\leq c}(t) = 1$ if activity j processes a job with counter $\leq c$ at time t (and $\sigma_{j,\leq c}(t) = 0$ otherwise). Since $\sigma_{j,\leq c}(t) = \sigma_{j,\leq c}(\lfloor t \rfloor)$ in synchronized networks, we write $\sigma_{j,\leq c} = \sigma_{j,\leq c}(0) = \sigma_{j,\leq c}(t)$ for $t \in [0, 1)$. The main maximality property we use in synchronized networks is that, under the LRFS policy, the schedule $[\sigma_{j,\leq c} I_{\mathcal{J}^{(h)}}(i_j)]$ is maximal with respect to $[Q_{i,\leq c}(0) I_{\mathcal{J}^{(h)}}(i)]$ for each component $\mathcal{J}^{(h)}$. Together with (3), this implies that for every partition $\mathcal{J}^{(h)}$,

$$2 \sum_{(i,\ell) \in \mathcal{J}^{(h)} \times \mathcal{J}^{(h)}} Z_{i\ell} Q_{\ell,\leq c}(0) \left(\rho_i + \varepsilon m_i - \sum_{j \in \mathcal{J}_i} E \left[\sigma_{j,\leq c} \mid \mathcal{F}(0), \mathcal{E}_{\text{LRFS}}^{(h)} \right] \right) \leq -\eta \sum_{i \in \mathcal{J}} Q_{i,\leq c}(0) + \mathcal{C}, \quad (38)$$

where we let $\mathcal{E}_{\text{LRFS}}^{(h)}$ denote the event that at time 0 the ε -LRFS policy for component $\mathcal{J}^{(h)}$ does not select a job for processing through step 3-1 (see its description in Definition 3.2) and all selected jobs are due the LRFS policy in step 4. We stress that in synchronized networks, every processor completes the service requirement of the job it processes at every integer time $t \in \mathbb{Z}_+$ and hence $\Pr \left[\mathcal{E}_{\text{LRFS}}^{(h_j)} \right] = 1 - \varepsilon$. Inequality (38) is analogous to (29), which concludes the first step in the non-synchronized case.

We proceed with the analog of the second step from the non-synchronized case, i.e., bounding $E \left[\widehat{W}_{i,\leq c}(1) - \widehat{W}_{i,\leq c}(0) \mid \mathcal{F}(0) \right]$. We stress that the definition of $\widehat{W}_{i,\leq c}(t)$ differs from the one used in non-synchronized networks, see Section 6.1.1. Scheduling decisions are only made at integer time epochs (i.e., $\sigma(t) = \sigma(\lceil t \rceil)$) in synchronized networks, so that

$$\begin{aligned} E \left[\widehat{W}_{i,\leq c}(1) - \widehat{W}_{i,\leq c}(0) \mid \mathcal{F}(0) \right] &= E \left[\widehat{A}_{i,\leq c} \mid \mathcal{F}(0) \right] - \sum_{j \in \mathcal{J}_i} E \left[\sigma_{j,\leq c} \mid \mathcal{F}(0) \right] \\ &\leq \rho_i - \sum_{j \in \mathcal{J}_i} E \left[\sigma_{j,\leq c} \mid \mathcal{F}(0) \right], \end{aligned}$$

where we again let $\widehat{A}_{i,\leq c}$ be the number of job arrivals contributing to an increase in $\widehat{Q}_{i,\leq c}(\cdot)$ during the time interval $[0, 1)$. On writing $\Delta_{i,\leq c} := \widehat{W}_{i,\leq c}(1) - \widehat{W}_{i,\leq c}(0)$, we have

$$E \left[\Delta_{i,\leq c} \mid \mathcal{F}(0) \right] \leq \rho_i - \sum_{j \in \mathcal{J}_i} E \left[\sigma_{j,\leq c} \mid \mathcal{F}(0) \right]. \quad (39)$$

The third step in the proof for the synchronized case is to prove the conclusion of Lemma 6.3. As in (37), it suffices to prove that

$$2 \sum_{(i,\ell) \in \mathcal{J} \times \mathcal{J}} m_\ell Z_{i\ell} Q_{\ell,\leq c}(0) E \left[\Delta_{i,\leq c} \mid \mathcal{F}(0) \right] \leq -\eta \sum_{i \in \mathcal{J}} m_i Q_{i,\leq c}(0) + \mathcal{C}.$$

Since $m_i = 1$ and $V_i(0) \leq 1$ in synchronized networks, this reduces to

$$2 \sum_{(i,\ell) \in \mathcal{J} \times \mathcal{J}} Z_{i\ell} Q_{\ell,\leq c}(0) E \left[\Delta_{i,\leq c} \mid \mathcal{F}(0) \right] \leq -\eta \sum_{i \in \mathcal{J}} Q_{i,\leq c}(0) + \mathcal{C}. \quad (40)$$

Combining (39) and (40), it suffices to show that

$$2 \sum_{(i,\ell) \in \mathcal{J} \times \mathcal{J}} Z_{i\ell} Q_{\ell,\leq c}(0) \left(\rho_i - \sum_{j \in \mathcal{J}_i} E \left[\sigma_{j,\leq c} \mid \mathcal{F}(0) \right] \right) \leq -\eta \sum_{i \in \mathcal{J}} Q_{i,\leq c}(0) + \mathcal{C}. \quad (41)$$

For activity $j \in \mathcal{J}$, writing $\mathcal{J}^{(h_j)}$ for the component of buffer i_j (i.e., $i_j \in \mathcal{J}^{(h_j)}$), we have

$$\begin{aligned} E[\sigma_{j,\leq c} \mid \mathcal{F}(0)] &\geq \Pr\left[\mathcal{E}_{\text{LRFS}}^{(h_j)}\right] E\left[\sigma_{j,\leq c} \mid \mathcal{F}(0), \mathcal{E}_{\text{LRFS}}^{(h_j)}\right] \\ &= (1 - \varepsilon) E\left[\sigma_{j,\leq c} \mid \mathcal{F}(0), \mathcal{E}_{\text{LRFS}}^{(h_j)}\right] \\ &\geq E\left[\sigma_{j,\leq c} \mid \mathcal{F}(0), \mathcal{E}_{\text{LRFS}}^{(h_j)}\right] - \varepsilon. \end{aligned}$$

Therefore, (41) follows after arguing that

$$2 \sum_{(i,\ell) \in \mathcal{J} \times \mathcal{J}} Z_{i\ell} Q_{\ell,\leq c}(0) \left(\rho_i + \varepsilon m_i - \sum_{j \in \mathcal{J}_i} E\left[\sigma_{j,\leq c} \mid \mathcal{F}(0), \mathcal{E}_{\text{LRFS}}^{(h_j)}\right] \right) \leq -\eta \sum_{i \in \mathcal{J}} Q_{i,\leq c}(0) + \mathcal{C}.$$

The above inequality follows from (38) and Condition *C1*. This completes the proof of Lemma 6.3 for synchronized networks.

6.1.5 Proof of Lemma 6.4

For notational convenience, we again restrict attention to the case $t = 0$ in the conclusion of Lemma 6.4, namely, we show that for some $D, T, \mathcal{C}, \gamma_1, \gamma_2 \in (0, \infty)$,

$$E\left[\mathcal{G}(Y(T))^2 \mid \mathcal{F}(0)\right] \leq \mathcal{G}(Y(0))^2 - \gamma_2 \|\mathbf{Q}_{>D}(0)\|_1 + \mathcal{C} \left(\|\mathbf{Q}_{\leq D}(0)\|_1 + \|\mathbf{V}(0)\|_1 + 1 \right),$$

where we recall that

$$\mathcal{G}(Y(t)) = M_1(t) + M_2(t) + \gamma_1 \|\mathbf{V}(t)\|_1.$$

All arguments are applicable for general $t \geq 0$ as well. First observe that $M(t) := M_1(t) + M_2(t)$ can only change through the following events for jobs of Type 1 and 2.

Arrivals. $M(t)$ increases when new external arrivals of Type 2 occur. Note that there are no such external arrivals for Type 1.

Routing. $M(t)$ may increase or decrease when a job with counter $\geq D$ (i.e., Type 1 or 2) is routed since the weight (i.e., future workload) of a job conditioned on the buffer to which it has been routed is different from the (unconditional) weight before it is routed. However, $M(t)$ does not change when a job with counter $< D$ (i.e., Type 2) is routed since it is routed deterministically.

Starting Service. $M(t)$ may increase or decrease when a job of Type 1 or 2 begins service, generating its service time at this point. Assuming the job is served from buffer i , then $M(t)$ increases if the random service time is larger than its mean m_i , and decreases otherwise.

Being in Service. $M(t)$ decreases when a job of Type 1 or 2 is currently being processed.

Now we express $M(t)$ as follows: for $t \in \mathbb{N}$,

$$M(t) = M(0) + M_{\text{arrival}}(t) + M_{\text{routing}}(t) + M_{\text{s-service}}(t) + M_{\text{b-service}}(t),$$

where $M_{\text{arrival}}(t)$, $M_{\text{routing}}(t)$, $M_{\text{s-service}}(t)$ and $M_{\text{b-service}}(t)$ describe the change in $M(t)$ in the time interval $[0, t]$ due to events of new arrivals, routing, starting service and being in service, respectively. Hence,

$$M_{\text{arrival}}(t) \geq 0 \quad \text{and} \quad M_{\text{b-service}}(t) \leq 0.$$

From our definition of the weights, one can further observe that

$$\begin{aligned} E[M_{\text{routing}}(t) \mid \mathcal{F}(0)] &= 0 & E[M_{\text{s-service}}(t) \mid \mathcal{F}(0)] &= 0 \\ E[M_{\text{arrival}}(t) \mid \mathcal{F}(0)] &\leq t \sum_{d=D}^{\infty} d m_{\max} \|P^d \boldsymbol{\alpha}\|_1, \end{aligned} \quad (42)$$

where we define $m_{\max} := \max_{i \in \mathcal{J}} m_i$.

By appropriately defining constants $\gamma, T, D, \gamma_1, \gamma_2 \in (0, \infty)$, we first prove the following.

$$E[\mathcal{G}(Y(T)) - \mathcal{G}(Y(0)) \mid \mathcal{F}(0)] \leq \begin{cases} -\gamma & \text{if } \mathbf{Q}_{>D}(0) \neq 0 \\ \mathcal{C} & \text{otherwise} \end{cases}, \quad (43)$$

for some constant $\mathcal{C} < \infty$. It is not hard to prove (43) for $\mathbf{Q}_{>D}(0) = 0$, and hence we only provide the proof for $\mathbf{Q}_{>D}(0) \neq 0$. Consider the two complementary events: $\|\mathbf{V}(0)\|_1 \geq 2JB_{\text{renewal}}$ and $\|\mathbf{V}(0)\|_1 < 2JB_{\text{renewal}}$, where B_{renewal} is some constant which will be determined later.

First case. On the event $\{\|\mathbf{V}(0)\|_1 \geq 2JB_{\text{renewal}}\}$, we observe that $\|\mathbf{V}(t)\|_1 = \sum_{j \in \mathcal{J}} V^j(t)$ and

$$E[V^j(T) \mid \mathcal{F}(0)] \leq \begin{cases} V^j(0) - \beta_j T & \text{if } V^j(0) > \beta_j T \\ B_{\text{renewal}} & \text{otherwise} \end{cases}, \quad (44)$$

where one can find an appropriate constant $B_{\text{renewal}} < \infty$ depending on the variances of the generic service times $\{\Gamma_i\}$ using the renewal theory (e.g., $B_{\text{renewal}} = \max_{i \in \mathcal{J}} E[\Gamma_i^2]/E[\Gamma_i]$ from the proof of Proposition 6.2 in [1]). Hence, in case there is no $j \in \mathcal{J}$ satisfying $V^j(0) > \beta_j T$,

$$\begin{aligned} E[\|\mathbf{V}(T)\|_1 \mid \mathcal{F}(0)] &= \sum_{j \in \mathcal{J}} E[V^j(T) \mid \mathcal{F}(0)] \\ &\leq JB_{\text{renewal}} \\ &\leq \|\mathbf{V}(0)\|_1 - JB_{\text{renewal}}, \end{aligned} \quad (45)$$

since $\|\mathbf{V}(0)\|_1 \geq 2JB_{\text{renewal}}$. On the other hand, if there exists a $j_0 \in \mathcal{J}$ satisfying $V^{j_0}(0) > \beta_{j_0} T$, we have

$$\begin{aligned} E[\|\mathbf{V}(T)\|_1 \mid \mathcal{F}(0)] &= \sum_{j \in \mathcal{J}} E[V^j(T) \mid \mathcal{F}(0)] \\ &= E[V^{j_0}(T) \mid \mathcal{F}(0)] + \sum_{j \in \mathcal{J} \setminus \{j_0\}} E[V^j(T) \mid \mathcal{F}(0)] \\ &\leq V^{j_0}(0) - \beta_{j_0} T + \sum_{j \in \mathcal{J} \setminus \{j_0\}} \max\{V^j(0) - \beta_j T, B_{\text{renewal}}\} \\ &\leq V^{j_0}(0) - \beta_{j_0} T + \sum_{j \in \mathcal{J} \setminus \{j_0\}} (V^j(0) + B_{\text{renewal}}) \\ &\leq \|\mathbf{V}(0)\|_1 + JB_{\text{renewal}} - \beta_{\min} T \\ &\leq \|\mathbf{V}(0)\|_1 - JB_{\text{renewal}}, \end{aligned} \quad (46)$$

where we use (44) and choose

$$T > \left\lceil \frac{2JB_{\text{renewal}}}{\beta_{\min}} + 1 \right\rceil. \quad (47)$$

Therefore, in both (45) and (46), we have

$$E[\|\mathbf{V}(T)\|_1 \mid \mathcal{F}(0)] \leq \|\mathbf{V}(0)\|_1 - JB_{\text{renewal}}.$$

Using this, it follows that

$$\begin{aligned} E[\mathfrak{G}(Y(T)) - \mathfrak{G}(Y(0)) \mid \mathcal{F}(0)] &\leq E[M(T) - M(0) \mid \mathcal{F}(0)] + \gamma_1 E[\|\mathbf{V}(T)\|_1 - \|\mathbf{V}(0)\|_1 \mid \mathcal{F}(0)] \\ &\leq E[M_{\text{arrival}}(T) \mid \mathcal{F}(0)] - \gamma_1 JB_{\text{renewal}} \\ &\leq T \sum_{d=D}^{\infty} d m_{\max} \|P^d \boldsymbol{\alpha}\|_1 - \gamma_1 JB_{\text{renewal}} \\ &\leq -\gamma, \end{aligned}$$

where we use (42) and now define D, γ_1 in terms of γ as follows:

$$D = D(\gamma) := \min \left\{ x \in \mathbb{N} : T \sum_{d=x}^{\infty} d m_{\max} \|P^d \boldsymbol{\alpha}\|_1 \leq \gamma \right\} \quad \gamma_1 = \gamma_1(\gamma) := \frac{2\gamma}{JB_{\text{renewal}}}.$$

We specify the value of γ at a later stage in the proof. This completes the proof of (43) on the first event $\{\|\mathbf{V}(0)\|_1 \geq 2JB_{\text{renewal}}\}$.

Second case. Now consider the second event $\{\|\mathbf{V}(0)\|_1 < 2JB_{\text{renewal}}\}$. From our choice of T in (47), we have

$$\max_{j \in \mathcal{J}} \frac{V^j(0)}{\beta_j} \leq \frac{\|\mathbf{V}(0)\|_1}{\beta_{\min}} < \frac{2JB_{\text{renewal}}}{\beta_{\min}} < T - 1,$$

which implies that before time $T - 1$, all activities have completed serving the jobs they were serving at time 0 (and they could have worked on other jobs as well). Thus, it follows that

$$E[\|\mathbf{V}(T)\|_1 \mid \mathcal{F}(0)] \leq JB_{\text{renewal}}. \quad (48)$$

Consider a job with the largest counter (and hence, contributing to $\mathbf{Q}_{>D}(0)$) at time 0 and let $\mathcal{J}^{(h_0)} \subset \mathcal{J}$ be the component the job belongs to. Define the event $\mathcal{E}_{\text{step 3-1}}$ that the ε -LRFS policy executes step 3-1 (see the description of ε -LRFS in Definition 3.2) for this component at least once before time $T - 1$. Let $\mathcal{E}_{\text{step 3-1}}^*$ be the subevent that the job identified in step 1 of the policy is selected for processing when step 3-1 is carried out for the first time, so that $\Pr[\mathcal{E}_{\text{step 3-1}}^* \mid \mathcal{E}_{\text{step 3-1}}] = \varepsilon$. On the event $\mathcal{E}_{\text{step 3-1}}^*$, let the random variable j^* denote the activity which is chosen to process this job, and let X^* be the associated service time. We then have that $E[X^* \mid j^*] = m_{i_{j^*}}$. Observe that

$$\begin{aligned} E[M_{\text{s-service}}(T) + M_{\text{b-service}}(T) \mid j^*, X^*, \mathcal{E}_{\text{step 3-1}}^*, \mathcal{F}(0)] \\ \leq \begin{cases} -m_{i_{j^*}} + X^* - \beta_{j^*} & \text{on the event } \{X^* \geq \beta_{j^*}\} \\ -m_{i_{j^*}} & \text{otherwise} \end{cases}. \end{aligned}$$

It thus follows that

$$\begin{aligned}
& E[M_{\text{s-service}}(T) + M_{\text{b-service}}(T) \mid \mathcal{E}_{\text{step 3-1}}^*, \mathcal{F}(0)] \\
& \leq E \left[-m_{i_{j^*}} + [X^* - \beta_{j^*}]_+ \mid \mathcal{E}_{\text{step 3-1}}^*, \mathcal{F}(0) \right] \\
& \leq E \left[-m_{i_{j^*}} + [X^* - \beta_{\min}]_+ \mid \mathcal{E}_{\text{step 3-1}}^*, \mathcal{F}(0) \right] \\
& = E \left[E \left[-m_{i_{j^*}} + [X^* - \beta_{\min}]_+ \mid j^*, X^*, \mathcal{E}_{\text{step 3-1}}^*, \mathcal{F}(0) \right] \mid \mathcal{E}_{\text{step 3-1}}^*, \mathcal{F}(0) \right] \\
& \leq \max_{i \in \mathcal{J}} E[-m_i + [\Gamma_i - \beta_{\min}]_+],
\end{aligned}$$

where Γ_i is a generic service time for buffer i and $[x]_+ = x$ if $x \geq 0$ and $[x]_+ = 0$ otherwise. It is easy to see that $E[-m_i + [\Gamma_i - \beta_{\min}]_+] < 0$ for all i since $E[\Gamma_i] = m_i > 0$ and $\beta_{\min} > 0$. Hence, the above inequality implies that

$$E[M_{\text{s-service}}(T) + M_{\text{b-service}}(T) \mid \mathcal{E}_{\text{step 3-1}}^*, \mathcal{F}(0)] \leq -\varepsilon\nu, \quad (49)$$

with

$$\nu := \min_{i \in \mathcal{J}} E[m_i - [\Gamma_i - \beta_{\min}]_+] > 0.$$

If the network is synchronized or if the second part of Condition $C2$ of Theorem 3.5 holds, the only way for the event $\mathcal{E}_{\text{step 3-1}}$ not to occur is that there exists a processor (for the component $\mathcal{J}^{(h_0)}$) processing a job constantly during the entire time interval $[\mathcal{T}^{(h_0)}(0), T-1]$ (i.e., the job starts service before time $\mathcal{T}^{(h_0)}(0)$ and is still in service at time $T-1$). Recall that before time $T-1$, every processor completes the service requirement of the job it was processing at time 0. Hence, in addition to (47), if we choose T to also satisfy

$$T > 2 \max_{i \in \mathcal{J}} m_i + 2,$$

then it follows that

$$\Pr[\mathcal{E}_{\text{step 3-1}} \mid \mathcal{F}(0)] \geq \left(\frac{1}{2}\right)^K, \quad (50)$$

where we use the fact that if a processor starts to process a new job in the time interval $(0, \mathcal{T}^{(h_0)}(0)] \subset (0, 1]$, its service requirement is at most $2 \max_{i \in \mathcal{J}} m_i$ with probability $1/2$ by the Markov inequality. From (49) and (50), we conclude that

$$E[M_{\text{s-service}}(T) + M_{\text{b-service}}(T) \mid \mathcal{F}(0)] \leq -\frac{\varepsilon\nu}{2^K}. \quad (51)$$

From (42), (48) and (51), the desired inequality (43) follows as

$$\begin{aligned}
& E[\mathcal{G}(Y(T)) - \mathcal{G}(Y(0)) \mid \mathcal{F}(0)] \\
& \leq E[M(T) - M(0) \mid \mathcal{F}(0)] + \gamma_1 E[\|\mathbf{V}(T)\|_1 - \|\mathbf{V}(0)\|_1 \mid \mathcal{F}(0)] \\
& \leq E[M_{\text{arrival}}(T) \mid \mathcal{F}(0)] + E[M_{\text{s-service}}(T) + M_{\text{b-service}}(T) \mid \mathcal{F}(0)] + \gamma_1 JB_{\text{renewal}} \\
& \leq T \sum_{d=D}^{\infty} d m_{\max} \|P^d \boldsymbol{\alpha}\|_1 - \frac{\varepsilon\nu}{2^K} + \gamma_1 JB_{\text{renewal}} \\
& \leq -\gamma,
\end{aligned}$$

where for the last inequality we define

$$\gamma := \frac{\varepsilon \nu}{2^{K+2}}.$$

Here we note that $D = D(\gamma) < \infty$ as long as $\gamma > 0$ (i.e., $\varepsilon > 0$). This completes the proof of (43) on the second event $\{\|\mathbf{V}(0)\|_1 < 2JB_{\text{renewal}}\}$.

Completing the proof of Lemma 6.4. Now (43) implies that

$$\begin{aligned} E[\mathfrak{G}(Y(T))^2 - \mathfrak{G}(Y(0))^2 \mid \mathcal{F}(0)] &\leq 2\mathfrak{G}(Y(0)) E[\mathfrak{G}(Y(T)) - \mathfrak{G}(Y(0)) \mid \mathcal{F}(0)] + \mathfrak{C} \\ &\leq 2 \left(-\gamma \mathbf{1}_{\|\mathbf{Q}_{>D}(0)\|_1}^+ + \mathfrak{C} \left(1 - \mathbf{1}_{\|\mathbf{Q}_{>D}(0)\|_1}^+ \right) \right) \mathfrak{G}(Y(0)) + \mathfrak{C} \\ &= 2 \left(-(\gamma + \mathfrak{C}) \mathbf{1}_{\|\mathbf{Q}_{>D}(0)\|_1}^+ + \mathfrak{C} \right) \mathfrak{G}(Y(0)) + \mathfrak{C} \end{aligned}$$

where one can check $E \left[(\mathfrak{G}(Y(T)) - \mathfrak{G}(Y(0)))^2 \mid \mathcal{F}(0) \right] < \infty$ for the first inequality and the precise value of \mathfrak{C} can be different from line to line. Finally, we define

$$\gamma_2 := m_{\min} \gamma$$

and the conclusion of Lemma 6.4 follows from

$$\begin{aligned} E[\mathfrak{G}(Y(T))^2 - \mathfrak{G}(Y(0))^2 \mid \mathcal{F}(0)] &\leq 2 \left(-(\gamma + \mathfrak{C}) \mathbf{1}_{\|\mathbf{Q}_{>D}(0)\|_1}^+ + \mathfrak{C} \right) \mathfrak{G}(Y(0)) + \mathfrak{C} \\ &\leq -2\gamma m_{\min} \|\mathbf{Q}_{>D}(0)\|_1 + \mathfrak{C} (\|\mathbf{Q}_{\leq D}(0)\|_1 + \|\mathbf{V}(0)\|_1 + 1) \\ &= -\gamma_2 \|\mathbf{Q}_{>D}(0)\|_1 + \mathfrak{C} (\|\mathbf{Q}_{\leq D}(0)\|_1 + \|\mathbf{V}(0)\|_1 + 1), \end{aligned}$$

where we use $\mathfrak{G}(Y(0)) \leq \mathfrak{C} (\|\mathbf{Q}_{\leq D}(0)\|_1 + \|\mathbf{V}(0)\|_1 + 1)$ for the case $\mathbf{Q}_{>D}(0) = 0$ and $\mathfrak{G}(Y(0)) \geq m_{\min} \|\mathbf{Q}_{>D}(0)\|_1$ for the other case $\mathbf{Q}_{>D}(0) \neq 0$ (i.e., $\|\mathbf{Q}_{>D}(0)\|_1 > 0$).

Acknowledgments

ABD gratefully acknowledges NSF grant EEC-0926308 for financial support.

References

- [1] S. Asmussen. Applied probability and queues. *Springer Verlag*, 2003.
- [2] M. Bramson. Stability of earliest-due-date, first-served queueing networks. *Queueing systems*, 39(1):79-102, 2001.
- [3] J. G. Dai. On positive Harris recurrence of multiclass queueing networks: a unified approach via fluid limit models. *Annals of Applied Probability*, 5:49-77, 1995.
- [4] J. G. Dai. Stability of open multiclass queueing networks via fluid models. *IMA Workshop on Stochastic Networks*, Springer-Verlag, New York, 71-90, 1995.
- [5] J. G. Dai and W. Lin. Asymptotic optimality of maximum pressure policies in stochastic processing networks. *Annals of Applied Probability*, 18(6):2239-2299, 2008.

- [6] J. G. Dai and B. Prabhakar. The throughput of data switches with and without speedup. *IEEE INFOCOM*, 556-564, Tel-Aviv, Israel, 2000.
- [7] J. G. Dai and G. Weiss. Stability and instability of fluid models for reentrant lines. *Mathematics of Operations Research*, 21:115-134, 1996.
- [8] A. B. Dieker and X. Gao. Positive recurrence of piecewise Ornstein-Uhlenbeck processes and common quadratic Lyapunov functions. Arxiv preprint arXiv:1107.2873, 2011.
- [9] P. Dupuis and R. J. Williams. Lyapunov functions for semimartingale reflecting Brownian motions. *Annals of Probability*, 22:680-702, 1994.
- [10] S. Foss and T. Konstantopoulos. An overview of some stochastic stability methods. *Journal of the Operations Research Society of Japan*, 47(4):275-303, 2004.
- [11] J. M. Harrison. Heavy traffic analysis of a system with parallel servers: Asymptotic optimality of discrete-review policies. *Annals of Applied Probability*, 8:822-848, 1998.
- [12] J. M. Harrison. Brownian models of open processing networks: Canonical representation of workload. *Annals of Applied Probability*, 10:75-103, MR1765204, 2000.
- [13] F. P. Kelly. Stochastic models of computer communication systems. *J. R. Statist. Soc B*, 47(3):379.395, 1985.
- [14] P. R. Kumar and S. P. Meyn. Stability of queueing networks and scheduling policies. *IEEE Transactions on Automatic Control*, 40(2):251-260, 1995.
- [15] P. R. Kumar and T. I. Seidman. Dynamic instabilities and stabilization methods in distributed real-time scheduling of manufacturing systems. *IEEE Transactions on Automatic Control*, 35:289298, 1990.
- [16] M. Leconte, J. Ni and R. Srikant. Improved Bounds on the Throughput Efficiency of Greedy Maximal Scheduling in Wireless Networks. *IEEE/ACM Transactions on Networking*, 19(3):709-720, 2011.
- [17] S. H. Lu and P. R. Kumar. Distributed scheduling based on due dates and buffer priorities. *IEEE Transactions on Automatic Control*, 36:14061416, 1991.
- [18] N. McKeown, V. Anantharan and J. Walrand. Achieving 100% throughput in an input-queued switch. *IEEE INFOCOM*, 296-302, San Francisco, 1996.
- [19] J. Mo and J. Walrand. Fair end-to-end window-based congestion control. *IEEE/ACM Transactions on Networking*, 8(5):556.567, 2000.
- [20] A. Rybko and A. Stolyar. On the ergodicity of stochastic processes describing open queueing networks. *Problemi Peredachi Informatsii*, 28:3-26, 1992.
- [21] S. Sanghavi, L. Bui and R. Srikant. Distributed link scheduling with constant overhead. *ACM SIGMETRICS Performance Evaluation Review*, 35(1):313-324, 2007.
- [22] D. Shah, J. Shin and P. Tetali. Medium Access using Queues. Annual IEEE Symposium on Foundations of Computer Science (FOCS), Palm Springs, California, 2011.

- [23] D. Shah and D. Wischik. Switched networks with maximum weight policies: Fluid approximation and multiplicative state space collapse. To appear in *Annals of Applied Probability*, Arxiv preprint arXiv:1004.1995, 2010.
- [24] 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.
- [25] X. Wu, R. Srikant and J. R. Perkins. Scheduling efficiency of distributed greedy scheduling algorithms in wireless networks. *Mobile Computing, IEEE Transactions on*, 6(6):595-605, 2007.