

SUFFICIENT CONDITIONS FOR STABILITY OF LONGEST QUEUE FIRST SCHEDULING: SECOND ORDER PROPERTIES USING FLUID LIMITS

ANTONIS DIMAKIS,* *University of California at Berkeley*

JEAN WALRAND,** *University of California at Berkeley*

Abstract

We consider the stability of *Longest Queue First* (LQF), a natural and low complexity scheduling policy, for a *generalized switch* model [1]. Contrary to common scheduling policies, the stability of LQF depends on the variance of the arrival processes in addition to their average intensities.

We identify new sufficient conditions for LQF to be throughput optimal for i.i.d. arrival processes. Deterministic fluid analogs, proved to be powerful in the analysis of stability in queueing networks, do not adequately characterize the stability of LQF. We combine properties of diffusion-scaled sample path functionals and local fluid limits into a sharper characterization of stability.

Keywords: Generalized switch, Longest Queue First, MaxWeight scheduling, throughput optimality, stability, local pooling, fluid limit, local fluid limit, second order conditions

2000 Mathematics Subject Classification: Primary 60K25

Secondary 90B15; 60G17

1. Introduction

We consider the scheduling of the *generalized switch* model described in [1], which includes as special cases models of multiuser data scheduling over a wireless medium, input-queued crossbar switches, and parallel server systems. The problem of throughput optimal scheduling is addressed in [1] where *MaxWeight* scheduling is proposed. This algorithm is based on that in [2] originally proposed in the context of data

* Postal address: Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, Berkeley, CA 94720, USA

scheduling in a multihop wireless network. A natural low complexity alternative to MaxWeight is *Longest Queue First* (LQF) scheduling, proposed in [4]. The main result of this paper is Theorem 1, which gives sufficient conditions for throughput optimality of LQF.

The achievable throughput of common scheduling policies in open queueing networks, e.g., first-in first-out, head-of-line processor sharing, or MaxWeight in a generalized switch is characterized by average arrival intensities, to the effect that the problem of stability is related to that of a deterministic fluid analog. The connection between the stochastic and the fluid system, based on the strong law of large numbers (SLLN), has been formalized in [9, 6, 8] and applied with great success in multiclass queueing networks [6] and generalized switch models [3, 1, 7, 5]. The premise is that in many cases it is easier to work with a deterministic system than with a stochastic one.

In LQF, on the other hand, variability in stochastic arrivals may affect stability, as will become evident in Sections 2 and 4. In this case, a deterministic analog is not detailed enough to infer stability since deterministic and nondeterministic arrivals with the same intensities, can lead to unstable and stable systems, respectively.

Previous work on LQF [10] considered sufficient conditions for stability by analyzing the associated deterministic systems. This led to stronger than necessary conditions on the average rate of arrivals for certain classes of systems. We show that in certain special cases, namely systems satisfying the rank condition in Theorem 1, LQF is throughput optimal under the assumption of nondeterministic arrivals. Although the systems satisfying the conditions in Theorem 1 are special, we believe that the techniques used in its proof are interesting on their own right. A novel feature of our analysis is that we combine diffusion-scale properties of the sample paths with the fluid limit framework of [6, 1].

The rest of the paper is organized as follows. In Section 2 we give an informal discussion of our model and provide some intuition for the proofs to be given in Section 4. In Section 3 we introduce the model and notation. Section 4 is devoted to proving the main result of this paper, Theorem 1. In Appendix C we discuss how our results extend to general service rates.

2. Discussion and examples

In this section we give an informal description of our results by means of examples. We start by giving a graph representation of our model. See Section 3 for a rigorous treatment.

A model is specified by a graph (V, E) , with V and E being the vertex and edge set, respectively. One can think of the vertices as queues where arrivals occur in discrete time at some given average rate. We consider i.i.d. arrival processes, although it is not hard to generalize our results to finite-Markov-modulated processes. When a queue is served, a unit of work is removed from that queue, provided it is nonempty. Not all queues/vertices can be served during the same time slot: if a queue is served, then its neighboring queues are not. Hence, at each time, the served queues form an independent set of the graph. LQF scheduling chooses this set iteratively, starting from the longest queue and proceeding in a decreasing manner. Queues that have one of their neighbors already selected are not considered in the next iteration step. When two or more queues under consideration have equal backlog, a tie-breaking rule must be specified. This procedure is stopped until no further (nonempty) queue can be included. At each time slot, the served queues form a maximal independent set of the subgraph consisting of the nonempty queues.

2.1. Three-Queue Example

As a first example, consider the graph with $V = \{1, 2, 3\}$ and $E = \{\{1, 2\}, \{2, 3\}\}$. Let λ_i be the arrival rate at queue i , for $i = 1, 2, 3$. In this paper, we are concerned with characterizing the set of vectors $\lambda = (\lambda_1, \lambda_2, \lambda_3)$ for which the queueing process is stable, i.e., positive recurrent, under LQF scheduling. Since either queue 2 or queues 1 and 3 can be served at any given time, we expect that a scheduler can be stable only if $\lambda_1 + \lambda_2 < 1$ and $\lambda_2 + \lambda_3 < 1$. Indeed, the scheduler cannot serve queues 1 and 2 at a total rate larger than 1 since it cannot serve both queues at the same time. The same observation applies to queues 1 and 3.

To see why LQF is expected to be stable when these conditions hold we argue that the longest queue decreases, on average. This is immediate if one queue i is the only longest queue for some interval of time. Indeed, in that case, queue i is constantly

served during that interval under LQF and its length tends to decrease because $\lambda_i < 1$. Now, it may happen that a subset L of queues alternate being the longest during some interval of time. For instance, assume that $L = \{1, 2\}$ so that the scheduler selects queue 2 some fraction of time during that interval and queues $\{1, 3\}$ the rest of the interval. During that interval of time, the scheduler always serves either queue 1 or queue 2, so that the total service rate of the queues 1 and 2 is equal to one. The total arrival rate of queues 1 and 2 is $\lambda_1 + \lambda_2 < 1$. Consequently, the length of the longest queue (which is that of queue 1 and also that of queue 2) decreases on average. The same argument can be made for any subset $L \subset \{1, 2, 3\}$. The key property is as follows. For any set of queues L that alternate being the longest, there is a subset F of L that is served at a constant total rate, independently of the selections that the scheduler makes, given that L is the set of longest queues. This constant service rate must be larger than the total arrival rate into the set F , otherwise the system would certainly be unstable. In that case, the longest queue, which is that of any queue in F , decreases on average. Note that this is a topological property of the graph. In our example, if $L = \{1, 2, 3\}$, one can choose $F = \{1, 2\}$. If $L = \{1, 3\}$, one can choose $F = \{1\}$, and so on.

The above condition is captured in our local pooling condition in Definition 1. A characterization of local pooling in terms of a linear program is given in Proposition 1. Using the fluid limit technique we are able to use an essentially deterministic argument to show stability under stochastic (and deterministic) arrivals (c.f. subsection 4.9).

Summing up, the argument shows that if the set of longest queues satisfies local pooling for some interval of time, then the longest queue tends to decrease during that interval. Consequently, if all sets satisfy local pooling, the system is stable. Thus, the stability conditions for systems that satisfy local pooling are only in terms of average intensities of arrivals. However, for other systems, the stability of LQF may depend on “second-order” properties. In particular, deterministic and nondeterministic arrival processes, with the same rates, may lead to unstable and stable behavior, respectively. We describe one such example next.

2.2. Six-Cycle

Consider the system specified by the 6-vertex cycle graph with $V = \{1, \dots, 6\}$ and $E = \{\{1, 2\}, \{2, 3\}, \dots, \{6, 1\}\}$ depicted in Figure 1. One can check that local pooling fails to hold. Indeed, if the set of longest queues is $L = \{1, 2, 3, 4, 5, 6\}$, there is no nonempty subset of L that LQF serves at a constant rate. For instance, LQF serves either 0 or 1 queue from the subset $\{1, 2\}$; it serves either 1 or 2 queue from $\{1, 2, 3\}$ and either 1 or 2 queues from $\{1, 2, 3, 4\}$, and similarly for all other subsets. Consequently, we cannot conclude the stability of this system using the argument in the three-queue example. In fact, the system may not be stable, as we explain next.

Assume that a constant (deterministic) amount λ of work arrives during each time slot at each queue. Furthermore, assume that the scheduler uses LQF and breaks ties by throwing an unbiased die with as many faces as there are tied queues. We claim that this system is unstable if $\lambda > 4/9$. What makes this fact interesting is that, as we show below, the same system with non-constant i.i.d. arrivals with average value λ is stable whenever $\lambda < 1/2$.

The maximal independent sets of this graph are $M_1 = \{1, 3, 5\}$, $M_2 = \{2, 4, 6\}$, $M_3 = \{1, 4\}$, $M_4 = \{2, 5\}$, and $M_5 = \{3, 6\}$. Assume the system starts with all queues having the same large backlog. Let us call the sets M_1 and M_2 “big matches” and the other sets $M_3 - M_5$ “small matches.” There are a number of possible sequences of choices that the LQF scheduler can make. These sequence correspond to choosing a small or big match at each step, whenever the choice is possible. The right part of the figure shows these sequences. Each rectangle represents a representative set of possible longest queues, up to a symmetric renumbering. For instance the queues start all equal; the scheduler chooses a small match, represented by S , with probability $1/3$; then it selects a big match B with probability $1/2$; then the scheduler is forced to pick a big match again, and so on. By analyzing this Markov chain that describes the choices of the scheduler, one finds that the queues become equal again after a random number of steps. Dividing the average number of times each queue is served by the average number of steps, we find that the scheduler serves each queue at the average rate of $4/9$. It follows that the system is unstable if $\lambda > 4/9$.

We now explain why this system is stable when $\lambda < 1/2$ if the arrivals are allowed




FIGURE 1: Left: A 6-cycle system; neighboring queues cannot be served at the same time. Middle: The separation argument. Right: Possible schedules.

to vary even slightly. We saw that the instability of the system with deterministic arrivals is caused by the locked step operations that cause the queue lengths to become all equal every so often, which occasionally leads LQF to serve only two queues. The randomness of the arrivals prevents the backlogs from being equal a significant fraction of the time.

The key idea is that the longest queue length tends to decrease whenever it is large. Specifically, for $n \gg 1$, there is some t_0 and some ϵ_* such that

$$\frac{1}{n} E[\max_i Q_i(nt_0) - \max_i Q_i(0) | Q(0)] \leq -\epsilon_* t_0 < 0, \text{ on } \{\max_i Q_i(0) > nh\}. \quad (1)$$

It follows that $\max_i Q_i(nt)$ is a Lyapunov function for the Markov chain, which implies its positive recurrence. To show (1), one considers a sequence of processes Q^n that behave as the original system but with initial conditions bounded by $O(n)$. One shows that $Q^n(nt)/n$ converges along a subsequence to $\bar{Q}(t)$, a deterministic system (called the “fluid limit”) that satisfies a set of differential equations. Moreover, one proves that

$$\bar{Q}(t) = 0 \text{ for } t \geq t_0. \quad (2)$$

This fact, together with the convergence, implies that

$$\frac{1}{n} \max_i Q_i(nt) \rightarrow 0, \text{ for } t \geq t_0.$$

With a uniform integrability argument, one concludes that $E[\max_i Q_i(nt)/n] \rightarrow 0$ for $t \geq t_0$, and this implies (1). This line of argument is standard. The novelty in the

paper is in the argument to show (2). The main step is to show that if $\max_i \bar{Q}_i(t) > 0$, then

$$\frac{d}{dt} \max_i \bar{Q}_i(t) \leq -\epsilon_* < 0.$$

If the set of longest queues satisfies the local pooling condition, then one obtains the inequality above as in the three-queue example. If it does not, then the argument is qualitatively different.

For the 6-cycle example, the only set of queues that does not satisfy the local pooling condition is the set of all queues. Assume then that all the queues are equally large at some time t , in this limit $\bar{Q}(t)$. This implies that all the queue lengths are of order $n\delta$ for the system Q^n at time nt . In that case, all the queues will remain nonempty for some interval of time of order $n\delta$, since there is at most one departure from any given queue in each step.

One then shows that the difference between the maximum and the minimum queue lengths is at least of order $6\Delta(n)$ for most of the times $s \in I := \{nt, nt + \Delta(n), nt + 2\Delta(n), \dots, nt + n\delta - \Delta(n)\}$, where $\Delta(n) = n^{1/6}$ and $t > 0$ is fixed. We explain the main steps of that argument below. Moreover, if that condition holds for some time $s = k\Delta(n) \in I$, there is a proper subset L of the queues that remain strictly longer than the others during the interval $I(s) := \{s, s + 1, \dots, s + \Delta(n) - 1\}$. (See Figure 2.) In addition, any proper subset of the queues satisfies the local pooling conditions. According to the argument we used for the three-queue example, it follows that the longest queue decreases, on average, during $I(s)$. Consequently, the longest queue decreases on average during most of the intervals $I(s)$ and, therefore, most of the time. This argument implies that

$$\frac{1}{n} E[\max_i Q_i(nt + n\delta) - \max_i Q_i(nt) | Q(nt)] \leq -\delta\epsilon_* < 0.$$

This fact implies the stability of the system.

We now explain why the queues remain separated for most of the times $s \in I$. The idea is that the difference between the maximum and the minimum queue lengths is lower-bounded by differences between cumulative arrivals. We approximate these differences by Gaussian random variables, after proper scaling. Using the Gaussian distribution, we find that these random variables cannot be small most of the time.

First, observe that if all the queues remain nonempty the number of queues served among queues $\{2, 3\}$ is always the same as among queues $\{5, 6\}$ under any service set M_i . Designate by $Q_i(t)$ the length of queue i at time t . If no queue empties, since the departures cancel out,

$$B(t) := Q_2(t) + Q_3(t) - Q_5(t) - Q_6(t) = B(0) + A_2(t) + A_3(t) - A_5(t) - A_6(t) =: B(0) + Z(t)$$

where $A_i(t)$ is the cumulative number of arrivals in queue i up to time t .

Second, note that $D(t) := \max_i Q_i(t) - \min_i Q_i(t) \geq |B(t)|/2$, so that

$$\frac{D(nt)}{\sqrt{n}} \geq \frac{1}{2\sqrt{n}} |B(0) + Z(nt)|.$$

Finally, since the arrivals are i.i.d., the random variable $Z(nt)/\sqrt{n}$ is approximately Gaussian and zero-mean, so that

$$\begin{aligned} E \sum_{m=1}^{n\delta} 1\{D(nt+m) \leq 6\Delta(n)\} &\approx n\delta P\left(\frac{D(nt)}{\sqrt{n}} \leq \frac{6\Delta(n)}{\sqrt{n}}\right) \\ &\lesssim n\delta P\left(\frac{|Z(nt)|}{2\sqrt{n}} \leq \frac{6\Delta(n)}{\sqrt{n}}\right) \lesssim O(\sqrt{n}\Delta(n)\delta), \end{aligned}$$

as $n \rightarrow \infty$. Choosing $\Delta(n)$ so that $\sqrt{n}\Delta(n) = o(n/\Delta(n))$ – for instance, $\Delta(n) = n^{1/6}$ – we conclude that the fraction of the values of $s \in I = \{nt, nt+\Delta(n), nt+2\Delta(n), \dots, nt+n\delta - \Delta(n)\}$ such that $D(s) \leq 6\Delta(n)$ is negligible, because $|I| = n\delta/\Delta(n)$. This step completes the argument.

Looking back, here are the main ideas of the argument. When the set of longest queues satisfies local pooling, the longest queue tends to decrease. Also, the set of all queues – for which local pooling fails – cannot remain longest for any significant fraction of time: the scheduler cannot fully compensate the fluctuations in the arrivals because the set of service vectors has only rank 4 in the 6-dimensional space of queue lengths. Summing up, the key stability condition is that when the set of longest queues does not satisfy local pooling, the rank of the corresponding service vectors is small (at most the number of queues minus two).

These ideas form the basis of Theorem 1, where we show that LQF is stable for any nondeterministic i.i.d. arrival processes, under a weakening of the local pooling condition.

When (the weakened) local pooling condition fails to hold, we have simulation evidence suggesting the system may be unstable, even for Bernoulli arrivals. One such

example is a system specified by an 8-vertex cyclic graph driven by Bernoulli arrivals at rate $\lambda_i = 0.4984$ for $i = 1, \dots, 8$; ties were broken as in the 6-vertex example.

Sufficient stability conditions for LQF in cyclic graphs, are considered in [10]. There, the analysis is based on the stability of the associated fluid system. It is found that if the total arrival rate of three consecutive queues on the cycle is strictly below 1, then the system is stable. Intuitively, if all queues alternate being the longest for some considerable amount of time, then the three queues satisfying the condition above must decrease because at least one of them is served at each timeslot. On the other hand, under nondeterministic arrivals, this condition on the rates is not necessary for the 6-cycle as described above. For systems not of the special form as the 6-cycle, i.e., those satisfying the rank condition in Theorem 1, it is possible to still derive stability by imposing stricter than feasibility, conditions on the arrival rates. [See remarks following Theorem 1.]

3. Model and notation

We consider a discrete time model of a set K of queues. Let $Q_k(t)$ denote the backlog of queue k at time slot $t \geq 0$ and $A_k(t)$ (resp., $D_k(t)$) the cumulative arrivals (resp., departures) of work at queue k up to time t . We assume that arrivals are mutually independent i.i.d. processes with $E[A_k(1)] = \lambda_k < \infty$ for every k .

The scheduler cannot serve all nonempty queues at a given time. The allowable schedules form a finite set $S[K] \subset \{0, 1\}^{|K|}$. At any time t , the scheduler selects a vector $m \in S[K]$ and removes $m_k \in \{0, 1\}$ unit of work from each queue k , except if that queue empties first. (We relax the $\{0, 1\}$ -assumption in Appendix C.)

By $M[K]$ we denote the maximal elements of $S[K]$ and we treat it as a matrix whose columns are its elements. For $L \subset K$, we use $S[L]$ to denote the restriction of $S[K]$ on L and $M[L]$ the maximal elements of $S[L]$. This model is a special case of the generalized switch in [1] and is similar to the model of the input-buffered packet switch in [5].

Given $M[K]$, the evolution of queue backlogs is completely determined by the arrival process $A(\cdot)$ and the scheduling policy. For $m \in M[K]$, $T_m(t)$ records the number of time slots up to and including t that the scheduling policy chooses m .

Thus, $(Q(\cdot), T(\cdot))$ satisfies the following equations:

$$Q(t) = Q(0) + A(t) - D(t), \quad (3)$$

$$D_k(t+1) - D_k(t) = \sum_{m \in M[K]} m_k (T_m(t+1) - T_m(t)), \text{ if } Q_k(t) > 0; 0 \text{ o.w.} \quad (4)$$

$$\sum_{m \in M[K]} T_m(t) = t, \quad \forall t \geq 0, \quad (5)$$

$$Q_k(t) \geq 0, \quad (6)$$

$$D_k(t), T_m(t) \geq 0, \text{ nondecreasing } \forall k \in K, m \in M[K], t \geq 0. \quad (7)$$

We consider all the above processes defined on noninteger times as well. For noninteger u , $D(u) = D(\lfloor u \rfloor)$. The other processes are extended similarly.

Note that (3)-(7) alone do not imply non-idleness. A *Longest Queue First* (LQF) scheduling policy is a stationary policy that fails to serve a nonempty queue only if it conflicts with a longer queue. That is, at every state the scheduler chooses a vector $m \in M[K]$ such that if $m_k = 0$ and $Q_k(t) > 0$, there exists a queue k' with $m_{k'} = 1$ and $n_k n_{k'} = 0$ for all $n \in M[K]$. This condition implies that k conflicts with k' and $Q_{k'}(t) \geq Q_k(t)$. To completely describe such a policy one has to specify (stationary) rules for breaking ties in the queue backlogs. Once such rules are fixed, the resulting process $(Q(t), t = 0, 1, 2, \dots)$ becomes a Markov chain.

Consider arrivals rates $\lambda \in \mathbb{R}_+^K$. If for this λ the process $Q(\cdot)$ has a stationary distribution, then the time-averaged service rates must exceed arrival rates. So by the ergodic theorem, $\lambda \leq \phi$, i.e., $\lambda_k < \phi_k$ for all $k \in K$, for some $\phi \in \text{Co}(M[K])$. [For more details on this, look [3].] A vector of arrival rates $\lambda \in \mathbb{R}_+^K$ is called feasible if there exists a $\phi \in \text{Co}(M[K])$ such that $\lambda < \phi$.

For a vector v we let v' be its transpose and by e we denote vectors whose coordinates are all equal to one. If $v \in \mathbb{R}^K$ and $L \subset K$ then we write v_L for the vector $(v_i, i \in L) \in \mathbb{R}^L$.

Definition 1. We say that $L \subset K$ satisfies *local pooling* if there exists a nonzero $\alpha \in \mathbb{R}_+^K$ such that $\alpha' \phi$ is a positive constant for all $\phi \in \text{Co}(M[L])$. We say that *local pooling* is satisfied if every $L \subset K$ satisfies local pooling.

For a matrix M , $\text{Co}(M)$ denotes the convex hull of its columns.

Note that if L satisfies local pooling, no vector in $\text{Co}(M[L])$ strictly dominates (in all coordinates) another vector in $\text{Co}(M[L])$. In the sequel we refer to the following remark.

Remark 1. If L satisfies local pooling, then for any λ feasible and $\phi \in \text{Co}(M[L])$, $\lambda_k < \phi_k$ for some $k \in L$.

For fixed feasible λ we define

$$\epsilon_* := \inf \left\{ \min_{k \in L} (\phi_k - \lambda_k) \mid L \subset K, L \text{ satisfies local pooling}, \phi \in \text{Co}(M[L]) \right\}. \quad (8)$$

By Remark 1, $\epsilon_* > 0$. The interpretation of ϵ_* is the excess of the service rate over the arrival rate that must exist at some queue whenever the set of longest queues satisfies local pooling.

The observation above yields an equivalent characterization of local pooling in terms of a linear program.

Proposition 1. *Consider the LP:*

$$\max_{c, \mu, \nu} c \quad (9)$$

$$s.t. \quad M[L]\mu \geq M[L]\nu + ce \quad (10)$$

$$e'\mu = 1 \quad (11)$$

$$e'\nu = 1 \quad (12)$$

$$\mu, \nu \in \mathbb{R}_+^{|L|}, c \in \mathbb{R}. \quad (13)$$

The set L satisfies local pooling if and only if the optimal value is $c^ = 0$.*

Proof. The proof is given in appendix A.

In the proof of Theorem 1 we make use of a large deviation bound on $A(\cdot)$. Henceforth we assume that for each $k \in K$ and $\epsilon > 0$,

$$P \left[\left| \frac{A_k(n)}{n} - \lambda_k \right| > \epsilon \right] \leq \beta \exp(-n\gamma(\epsilon)), \quad \text{for all } n \geq 1, \quad (14)$$

for some $\gamma(\epsilon) > 0$ and $\beta > 0$.

4. Stability

In this section, we first state the main result and we outline its proof. We then establish the main properties we need and we proceed with the formal proof. We conclude the section with a remark about stability when local pooling holds.

4.1. Stability Result

Theorem 1. *Consider a system $M[K]$ such that every $L \subset K$ with $\text{rank } M[L] \geq |L| - 1$ satisfies local pooling. Assume that the arrival processes $A_k(\cdot)$ for each queue $k \in K$ are i.i.d., mutually independent, satisfy (14), and have nonzero variance.*

Then, the system is stable under LQF for all feasible λ .

The condition of the theorem means that if a subset L does not satisfy local pooling, then the rank of the service vectors is too small for the scheduler to be able to keep the queue lengths from diverging.

In applications, it may be useful to infer stability for some given feasible λ rather than for all feasible rates, as in Theorem 1. In these cases, Theorem 1 implies stability under the same conditions on the arrival processes, when queues L that do not satisfy local pooling (having necessarily $\text{rank } M[L] \geq |L| - 1$) need to satisfy the following instead: The optimum c of the linear program

$$\begin{aligned} & \max_{c, \nu} c \\ & \text{s.t. } \lambda_L \geq M[L]\nu + ce \\ & e'\nu = 1 \\ & \nu \in \mathbb{R}_+^{|L|}, c \in \mathbb{R}. \end{aligned}$$

is negative. This guarantees (by Lemma 4 below) that when queues in L are the longest, they must decrease at a certain rate even though local pooling does not hold for L .

4.2. Proof Outline

The strategy of the proof is as follows.

1. We consider $Q^n(nt)/n$ where $Q^n(\cdot)$ is the original system with initial backlogs scaled by n .

2. We show that $E(|Q^n(nt)/n|) \rightarrow 0$ for $t \geq t_0$. This result provides a Lyapunov function for the stability of Q^n . Indeed, the limit implies that the expected queue lengths tend to decrease.
3. To prove the previous result, we use uniform integrability and the fact that $Q^n(nt)/n \rightarrow 0$ for $t \geq t_0$. To get that fact, we prove that $Q^n(nt)/n$ converges along a subsequence to $\bar{Q}(t)$, a fluid system described by a system of differential equations, and that $\bar{Q}(t) = 0$ for $t \geq t_0$, for some t_0 .
4. To prove that $\bar{Q}(t) = 0$ for $t \geq t_0$, we show that, as long as $\bar{Q}(t) > 0$, it must decrease at least at a given rate. Specifically, we show that if $\max_i \bar{Q}_i(t) > 0$ and t is a regular time, then $d \max_i \bar{Q}_i(t)/dt \leq -\epsilon_* < 0$.
5. To derive the rate of decrease, we use the observation that while the set of longest queues is fixed and satisfies local pooling, then the queue lengths decrease at least at a given rate.
6. Finally, we show that, most of the time, a set of queues that satisfies local pooling dominates the other queues. We show that fact by proving that if L does not satisfy local pooling, then a proper subset of L must dominate the other queues most of the time. That argument has two parts:
 - (a) We show that the queues must be separated by some $\beta\Delta|L|$ at most time steps that are multiples of Δ , as we did in the 6-cycle argument.
 - (b) We then use a uniform bound on the arrivals to conclude the domination by a subset of L during the subsequent $\Delta - 1$ steps.

The first four ideas of the proof are identical to [6] and [7]. However, the subsequent steps are quite different since, as we explained in the examples, they involve second-order properties of the processes.

This section is organized as follows.

1. We first define the sequence of systems Q^n in Definition 2 and state the convergence to the fluid limit $\bar{Q}$ in Proposition 2.
2. We prove the separation of the queues at multiples of Δ in Lemma 1.

3. We show that queues remain separated for the subsequent $\Delta - 1$ steps in Lemma 3.
4. We derive the rate of decrease when the longest queues satisfy longest pooling in Lemma 4.
5. We show that the fluid limits of the queue lengths decrease with an upper-bounded rate in Lemma 5.
6. Finally, we conclude the proof of Theorem 1 in Section 4.8.

4.3. Fluid Limit

We define a sequence of systems scaled in space and time by a factor n . The proposition shows that they converge to a fluid limit.

Definition 2. (*Sequence of Systems.*)

On the same probability space we define a sequence of systems with the same $M[K]$, indexed by $n = 1, 2, \dots$. The initial conditions $Q^n(0)$ for the n -th system satisfy $\sum_{i \in K} Q_i^n(0)/n \leq 1$. The arrival process is the same for all elements of the sequence, i.e., $A^n(\cdot) = A(\cdot)$. Let $Q^n(\cdot)$, $T^n(\cdot)$, $D^n(\cdot)$ be the resulting queueing, service and departure processes respectively under LQF. (All systems use the same LQF scheduling policy.)

We use the fluid limits $(\bar{Q}(\cdot), \bar{T}(\cdot), \bar{D}(\cdot))$ to study the stability of the original (pre-limit) processes. The existence and properties of these limits are stated in the following proposition whose proof is similar to [6, 5, 8], so we omit it.

Proposition 2. (Fluid Limit.) *A limit $(\bar{Q}(\cdot), \bar{T}(\cdot), \bar{D}(\cdot))$ of $(Q^n(\cdot), T^n(\cdot), D^n(\cdot))$ as $n \rightarrow \infty$ exists a.s., in the topology of uniform convergence over compact sets, along*

some subsequence, and satisfies the following properties:

$$\begin{aligned}
\bar{Q}(t) &= \bar{Q}(0) + \lambda t - \bar{D}(t), \quad \forall t \geq 0; \\
\sum_{m \in M[K]} \dot{\bar{T}}_m(t) &= 1, \quad \forall t \geq 0; \\
\dot{\bar{D}}(t) &= \sum_{m \in M[K]} m_k \dot{\bar{T}}_m(t), \quad \text{if } \bar{Q}_k(t) > 0; \\
\bar{Q}_k(t) &\geq 0; \\
\bar{D}_k(t), \bar{T}_m(t) &\geq 0, \quad \text{nondecreasing } \forall k \in K, m \in M[K], t \geq 0.
\end{aligned}$$

Moreover, $\bar{Q}(\cdot), \bar{T}(\cdot), \bar{D}(\cdot)$ are absolutely continuous. Times $t \geq 0$ for which the derivatives of $\bar{Q}(t), \bar{T}(t), \bar{D}(t)$ exist will be called regular times.

4.4. Separation of the Queues at Multiples of Δ

Consider any nonempty $L \subset K$ that does not satisfy local pooling.

Define $\Delta(n) = n^{1/6}$. The following lemma, whose proof is in Appendix B, states that the queue lengths remain separated at most of the multiples of $\Delta(n)$.

Lemma 1. (Queue Separation.) *Let nonempty $L \subset K$ not satisfy local pooling. For all $1 \geq \delta > \delta_0 > 0, \alpha > 0$,*

$$\frac{1\{\min_{i \in L} Q_i^n(nt) > n\delta\}}{n^{\frac{1}{2} + \alpha} \Delta(n)} \sum_{m=n\delta_0}^{n\delta} 1\{\max_{i \in L} Q_i^n(nt+m) - \min_{i \in L} Q_i^n(nt+m) \leq \gamma \Delta(n)\} \rightarrow 0 \quad (15)$$

in probability, as $n \rightarrow \infty$ with $\gamma := 1.1 |L| (\max_{i \in K} \lambda_i \vee 1)$.

Therefore, with probability one, for any subsequence (n_k^0) , there is a further subsequence (n_k^1) along which the limit in (15) holds simultaneously for all sets L that do not satisfy local pooling, all $t \in \mathbb{Q}_+$, and all $\delta, \delta_0 \in \mathbb{Q}$ with $1 \geq \delta > \delta_0 > 0$.

4.5. Separation of the Queues during Intervals

To show that a set of queues that satisfies local pooling dominates the other queues most of the time, we use a bound on the fluctuations of the queue lengths. This bound relies on the following property of the arrivals.

Lemma 2. (Property of Arrivals.)

The following event has probability one: For all $T \in \mathbb{Q} \cap (0, +\infty)$,

$$\limsup_n \max_{0 \leq i \leq n^{6/7}T, k \in K} \left| \frac{A_k((i+1)n^{1/7}) - A_k(n^{1/7})}{n^{1/7}} - \lambda_k \right| = 0. \quad (16)$$

Proof. Fix any $T, \epsilon > 0$ and note that

$$\begin{aligned} & \sum_{n=1}^{\infty} P \left[\max_{0 \leq i \leq n^{6/7}T, k \in K} \left| \frac{A_k((i+1)n^{1/7}) - A_k(n^{1/7})}{n^{1/7}} - \lambda_k \right| > \epsilon \right] \\ & \leq \beta T \sum_{n=1}^{\infty} n^{6/7} \exp(-n^{1/7} \gamma(\epsilon)) < \infty. \end{aligned}$$

The result now follows by the Borel-Cantelli lemma.

Definition 3. (*Domination.*)

Let $I \subset \{0, 1, \dots\}$ and $L, L' \subset K$. We say that L is not dominated by L' during I for the system Q^n if there exist $u \in I, i \in L, j \in L'$ such that $Q_i^n(u) \geq Q_j^n(u)$.

Fix any sequence $(n_k^0) \subset \mathbb{Z}_+$ with $n_k^0 \rightarrow \infty$ as $k \rightarrow \infty$. Pick a subsequence (n_k) of (n_k^0) such that the convergence in Lemma 1 is achieved a.s. From now on we fix ω such that that limit is attained and (16) holds.

The following result holds.

Lemma 3. For $n \geq 1$, let $g(n) = n^{\frac{1}{2}+\alpha} \Delta(n)$ for $\alpha > 0$ such that $g(n) \Delta(n) = o(n)$ as $n \rightarrow \infty$. Assume that there exist $\delta \in \mathbb{Q} \cap (0, 1)$, $K' \subset K$ s.t. K' dominates $K \setminus K'$ during $\{n_k t, \dots, n_k(t + \delta)\}$ for all sufficiently large k . Also, assume $Q_i^{n_k}(n_k t) > n_k \delta$ for all $i \in K'$. Fix $\delta_0 \in \mathbb{Q}$ with $0 < \delta_0 < \delta$.

Divide $[n_k t, n_k(t + \delta)]$ into intervals of size $\Delta(n_k)$:

$$I_j^{n_k} = [n_k t + (j-1)\Delta(n_k), n_k t + j\Delta(n_k)]$$

with $j \in \{1, 2, \dots, \delta n_k / \Delta(n_k)\}$.

The number of intervals not dominated by some L that satisfies local pooling for the system Q^{n_k} does not exceed $g(n_k) + \delta_0 n_k / \Delta(n_k)$ as $k \rightarrow \infty$.

Proof. If K satisfies local pooling, then there is nothing to prove. Let nonempty $L \subset K$ not satisfy local pooling. Using Lemma 1, the number of intervals $I_j^{n_k}$ where

$$\max_{i \in L} Q_i^{n_k}(t) - \min_{i \in L} Q_i^{n_k}(t) \leq \gamma \Delta(n_k)$$

at the left point t of the interval is $o(g(n_k))$, and consequently does not exceed $g(n_k) + \delta_0 n_k / \Delta(n_k)$ as $k \rightarrow \infty$.

For k large enough, the arrivals at each queue j in all the intervals are bounded by $1.1\lambda_j$, since ω satisfies (16). Moreover, there is at most one departure from each queue in each time step. It is easy to check that if the queue lengths are separated by more than $\gamma\Delta(n_k)$ at the left point of an interval $I_j^{n_k}$ with size $\Delta(n_k)$, there must be a proper subset L' of L that dominates the other queues during that interval. (This fact is illustrated in the right part of the figure.)

Consequently, the number of intervals such that any set L that does not satisfy local pooling is not dominated by some $L' \subsetneq L$ does not exceed $g(n_k) + \delta_0 n_k / \Delta(n_k)$ as $k \rightarrow \infty$. Since all subsets of L with less than four elements satisfy local pooling, the number of intervals not dominated by some L that satisfies local pooling does not exceed $g(n_k) + \delta_0 n_k / \Delta(n_k)$ in the limit.

4.6. Rate of Decrease under Local Pooling

Our main objective is to prove Lemma 4. We use the local fluid limit technique in [7].

Fix $L \subset K$. Let $(r_k), (n_k) \subset \mathbb{Z}_+$ be sequences for which $r_k/n_k \rightarrow r < \infty$, $Q_i^{n_k}(r_k + u\Delta(n_k)) > Q_j^{n_k}(r_k + u\Delta(n_k))$, for all $u \in [0, 1], i \in L, j \notin L$ and the ordering π on L given by $(Q_i^{n_k}(r_k))_{i \in L}$ kept fixed for all k . Let i_0 be a maximum element of L according to π . We define

$$\hat{Q}_i^{n_k}(u) = \frac{Q_i^{n_k}(r_k + u\Delta(n_k)) - Q_{i_0}^{n_k}(r_k)}{\Delta(n_k)}, u \in [0, 1], i \in L. \quad (17)$$

For $u \in [0, 1]$ we also define $\hat{F}^{n_k}(u) = (F^{n_k}(r_k + u\Delta(n_k)) - F^{n_k}(r_k))/\Delta(n_k)$ for $F^{n_k} \in \{A^{n_k}, D^{n_k}, T^{n_k}\}$. Limits will be taken with respect to the topology of uniform convergence on $[0, 1]$.

Proposition 3. *The limit $\hat{Q}_L^\infty(\cdot)$ as $k \rightarrow \infty$ of (17) exists over some subsequence and satisfies*

$$\hat{Q}_i^\infty(u) = \hat{Q}_i^\infty(0) + \lambda_i u - \hat{D}_i^\infty(u) \quad \forall u \in [0, 1] \quad (18)$$

$$\hat{D}_i^\infty(\cdot) \text{ is nonnegative, nondecreasing, } \hat{D}_i^\infty(0) = 0, \quad (19)$$

$$u^{-1} \hat{D}_L^\infty(u) \in \text{Co}(M[L]), \forall u \in (0, 1]. \quad (20)$$

Proof. Note that $\hat{Q}_i^{n_k}(0) \leq 0$ for all $i \in L$ and large k , so over a subsequence $\hat{Q}^\infty(0)$ exists in $[-\infty, 0]$. Then we approximate $T^{n_k}(\cdot)$ by continuous processes such that the uniform error in $[r_k, r_k + \Delta(n_k)]$ is no more than n_k^{-1} . (For example, on noninteger times approximate by the linear interpolation of the values on the two neighboring integers.) These processes are Lipschitz continuous so, by Arzela-Ascoli, the limit $\hat{T}^\infty(\cdot)$ as $k \rightarrow \infty$ of the original processes exists over some subsequence of (n_k) .

Properties (18),(19) follow by uniform convergence of $\hat{A}^{n_k}(\cdot)$ (which in turn follows by (16)), continuity and (3),(7).

Since queues in L are given higher priority by LQF than those in $K \setminus L$ during $[r_k, r_k + \Delta(n_k)]$, and no queue in L is empty during the same interval, by eqs. (4),(5) $(D_L^{n_k}(r_k + u\Delta(n_k)) - D_L^{n_k}(r_k))/u\Delta(n_k) \in \text{Co}(M[L])$.

Lemma 4. *Fix $L \subset K$. If local pooling holds for L then for any ordering π of elements of L any local limit $\hat{Q}_L^\infty(\cdot)$ of (17) satisfies*

$$\max_{i \in L} \hat{Q}_i^\infty(1) - \max_{i \in L} \hat{Q}_i^\infty(0) \leq -\epsilon_* .$$

Proof. Without loss of generality consider a regular time $t \in (0, 1)$ where $\hat{Q}_i^\infty(t) = \hat{Q}_j^\infty(t) > -\infty$, $\dot{\hat{Q}}_i^\infty(t) = \dot{\hat{Q}}_j^\infty(t)$ for all $i, j \in \arg \max_l \bar{Q}_l =: L$. Fix small enough $\delta > 0$ such that $\hat{Q}_i^\infty(u) > \hat{Q}_j^\infty(u)$ for all $i \in L, j \in K \setminus L, u \in [t, t + \delta]$.

Fix an element of the probability space where the fluid limit exists, by Proposition 3. Let (n_k) be any sequence along which $(\hat{Q}_L^{n_k}(\cdot), \hat{A}_L^{n_k}(\cdot), \hat{T}_L^{n_k}(\cdot), \hat{D}_L^{n_k}(\cdot))$ converges to a fluid limit u.o.c. For sufficiently large k , $\hat{Q}_i^{n_k}(u) > \hat{Q}_j^{n_k}(u)$ for all $i \in L, j \in K \setminus L, u \in [n_k t, n_k(t + \delta)]$. Since queues in L are given higher priority by LQF than those in $K \setminus L$ during $[n_k t, n_k(t + \delta)]$, by eqs. (4),(5)

$$\frac{\hat{D}_L^{n_k}(t + \delta) - \hat{D}_L^{n_k}(t)}{n_k \delta} \in \text{Co}(M[L]) .$$

Letting $k \rightarrow \infty$ and then $\delta \downarrow 0$, $\dot{\hat{D}}_L^\infty(t) \in \text{Co}(M[L])$.

Now since local pooling holds for L and λ is feasible, by Remark 1 there exists $i_0 \in L$ with $\lambda_{i_0} - \dot{\hat{D}}_{i_0}^\infty(t) \leq -\epsilon_* < 0$. But $\dot{\hat{Q}}_i^\infty(t) = \dot{\hat{Q}}_{i_0}^\infty(t)$ for all $i \in L$, so $d \max_j \hat{Q}_j^\infty(t)/dt \leq -\epsilon_* < 0$ for all $t \in (0, 1)$ except some t of Lebesgue measure 0.

4.7. Rate Decrease of Fluid Limit

Pick a subsequence (n_k^1) of (n_k^0) such that the convergence in Lemma 1 is achieved a.s. From now on we fix ω such that that limit holds and such that (16) holds.

Consider a further subsequence (n_k) of (n_k^1) along which the fluid limit $(\bar{Q}(\cdot), \bar{T}(\cdot), \bar{D}(\cdot))$ exists. Since $\bar{Q}(\cdot)$ is absolutely continuous, consider any regular time $t \in \mathbb{Q}_+$ for which $\bar{Q}(t) \neq 0$ and the derivative of $\max_i \bar{Q}_i(\cdot)$ exists.

The following result holds.

Lemma 5. (Decrease Rate of Fluid Limit.) *If $\bar{Q}(t) \neq 0$, one has*

$$\frac{d}{dt} \max_{i \in K} \bar{Q}_i(t) \leq -\epsilon_*$$

where ϵ_* was defined in (8).

Proof. We proceed using contradiction. Assume

$$\frac{d}{dt} \max_{i \in K} \bar{Q}_i(t) > -\epsilon_* .$$

Then for some $\delta \in \mathbb{Q} \cap (0, 1)$ and all large enough k ,

$$\max_{i \in K} Q_i^{n_k}(n_k(t + \delta)) - \max_{i \in K} Q_i^{n_k}(n_k t) > -\delta \epsilon_* n_k . \quad (21)$$

Divide $[n_k t, n_k(t + \delta)]$ into intervals of size $\Delta(n_k)$:

$$I_j^{n_k} = [n_k t + (j - 1)\Delta(n_k), n_k t + j\Delta(n_k)]$$

with $j \in \{1, 2, \dots, \delta n_k / \Delta(n_k)\}$.

Since $\bar{Q}(t) \neq 0$, there exists nonempty $K' \subset K$ s.t. K' dominates $K \setminus K'$ during $\{n_k t, \dots, n_k(t + \delta)\}$, and $Q_i^{n_k}(n_k t) > n_k \delta$ for all $i \in K'$, by an appropriate choice of δ above, and noting that ω satisfies (16) and the departure rate is bounded for each queue. From Lemma 3, the number of intervals not dominated by any subset of queues that satisfies local pooling does not exceed $g(n_k) + \delta_0 n_k / \Delta(n_k)$ as $k \rightarrow \infty$.

Hence by (21), there exists an interval $I_{j_k}^{n_k}$ dominated by some L that satisfies local pooling and

$$\max_{i \in L} Q_i^{n_k}(n_k t + j_k \Delta(n_k)) - \max_{i \in L} Q_i^{n_k}(n_k t + (j_k - 1)\Delta(n_k)) > -\Delta(n_k) \epsilon_* ,$$

for all sufficiently large k . Otherwise for arbitrarily large k ,

$$\begin{aligned} \max_{i \in L} Q_i^{n_k}(n_k(t + \delta)) - \max_{i \in L} Q_i^{n_k}(n_k t) \leq \\ -\epsilon_* \Delta(n_k) \left[\frac{n_k \delta}{\Delta(n_k)} - \left(g(n_k) + \frac{\delta_0 n_k}{\Delta(n_k)} \right) \right] + \left(g(n_k) + \frac{\delta_0 n_k}{\Delta(n_k)} \right) \gamma \Delta(n_k), \end{aligned}$$

for some fixed $\gamma > 0$, which contradicts (21) for small enough δ_0 .

Since for all large k , $I_{j_k}^{n_k}/n_k \subset [t, t + \delta]$, there exist further subsequences (n'_k) of (n_k) , (j'_k) of (j_k) such that: a) $(n'_k t + (j'_k - 1)\Delta(n'_k))/n'_k \rightarrow \infty$, b) the order of elements of L implied by the ordering of coordinates in $Q_L^{n'_k}(n'_k t + (j'_k - 1)\Delta(n'_k))$ is fixed for all sufficiently large k , and c) the limit of $\hat{Q}^{n'_k}(\cdot)$ in $[0, 1]$ as $k \rightarrow \infty$ exists and satisfies

$$\max_{i \in L} \hat{Q}_i^\infty(1) - \max_{i \in L} \hat{Q}_i^\infty(0) > -\epsilon_* . \quad (22)$$

Since L satisfies local pooling then by Lemma 4 we arrive at a contradiction.

4.8. Proof of Theorem 1

Fix any sequence $(n_k^0) \subset \mathbb{Z}_+$ with $n_k^0 \rightarrow \infty$ as $k \rightarrow \infty$.

From Lemma 5, there exists $t_0 \geq 0$ such that $\bar{Q}(t) = 0$ for all $t \geq t_0$, and since (n_k) was an arbitrary subsequence of (n_k^1) so long as it satisfied the prescribed properties, $Q^{n_k^1}(n_k^1 t_0)/n_k^1 \rightarrow 0$ a.s. as $k \rightarrow \infty$. After some uniform integrability analysis for $(Q^{n_k^1}(n_k^1 t_0)/n_k^1)_k$, e.g., cf. [6], we get $E[\|Q^{n_k^1}(n_k^1 t_0)\|_1]/n_k^1 \rightarrow 0$ as $k \rightarrow \infty$. Since (n_k^1) could have been any arbitrary subsequence of (n_k^0) with the prescribed properties, the limit over (n_k^0) is zero as well.

4.9. Stability Under Local Pooling

The same proof can be used (virtually unchanged) to show that a system that satisfies local pooling is stable under i.i.d. (possibly constant) arrivals. For systems $M[K]$ that form trees, in the graph notation of section 2, it is easy to see that local pooling holds: at all times either a leaf or it's parent vertex are served. Hence, in such tree $M[K]$, LQF is stable under any i.i.d. arrivals.

Appendix A. Proof of Proposition 1

Let $M = M[L]$ and assume $c^* = 0$ first. Then for any $M\mu^* \in \text{Co}(M[L])$ (with $e'\mu^* = 1$), $(c, \mu, \nu) = (0, \mu^*, \mu^*)$ must be a solution since (10) is satisfied with equality

for $\nu = \mu$ and $c = 0$. Hence complementary slackness and dual feasibility must hold for some dual variables $\alpha^*, \beta^*, \gamma^*$ corresponding to (10),(11),(12). By complementary slackness,

$$(\alpha^{*'}M - \beta^*e')\mu^* = -(\alpha^{*'}M + \gamma^*e')\mu^* = 0. \quad (23)$$

From (23) and by noting that $(0, \mu^*, \mu^*)$ is feasible we have $\beta^* = -\gamma^*$. Dual feasibility gives $\alpha^* \geq 0$, $\alpha^{*'}e = 1$, and

$$-\gamma^*e \leq \alpha^{*'}M \leq \beta^*e.$$

Since $\beta^* = -\gamma^*$, the last display implies

$$\alpha^{*'}M = \beta^*e. \quad (24)$$

$\alpha^*e = 1$, so $\alpha \geq 0$ is nonzero and $\beta^* > 0$ by nondegeneracy of M . But $M\mu^* \in \text{Co}(M[L])$ was arbitrary so (24) and feasibility imply local pooling holds for L with vector α^* as in Def. 1.

For the converse, it suffices to note that under local pooling for L , no vector can dominate another in $\text{Co}(M[L])$.

Appendix B. Proof of Queue Separation Lemma

Let S_m^n be the partial sums of a triangular array $((X_m^n)_{m=1,\dots,n})_{n \geq 1}$ of i.i.d. r.v.s with nonzero variance and finite third moment. Let (Y_n) be a sequence of r.v.s such that for each n , Y_n is independent of $(S_m^n, m = 1, \dots, n)$.

Lemma 6. *For all $1 \geq \delta > \delta_0 > 0, \alpha > 0$,*

$$\frac{1}{n^{\frac{1}{2}+\alpha}\Delta(n)} \sum_{m=n\delta_0}^{n\delta} 1\{|Y_n + S_m^n| \leq \Delta(n)\} \rightarrow 0 \quad (25)$$

in probability, as $n \rightarrow \infty$.

Proof. Let Z be a standard normal random variable, then by Berry-Esseen (e.g., cf. Theorem 2.4.9 of Durrett [11]),

$$\begin{aligned} \frac{P[|S_m^n - mEX_1^1| \leq \Delta(n)]}{\sup_y P[|y + S_m^n| \leq \Delta(n)]} &\geq \frac{P[|Z| \leq \Delta(n)/\sqrt{m}] - c_1/\sqrt{m}}{P[|Z| \leq \Delta(n)/\sqrt{m}] + c_1/\sqrt{m}} \\ &\geq 1 - \frac{2}{c_2\Delta(m) + 1}, \end{aligned}$$

for some constants $c_1, c_2 > 0$ and large m with $n\delta_0 \leq m \leq n\delta$.

Hence,

$$\begin{aligned} \sum_{m=n\delta_0}^{n\delta} P[|Y_n + S_m^n| \leq \Delta(n)] &\leq c_3 \sum_{m=n\delta_0}^{n\delta} P\left[\left|\frac{S_m^n - mEX_1^1}{\sqrt{m}}\right| \leq \frac{\Delta(n)}{\sqrt{n\delta_0}}\right] \\ &\leq c_3 n P\left[|Z| \leq \frac{\Delta(n)}{\sqrt{n\delta_0}}\right] + nO\left(\frac{1}{\sqrt{n}}\right) \\ &= \Theta(\sqrt{n}\Delta(n)), \quad \text{as } n \rightarrow \infty, \end{aligned}$$

for some constant $c_3 > 0$. The result follows by Markov's inequality.

By the assumption of the theorem, there must exist some nonzero row vector $v^L \in \mathbb{R}^{|L|}$ such that $v^L e = 0, v^L M[L] = 0$, and $\max_{i \in L} |v_i^L| = 1$. We can apply the lemma above to

$$S_m^n = \gamma^{-1} v^L [A_L(nt + m) - A_L(nt)], \quad (26)$$

with $\gamma := 1.1 |L| (\max_{i \in K} \lambda_i \vee 1)$ and

$$Y_n = \gamma^{-1} v^L Q_L^n(nt), \quad (27)$$

for $n = 1, 2, \dots$

Now, provided that $\min_{i \in L} Q_i^n(nt) > n\delta$, so queues in L do not empty during $\{nt, \dots, n(t + \delta)\}$,

$$v^L Q_L^n(nt + m) = v^L Q_L^n(nt) + v^L [A_L(nt + m) - A_L(nt)].$$

We conclude that for all $1 \geq \delta > \delta_0 > 0, \alpha > 0$,

$$\frac{1\{\min_{i \in L} Q_i^n(nt) > n\delta\}}{n^{\frac{1}{2} + \alpha} \Delta(n)} \sum_{m=n\delta_0}^{n\delta} 1\{|v^L Q_L^n(nt + m)| \leq \gamma \Delta(n)\} \rightarrow 0 \quad (28)$$

in probability, as $n \rightarrow \infty$.

Moreover, we can find some constant $\kappa > 0$, that depends on v^L , such that

$$\max_{i \in L} Q_i^n(m) - \min_{i \in L} Q_i^n(m) \geq \kappa |v^L Q_L^n(m)|, m \geq 0.$$

Consequently, for all $1 \geq \delta > \delta_0 > 0, \alpha > 0$,

$$\frac{1}{n^{\frac{1}{2} + \alpha} \Delta(n)} \sum_{m=n\delta_0}^{n\delta} 1\{\max_{i \in L} Q_i^n(nt + m) - \min_{i \in L} Q_i^n(nt + m) \leq \gamma \Delta(n)\} \rightarrow 0 \quad (29)$$

in probability, as $n \rightarrow \infty$.

Appendix C. Extension to general service rates

The model can be generalized to all nonzero nonnegative matrices $M[K]$ in a straightforward manner. We do not consider that level of generality, in this paper, in order to avoid overloading notation.

For nonintegral $M[K]$, one needs to define LQF appropriately since the definition given above depends on the fact that elements of $M[K]$ can take values in $\{0, 1\}$ only. A stationary scheduling policy assigns service vectors, i.e. columns of $M[K]$, to queue lengths, i.e. vectors in $\mathbb{R}_+^K$. LQF will do this by “prioritizing” queues according to their backlog size. The allocation of service rates (more general than 0-1) should be based on these priorities. Such considerations motivate the following definition: For a service matrix $M[K]$ let $\mathcal{S}$ be the set of possible service vectors, i.e., set of column vectors of $M[K]$. A stationary policy $\Phi : \mathbb{R}_+^K \rightarrow \mathcal{S}$ is an LQF policy for $M[K]$, if $\Phi = \Phi_2 \circ \Phi_1$, where

- Φ_1 is a (possibly random) mapping to the space of permutations of K so that Φ_1 reorders the coordinates of its argument in decreasing order. So if $\Phi_1(q) = \pi$, then $q_{\pi(1)} \geq q_{\pi(2)} \geq \dots$
- For a permutation π , $\Phi_2(\pi)$ is obtained as follows. First π reorders the rows of $M[K]$. Let $M^{\pi\cdot}$ be the resulting matrix. The column $j(\pi)$ indexed by Φ_2 is the (unique) element of $\mathcal{S}$ such that the vector $(M_{1j(\pi)}^{\pi\cdot}, M_{2j(\pi)}^{\pi\cdot}, \dots, M_{kj(\pi)}^{\pi\cdot})$ is maximal across all columns, for all k with $1 \leq k \leq K$.

Our results carry over to this more general rule (although not included in the proofs) because the essential property of LQF we make use of is that shorter queues do not affect longer ones, i.e., in each time slot the service rate of a (nonzero) queue does not depend on the queue length of any shorter queues.

References

- [1] STOLYAR, A. (2004). MaxWeight scheduling in a generalized switch: State space collapse and workload minimization in heavy traffic. *Ann. Appl. Probab.*, **14**, 1–53.
- [2] TASSIULAS, L. AND EPHREMIDES, A. (1992). Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop

- radio networks. *IEEE Trans. Automat. Control*, **47**, 1936–1948.
- [3] ANDREWS, M., KUMARAN, K., RAMANAN, K., STOLYAR, A., VIJAYAKUMAR, R. AND WHITING, P. (2004). Scheduling in a queueing system with asynchronously varying service rates. *Probab. Engrg. Inform. Sci.*, **18**, 191–217.
- [4] MCKEOWN, N. (1995). Scheduling Algorithms for Input-Queued Cell Switches, PhD Thesis, Univ. of Cal. at Berkeley.
- [5] DAI, J. AND PRABHAKAR, B. (2000). The throughput of data switches with and without speedup. In *Proc. of IEEE INFOCOM 2000*. Spring 556–564. IEEE. Tel Aviv, Israel.
- [6] DAI, J. (1995). On positive Harris recurrence of multiclass queueing networks: a unified approach via fluid limit models. *Ann. Appl. Probab.*, **5**, 49–77.
- [7] SHAKKOTTAI, S. AND STOLYAR, A. (2002). Scheduling for multiple flows sharing a time-varying channel: The exponential rule. In *Analytic Methods in Applied Probability*. 185–202. AMS, Providence, RI.
- [8] STOLYAR, A. (1995). On the Stability of Multiclass Queueing Networks: a Relaxed Sufficient Condition via Limiting Fluid Processes. *Markov Process. Related Fields*, **1**, 491–512.
- [9] MALYSHEV, V.A. (1993). Networks and dynamical systems. *Adv. Appl. Prob.*, **25**, 140–175.
- [10] KUMAR, S. GIACCONE, P. AND LEONARDI, E. (2002). Rate Stability of Stable Marriage Scheduling Algorithms in Input-Queued Switches. In *Proc. of 40th Annual Allerton Conference*.
- [11] DURRETT, R. (1995). *Probability: Theory and Examples*. 3rd Edition, Duxbury Press.