

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

ANTONIS DIMAKIS * ** AND

JEAN WALRAND,* *University of California, Berkeley*

Abstract

We consider the stability of the *longest-queue-first* scheduling policy (LQF), a natural and low-complexity scheduling policy, for a *generalized switch* model. Unlike that of 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 independent, identically distributed 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 scheduling; MaxWeight scheduling; throughput optimality; stability; local pooling; fluid limit; local fluid limit; second-order condition

2000 Mathematics Subject Classification: Primary 60K25

Secondary 90B15; 60G17

1. Introduction

We consider the scheduling of the *generalized switch* model described in [9], which includes as special cases models of multi-user data scheduling over a wireless medium, input-queued crossbar switches, and parallel server systems. The problem of throughput optimal scheduling was addressed in [9], where *MaxWeight* scheduling was proposed. This algorithm is based on that in [11] originally proposed in the context of data scheduling in a multihop wireless network. A natural low-complexity alternative to *MaxWeight* is *longest-queue-first* scheduling (LQF), proposed in [7]. 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 the 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 systems, based on the strong law of large numbers, has been formalized in [6], [2], and [10] and applied with great success in multiclass queueing networks [2] and generalized switch models [1], [9], [8], [3]. The premise is that in many cases it is easier to work with a deterministic system than with a stochastic one.

Received 14 June 2005; revision received 3 February 2006.

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

**Email address: dimakis@eecs.berkeley.edu

In LQF, however, 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 for us to infer stability, since deterministic and nondeterministic arrivals with the same intensities can lead to unstable and stable systems, respectively.

Previous work on LQF [5] 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 system. We show that in certain special cases, namely systems satisfying the rank condition of Theorem 1, LQF is throughput optimal under the assumption of nondeterministic arrivals. Although the systems satisfying the conditions in Theorem 1 are special, the authors 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 [2] and [9].

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 our main result, 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 sets, respectively. One can think of the vertices as queues in which arrivals occur in discrete time at some given average rate. We consider independent, identically distributed (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 that 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 decreasing order of queue length. Any queue with a neighbor that has already been selected is 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 repeated 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 scheduling policy (or 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 on average the longest queue decreases. This is immediate if one queue, say queue i , is the unique longest queue for some interval of time. Indeed, in that case, queue i is constantly served under LQF during that interval, and its length tends to decrease because $\lambda_i < 1$. Now, it may happen that

a subset, L , of queues are variously the longest during some interval of time. For instance, assume that $L = \{1, 2\}$, meaning that during the interval the scheduler selects queue 2 some fraction of the time and queues $\{1, 3\}$ the rest of the time. During that interval of time, the scheduler always serves either queue 1 or queue 2, so the total service rate of 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 both queue 1 and queue 2) decreases on average. The same argument can be made for any subset $L \subset \{1, 2, 3\}$ (we use ‘ \subset ’ in the sense of nonstrict inclusion). The key property is as follows. For any set, L , of queues that are variously 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. 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 length of 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 $L = \{1, 2, 3\}$ we can choose $F = \{1, 2\}$, if $L = \{1, 3\}$ we can choose $F = \{1\}$, and so on.

This behavior 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 (see Subsection 4.9).

To summarize, 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 in length during that interval. Consequently, if all sets satisfy local pooling then the system is stable. Thus, the stability conditions for systems that satisfy local pooling are only formulated in terms of the 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 six-vertex cycle graph with $V = \{1, \dots, 6\}$ and $E = \{\{1, 2\}, \{2, 3\}, \dots, \{6, 1\}\}$ depicted in Figure 1. It can be checked 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 zero or one queue from the subset $\{1, 2\}$; serves either one or two queues from $\{1, 2, 3\}$ and either one or two 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 of work, w , 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 $w > \frac{4}{9}$. What makes this fact interesting is that, as we show below, the same system with nonconstant i.i.d. arrivals with average value w is stable whenever $w < \frac{1}{2}$.

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 that 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 ‘small matches’. There are a number of possible sequences of choices that the LQF scheduler can make. These sequences correspond to choosing a small or big match at each step, whenever the choice is possible. The right-hand part of Figure 1 shows these sequences. Each rectangle




FIGURE 1: *Left*: a six-cycle system. Neighboring queues cannot be served at the same time. *Middle*: the separation argument. *Right*: possible schedules.

represents a representative set of possible longest queues, up to a symmetric renumbering. For instance, the queues are initially all equal. First the scheduler chooses a small match, represented by S , with probability $\frac{1}{3}$; then selects a big match, B , with probability $\frac{1}{2}$; then is forced to pick a big match again; and so on. By analyzing this Markov chain describing the choices of the scheduler, we find that the queues become equal again after a random number of steps. By 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 $\frac{4}{9}$. It follows that the system is unstable if $w > \frac{4}{9}$.

We now explain why this system is stable when $w < \frac{1}{2}$ if the arrivals are allowed 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 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.

Denote by $Q_i(t)$ the length of queue i at time t . The key idea is that the longest queue length tends to decrease whenever it is large. Specifically, for scale parameter $n \gg 1$, there exist some t_0 and some ε_* such that

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

It follows that $\max_i Q_i(nt)$ is a Lyapunov function for the Markov chain, which implies its positive recurrence. To demonstrate (1), we consider a sequence of processes Q^n that behave like the original system, but with initial conditions bounded by $O(n)$. We show 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, we prove that

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

This fact, together with the convergence, implies that

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

Using a uniform integrability argument, we conclude that

$$E \left[\max_i Q_i(nt)/n \right] \rightarrow 0 \quad \text{for} \quad t \geq t_0,$$

which implies (1). This line of argument is standard. The novel contribution of this paper is in the argument demonstrating (2). The main step is to show that if $\max_i \bar{Q}_i(t) > 0$, then

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

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

In the six-cycle example, the only set of queues that does not satisfy the local pooling condition is the set of all queues. Thus, assume 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.

We then show 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 this condition holds for some time $s = k\Delta(n) \in I$, then there is a proper subset, L , of queues that remain strictly longer than the others during the interval

$$I(s) := \{s, s + 1, \dots, s + \Delta(n) - 1\}.$$

(See Figure 1.) 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

$$\frac{1}{n} \mathbb{E} \left[\max_i Q_i(nt + n\delta) - \max_i Q_i(nt) \mid Q(nt) \right] \leq -\delta\varepsilon_* < 0.$$

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 bounded below 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, then the number of queues served in $\{2, 3\}$ is always the same as the number served in $\{5, 6\}$, under any service set M_i . If no queue empties, since the departures cancel out, we have

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

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

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

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

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

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

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 $s \in I$ 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 in length. 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 for the fluctuations in the arrivals because the set of service vectors only has rank four in the six-dimensional space of queue lengths. In summary, 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 is sufficient, as will become evident in the proof of Lemma 1 in Appendix B).

These ideas form the basis of Theorem 1, in which 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 that the system may be unstable, even for Bernoulli arrivals. One such example is a system specified by an eight-vertex cyclic graph driven by Bernoulli arrivals at rate $\lambda_i = 0.4984$ for $i = 1, \dots, 8$; ties are broken as in the six-vertex example.

Sufficient stability conditions for LQF in cyclic graphs were considered in [5]. There the analysis was based on the stability of the associated fluid system. It was found that if the total arrival rate at the cycle of three consecutive queues is strictly less than 1, then the system is stable. Intuitively, if all queues are variously the longest for some considerable amount of time, then the lengths of the three queues satisfying the condition above must decrease, because at least one of them is served at each time slot. On the other hand, under nondeterministic arrivals, this condition on the rates is not necessary for the six-cycle described above. For systems not of the same special form as the six-cycle, i.e. those satisfying the rank condition in Theorem 1, it is still possible to achieve stability, by imposing stricter conditions on the arrival rates. (See the 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 $k \in K$ at time slot $t \geq 0$, and let $A_k(t)$ and $D_k(t)$ respectively denote the cumulative arrivals and departures of work at queue k up to time t . We assume that arrivals are i.i.d. processes with $\mathbb{E}[A_k(1)] = \lambda_k < \infty$ for every k . Additionally, the processes $A_k(\cdot)$ are independent for each k .

The scheduler cannot serve all nonempty queues at any 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 $m_k \in \{0, 1\}$ units of work from each queue k , unless queue k empties first. (We relax this ‘ $\{0, 1\}$ -assumption’ in Appendix C.)

By $M[K]$ we denote the maximal elements of $S[K]$, and treat it as a matrix whose columns are its elements. For $L \subset K$, we use $S[L]$ to denote the restriction of $S[K]$ to L and $M[L]$ to denote the maximal elements of $S[L]$. This model is a special case of the generalized switch of [9] and is similar to the model of the input-buffered packet switch of [3].

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 , during which the scheduling policy chooses m .

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

$$Q(t) = Q(0) + A(t) - D(t), \tag{3}$$

$$D_k(t + 1) - D_k(t) = \begin{cases} \sum_{m \in M[K]} m_k(T_m(t + 1) - T_m(t)) & \text{if } Q_k(t) > 0, \\ 0 & \text{otherwise,} \end{cases} \tag{4}$$

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

$$Q_k(t) \geq 0 \quad \text{for all } k \in K \text{ and } t \geq 0,$$

$$D_k(t) \text{ and } T_m(t) \text{ are nonnegative and nondecreasing for all } k \in K, m \in M[K], \text{ and } t \geq 0. \tag{6}$$

We consider all the above processes to be defined for noninteger times as well: for a noninteger u , $D(u) = D(\lfloor u \rfloor)$, and the other processes are extended similarly (here $\lfloor u \rfloor$ denotes the greatest integer not exceeding u).

LQF is a stationary policy that at any time t selects a vector $n \in M[K]$ as follows. First, the queues are sorted in decreasing order according to their backlogs. Given such an ordering $k(1), \dots, k(|K|)$ (where $Q_{k(1)}(t) \geq \dots \geq Q_{k(|K|)}(t)$), n is constructed iteratively by considering the ‘longest queues first’. In the first iteration, we start with $M_1 := \{m \in M[K] : m_{k(1)} = 1\}$, i.e. all service vectors activating the longest queue. If $M[K]$ is nondegenerate then $M_1 \neq \emptyset$. Moreover, we set $n_{k(1)} = 1$. In the i th iteration (for $i > 1$), the search is confined to vectors $M_i \subset M_{i-1}$ that activate queue $k(i)$ if possible; i.e. $M_i := \{m \in M_{i-1} : m_{k(i)} = 1\}$ when $M_i \neq \emptyset$ (in which case we set $n_{k(i)} = 1$). If no such vectors exist (i.e. $m_{k(i)} = 0$ for all $m \in M_{i-1}$), then we set $n_{k(i)} = 0$, and $M_i := M_{i-1}$. Having fixed $n_{k(i)}$, we continue with the $(i + 1)$ th iteration until we have enumerated all queues in K and chosen all components of the service vector n . To completely describe such a policy we must 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 a vector of arrival 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 the arrival rates. Thus, 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, see [1]. For a matrix M , $\text{Co}(M)$ denotes the convex hull of its columns.) The vector $\lambda \in \mathbb{R}_+^K$ is called feasible if there exists a $\phi \in \text{Co}(M[K])$ such that $\lambda < \phi$.

For a vector v we let v^\top denote its transpose, and by e denote the vector whose coordinates all equal 1. If $v \in \mathbb{R}^K$ and $L \subset K$ then we write v_L for the vector $(v_i, i \in L) \in \mathbb{R}^L$. We use \mathbb{Q} to denote the rationals and \mathbb{Q}_+ to denote the nonnegative rationals. Finally, we let $a \vee b = \max(a, b)$.

Definition 1. We say that $L \subset K$ satisfies local pooling if there exists a nonzero vector $\alpha \in \mathbb{R}_+^K$ such that $\alpha^\top \phi$ 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.

Note that if L satisfies local pooling, no vector in $\text{Co}(M[L])$ strictly dominates (in all coordinates) any other vector in $\text{Co}(M[L])$.

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

For a fixed, feasible λ we define

$$\varepsilon_* := \inf \left\{ \max_{k \in L} (\phi_k - \lambda_k) : L \subset K, L \text{ satisfies local pooling, } \phi \in \text{Co}(M[L]) \right\}. \tag{7}$$

By Remark 1, $\varepsilon_* > 0$. We interpret ε_* as 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 following linear program, where r is the number of columns of $M[L]$:

$$\begin{aligned} & \text{find } \max_{c, \mu, v} c \text{ such that} \\ & M[L]\mu \geq M[L]v + ce, \end{aligned} \tag{8}$$

$$e^\top \mu = 1, \tag{9}$$

$$e^\top v = 1, \tag{10}$$

$$\mu, v \in \mathbb{R}_+^r, \quad c \in \mathbb{R}.$$

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 $\varepsilon > 0$,

$$\mathbb{P} \left(\left| \frac{A_k(n)}{n} - \lambda_k \right| > \varepsilon \right) \leq \beta \exp(-n\gamma(\varepsilon)) \quad \text{for all } n \geq 1, \tag{11}$$

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

4. Stability

In this section, we first state the main result and outline its proof. We then establish the main properties we need and 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 independent arrival processes $A_k(\cdot)$, one for each queue $k \in K$, are i.i.d., satisfy (11), and have nonzero variance. The system is then stable under LQF for all feasible vectors λ .

According to the condition of the theorem, 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 prevent 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 and do *not* have rank $M[L] \geq |L| - 1$ are instead such that the optimum of following the linear program is negative, where r is again the number of columns of $M[L]$:

$$\text{find } \max_{c, v} c \quad \text{such that } \lambda_L \geq M[L]v + ce \text{ and } e^\top v = 1 \text{ for } v \in \mathbb{R}_+^r \text{ and } c \in \mathbb{R}.$$

This guarantees (according to Lemma 4, below) that when queues in L are the longest, they must decrease in length 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 vector of queue lengths in a system with initial backlogs not exceeding n , i.e. with $\sum_k Q_k^n(0) \leq n$.
2. We show that $E[\sum_k Q_k^n(nt)/n] \rightarrow 0$ 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 prove this claim, we show 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 $\max_i \bar{Q}_i(t)$ is positive 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 $d \max_i \bar{Q}_i(t)/dt \leq -\varepsilon_* < 0$.
5. To derive the rate of decrease, we use the observation that while the set of longest queues is fixed and satisfies local pooling, 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 demonstrate this 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, which argument has two parts:
 - (a) We show that the queues must differ in length by some amount $\beta \Delta(n)|L|$ at most time steps that are multiples of $\Delta(n)$, as we did in the six-cycle argument.
 - (b) We then use a uniform bound on the arrivals to conclude the domination by a subset of L during the subsequent $\Delta(n) - 1$ steps.

The first four ideas of the proof are identical to the general programme followed in [2] and [8]. However, the subsequent steps are quite different since, as explained in the examples, they involve second-order properties of the processes.

This section is organized as follows.

- We define the sequence of systems Q^n in Definition 2, and state the convergence to the fluid limit \bar{Q} in Proposition 2.
- We demonstrate the required queue length differences at times that are multiples of $\Delta(n)$ in Lemma 1.

- We show that queues remain separated (i.e. appropriately different in length) for the subsequent $\Delta(n) - 1$ steps in Lemma 3.
- We derive the rate of decrease when the local queues satisfy local pooling in Lemma 4.
- In Lemma 5 we show that the fluid limits of the queue lengths decrease with a rate that is bounded from above.
- 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 . Proposition 2 shows that they converge to a fluid limit.

Definition 2. (*Sequence of scaled systems.*) On the same probability space implicitly used above, we define a sequence of systems with the same $M[K]$ and indexed by $n = 1, 2, \dots$. The initial values, $Q^n(0)$, 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 $Q^n(\cdot)$, $T^n(\cdot)$, and $D^n(\cdot)$ respectively be the resulting queueing, service, and departure processes under LQF. (All systems use the same LQF scheduling policy.)

We use the fluid limits (respectively $\bar{Q}(\cdot)$, $\bar{T}(\cdot)$, and $\bar{D}(\cdot)$) to study the stability of the original (prelimit) processes. The existence and properties of these limits are stated in the following proposition, whose proof, as it is similar to that of Theorem 4.1 of [2] and that of Theorem 3 of [3], is omitted.

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 $n \rightarrow \infty$ almost surely exists, in the topology of uniform convergence over compact sets, along some subsequence, and satisfies the following properties (where d_t denotes ordinary differentiation with respect to time):*

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

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

4.4. Separation of the queues at multiples of $\Delta(n)$

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

Lemma 1. (Queue separation.) Consider any nonempty set $L \subset K$ that does not satisfy local pooling. For all δ and δ_0 , $1 \geq \delta > \delta_0 > 0$, and $\alpha > 0$, we have

$$\frac{\mathbf{1}\{\min_{i \in L} Q_i^n(nt) > n\delta\}}{n^{1/2+\alpha} \Delta(n)} \sum_{m=n\delta_0}^{n\delta} \mathbf{1}\left\{\max_{i \in L} Q_i^n(nt+m) - \min_{i \in L} Q_i^n(nt+m) \leq \gamma \Delta(n)\right\} \rightarrow 0 \tag{12}$$

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

Therefore, with probability 1, for any subsequence (n_k^0) there is another subsequence, (n_k^1) , along which the convergence in (12) holds simultaneously for all sets L that do not satisfy local pooling, for all $t \in \mathbb{Q}_+$, and for 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 1:

$$\left\{ \limsup_n \max_{0 \leq i \leq n^{6/7}T, k \in K} \left| \frac{A_k((i+1)n^{1/7}) - A_k(n^{1/7})}{n^{1/7}} - \lambda_k \right| = 0 \text{ for all } T \in \mathbb{Q} \cap (0, \infty) \right\}. \tag{13}$$

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

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

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 $u \in I$, $i \in L$, and $j \in L'$ 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 holds almost surely. From now on we fix a sample point, ω , of the underlying probability space such that that limit is attained and (13) holds.

We have the following result.

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

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

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

with $j \in \{1, 2, \dots, \delta n_k / \Delta(n_k)\}$. The number of intervals not dominated by some set 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, so consider a nonempty set $L \subset K$ that does not satisfy local pooling. From Lemma 1, the number of intervals $I_j^{n_k}$ in which

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

at the left-hand endpoint, 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 sufficiently large k , the arrivals at each queue j in all the intervals are bounded by $1.1 \lambda_j$, since ω satisfies (13). 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-hand endpoint of an interval of size $\Delta(n_k)$, then there must be a proper subset, L' , of L that dominates the other queues during that interval. (This fact is illustrated in the middle part of Figure 1.)

Consequently, the number of intervals in which 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 fewer than four elements satisfy local pooling, the number of intervals not dominated by some L that does satisfy 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

Here our main objective is to prove Lemma 4, which concerns the rate of decrease of a local fluid limit. The local fluid limit idea comes from [8], and we will use it to prove Lemma 5.

Fix $L \subset K$. Assume that $(r_k), (n_k) \subset \mathbb{Z}_+$ are sequences for which $r_k/n_k \rightarrow r < \infty$ and

$$Q_i^{n_k}(r_k + u \Delta(n_k)) > Q_j^{n_k}(r_k + u \Delta(n_k))$$

for all $u \in [0, 1]$, $i \in L$, and $j \notin L$, and that the ordering π in L given by $(Q_i^{n_k}(r_k))_{i \in L}$ is kept fixed for all k . Let i_0 be a maximum element of L according to π . We define

$$\hat{Q}_i^{n_k}(u) = \frac{Q_i^{n_k}(r_k + u \Delta(n_k)) - Q_{i_0}^{n_k}(r_k)}{\Delta(n_k)}, \quad u \in [0, 1], i \in L. \tag{14}$$

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, $(\hat{Q}_L^\infty(\cdot), \hat{A}_L^\infty(\cdot), \hat{D}_L^\infty(\cdot), \hat{T}^\infty(\cdot))$, of $(\hat{Q}_L^{n_k}(\cdot), \hat{A}_L^{n_k}(\cdot), \hat{D}_L^{n_k}(\cdot), \hat{T}^{n_k}(\cdot))$ as $k \rightarrow \infty$ exists over some subsequence and satisfies*

$$\hat{Q}_i^\infty(u) = \hat{Q}_i^\infty(0) + \lambda_i u - \hat{D}_i^\infty(u) \quad \text{for all } u \in [0, 1], \tag{15}$$

$$\hat{D}_i^\infty(\cdot) \text{ is nonnegative and nondecreasing,} \quad \hat{D}^\infty(0) = 0, \tag{16}$$

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

Proof. Note that $\hat{Q}_i^{n_k}(0) \leq 0$ for all $i \in L$ and large k , so $\hat{Q}^\infty(0)$ exists in $[-\infty, 0]$ over a subsequence. 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, at noninteger times approximate by the linear interpolation of the values at the two neighboring integer times.) These processes are Lipschitz continuous so, by the Arzela-Ascoli theorem, the limit, $\hat{T}^\infty(\cdot)$, of $\hat{T}^{n_k}(\cdot)$ as $k \rightarrow \infty$ exists over some subsequence of (n_k) .

Properties (15) and (16) follow by uniform convergence of $\hat{A}^{n_k}(\cdot)$ (which in turn follows from (13)), continuity, and (3) and (6).

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, from (4) and (5) we obtain

$$\frac{D_L^{n_k}(r_k + u\Delta(n_k)) - D_L^{n_k}(r_k)}{u\Delta(n_k)} \in \text{Co}(M[L]).$$

Lemma 4. *If local pooling holds for L then, for any ordering, π , of the elements of L , any local limit $\hat{Q}_L^\infty(\cdot)$ of $\hat{Q}_i^{n_k}(\cdot)$ satisfies*

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

Proof. Without loss of generality, consider a regular time $t \in (0, 1)$ at which $\hat{Q}_i^\infty(t) = \hat{Q}_j^\infty(t) > -\infty$ and $d_t \hat{Q}_i^\infty(t) = d_t \hat{Q}_j^\infty(t)$ for all $i, j \in \text{argmax}_{l \in L} \hat{Q}_l(t) =: \hat{L}$. Choose a $\delta > 0$ sufficiently small that $\hat{Q}_i^\infty(u) > \hat{Q}_j^\infty(u)$ for all $i \in \hat{L}, j \in L \setminus \hat{L}$, and $u \in [t, t + \delta]$.

Fix an element of the probability space where, by Proposition 3, the fluid limit exists. 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 uniformly on compact sets. For sufficiently large k , $\hat{Q}_i^{n_k}(u) > \hat{Q}_j^{n_k}(u)$ for all $i \in \hat{L}, j \in L \setminus \hat{L}$, and $u \in [n_k t, n_k(t + \delta)]$. Since queues in \hat{L} are given higher priority by LQF than are those in $K \setminus \hat{L}$ during $[n_k t, n_k(t + \delta)]$, from (4) and (5) we obtain

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

By letting $k \rightarrow \infty$ and then $\delta \downarrow 0$, we see that $d_t \hat{D}_{\hat{L}}^\infty(t) \in \text{Co}(M[\hat{L}])$.

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

4.7. Rate of decrease of fluid limit

Pick a subsequence, (n_k^1) , of (n_k^0) such that the convergence in Lemma 1 holds almost surely. From now on we fix ω such that that convergence holds and such that (13) holds. Now consider a 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. (Rate of decrease of fluid limit.) *If $\bar{Q}(t) \neq 0$ then*

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

where ε_* is as defined in (7).

Proof. We proceed using contradiction. Assume that

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

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

$$\max_{i \in K} Q_i^{n_k}(n_k(t + \delta)) - \max_{i \in K} Q_i^{n_k}(n_k t) > -\delta \varepsilon_* n_k. \tag{17}$$

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

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

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

Since $\bar{Q}(t) \neq 0$, there exists a nonempty set $K' \subset K$ such that 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 , noting that ω satisfies (13) and that 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 (17), there exists an interval $I_{j_k}^{n_k}$ dominated by some L that satisfies local pooling, and

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

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

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

for some fixed $\gamma > 0$, which contradicts (17) for sufficiently small δ_0 .

Since $I_{j_k}^{n_k} \subset [n_k t, n_k(t + \delta)]$ for all large k , there exist further subsequences (n'_k) of (n_k) and (j'_k) of (j_k) such that $(n'_k t + (j'_k - 1)\Delta(n'_k)) / n'_k \rightarrow \infty$, 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 the limit of $\hat{Q}^{n'_k}(\cdot)$ in $[0, 1]$ as $k \rightarrow \infty$ exists and satisfies

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

Since L satisfies local pooling, by Lemma 4 we arrive at a contradiction, proving Lemma 5.

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 and the condition on $\bar{Q}(0)$ in Proposition 2, there exists a $t_0 \geq 0$ such that $\bar{Q}(t) = 0$ for all $t \geq t_0$, and since (n_k) is an arbitrary subsequence of (n_k^1) so long as it satisfies the prescribed properties, $Q^{n_k^1}(n_k^1 t_0) / n_k^1 \rightarrow 0$ almost surely as $k \rightarrow \infty$. After some uniform integrability analysis for $(Q^{n_k^1}(n_k^1 t_0) / n_k^1)_k$ (see, e.g. [2]), we find that

$$\frac{1}{n_k^1} \mathbb{E} \left[\sum_{i \in K} Q_i^{n_k^1}(n_k^1 t_0) \right] \rightarrow 0 \quad \text{as } k \rightarrow \infty.$$

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 also 0.

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 its parent vertex is served.

Appendix A. Proof of Proposition 1

Let $M = M[L]$ and first assume that $c^* = 0$. Then, for any $M\mu^* \in \text{Co}(M[L])$ (with $e^\top \mu^* = 1$), $(c, \mu, \nu) = (0, \mu^*, \mu^*)$ must be a solution to the linear program of Proposition 1, since (8) is satisfied with equality for $\nu = \mu$ and $c = 0$. Hence, complementary slackness and dual feasibility must hold for some dual variables, α^* , β^* , and γ^* , corresponding to (8), (9), and (10). By complementary slackness, we have

$$(\alpha^{*\top} M - \beta^* e^\top) \mu^* = -(\alpha^{*\top} M + \gamma^* e^\top) \mu^* = 0. \tag{18}$$

From (18) and by noting that $(0, \mu^*, \mu^*)$ is feasible, we have $\beta^* = -\gamma^*$. Dual feasibility yields $\alpha^* \geq 0$, $\alpha^{*\top} e = 1$, and

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

Since $\beta^* = -\gamma^*$, the last display implies that $\alpha^{*\top} M = \beta^* e^\top$. Since $\alpha^{*\top} e = 1$, $\alpha^* \geq 0$ is nonzero, and $\beta^* > 0$ by the nondegeneracy of M . This implies that local pooling holds for L , according to Definition 1 with α^* as the required nonzero vector.

To prove 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 the 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. random variables with nonzero variance and finite third moment. Let (Y_n) be a sequence of random variables such that for each n , Y_n is independent of $(S_m^n, m = 1, \dots, n)$.

Lemma 6. *For all δ and δ_0 , $1 \geq \delta > \delta_0 > 0$, and $\alpha > 0$, we have*

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

in probability as $n \rightarrow \infty$.

Proof. Let Z be a standard normal random variable. Then, by the Berry–Esseen theorem (see, e.g. Theorem 2.4.9 of Durrett [4, p. 124]), for all sufficiently large n we have

$$\begin{aligned} \frac{\mathbb{P}(|S_m^n - m \mathbb{E}[X_1^n]| \leq \Delta(n))}{\sup_y \mathbb{P}(|y + S_m^n| \leq \Delta(n))} &\geq \frac{\mathbb{P}(|Z| \leq \Delta(n)/\sqrt{m}) - c_1/\sqrt{m}}{\mathbb{P}(|Z| \leq \Delta(n)/\sqrt{m}) + c_1/\sqrt{m}} \\ &\geq 1 - \frac{2}{c_2 \Delta(m) + 1}, \end{aligned}$$

for some constants $c_1, c_2 > 0$ and an m such that $n\delta_0 \leq m \leq n\delta$.

Hence,

$$\begin{aligned} \sum_{m=n\delta_0}^{n\delta} \mathbb{P}(|Y_n + S_m^n| \leq \Delta(n)) &\leq c_3 \sum_{m=n\delta_0}^{n\delta} \mathbb{P}\left(\left|\frac{S_m^n - m \mathbb{E}[X_1^1]}{\sqrt{m}}\right| \leq \frac{\Delta(n)}{\sqrt{n\delta_0}}\right) \\ &\leq c_3 n \mathbb{P}\left(|Z| \leq \frac{\Delta(n)}{\sqrt{n\delta_0}}\right) + nO\left(\frac{1}{\sqrt{n}}\right) \\ &= O(\sqrt{n}\Delta(n)) \quad \text{as } n \rightarrow \infty, \end{aligned}$$

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

By the rank assumption of Theorem 1, 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 $\max_{i \in L} |v_i^L| = 1$. Now, provided that $\min_{i \in L} Q_i^n(nt) > n\delta$, queues in L do not empty during $\{nt, \dots, n(t + \delta)\}$; thus,

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

We can apply the lemma above to

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

with $\gamma := 1.1|L|(\max_{i \in K} \lambda_i \vee 1)$, and $Y_n = \gamma^{-1} v^L Q_L^n(nt)$, for $n = 1, 2, \dots$, to conclude that for all δ and δ_0 , $1 \geq \delta > \delta_0 > 0$, and $\alpha > 0$,

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

in probability as $n \rightarrow \infty$.

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

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

Consequently, for all δ and δ_0 , $1 \geq \delta > \delta_0 > 0$, and $\alpha > 0$,

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

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 overcomplicating the notation.

For nonintegral $M[K]$, we must redefine 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 or 1) should be based on these priorities. Such considerations motivate the following definition. For a service matrix $M[K]$, let \mathcal{A} be the set of possible

service vectors, i.e. columns of $M[K]$. A stationary policy $\Phi: \mathbb{R}_+^K \rightarrow \mathcal{S}$ is an LQF policy for $M[K]$ if $\Phi = \Phi_2 \circ \Phi_1$. Here Φ_1 is a (possibly random) mapping to the space of permutations of K that reorders the coordinates of its argument in decreasing order: 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 policy (although it is not included in the proofs), because the essential property of LQF we make use of is that shorter queues do not affect longer ones; in each time slot the service rate of a (nonzero) queue does not depend on the queue length of any shorter queues.

References

- [1] ANDREWS, M. *et al.* (2004). Scheduling in a queuing system with asynchronously varying service rates. *Prob. Eng. Inf. Sci.* **18**, 191–217.
- [2] DAI, J. G. (1995). On positive Harris recurrence of multiclass queueing networks: a unified approach via fluid limit models. *Ann. Appl. Prob.* **5**, 49–77.
- [3] DAI, J. G. AND PRABHAKAR, B. (2000). The throughput of data switches with and without speedup. In *Proc. INFOCOM 2000*, Vol. 2, IEEE, Piscataway, NJ, pp. 556–564.
- [4] DURRETT, R. (2004). *Probability: Theory and Examples*, 3rd edn. Duxbury, Belmont, CA.
- [5] KUMAR, S., GIACCONE, P. AND LEONARDI, E. (2002). Rate stability of stable marriage scheduling algorithms in input-queued switches. In *Proc. 40th Annual Allerton Conference on Computers, Communication, and Control*, University of Illinois, Urbana-Champaign, IL. Available at <http://www.stanford.edu/~skumar/preprints.htm>.
- [6] MALYSHEV, V. A. (1993). Networks and dynamical systems. *Adv. Appl. Prob.* **25**, 140–175.
- [7] MCKEOWN, N. (1995). Scheduling algorithms for input-queued cell switches. Doctoral Thesis, University of California, Berkeley.
- [8] SHAKKOTTAI, S. AND STOLYAR, A. (2002). Scheduling for multiple flows sharing a time-varying channel: the exponential rule. In *Analytic Methods in Applied Probability*, ed. Yu. M. Suhov, American Mathematical Society, Providence, RI, pp. 185–202.
- [9] STOLYAR, A. (2004). MaxWeight scheduling in a generalized switch: state space collapse and workload minimization in heavy traffic. *Ann. Appl. Prob.* **14**, 1–53.
- [10] STOLYAR, A. L. (1995). On the stability of multiclass queueing networks: a relaxed sufficient condition via limiting fluid processes. *Markov Process. Relat. Fields* **1**, 491–512.
- [11] TASSIULAS, L. AND EPHREMIDES, A. (1992). Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Trans. Automatic Control* **37**, 1936–1948.